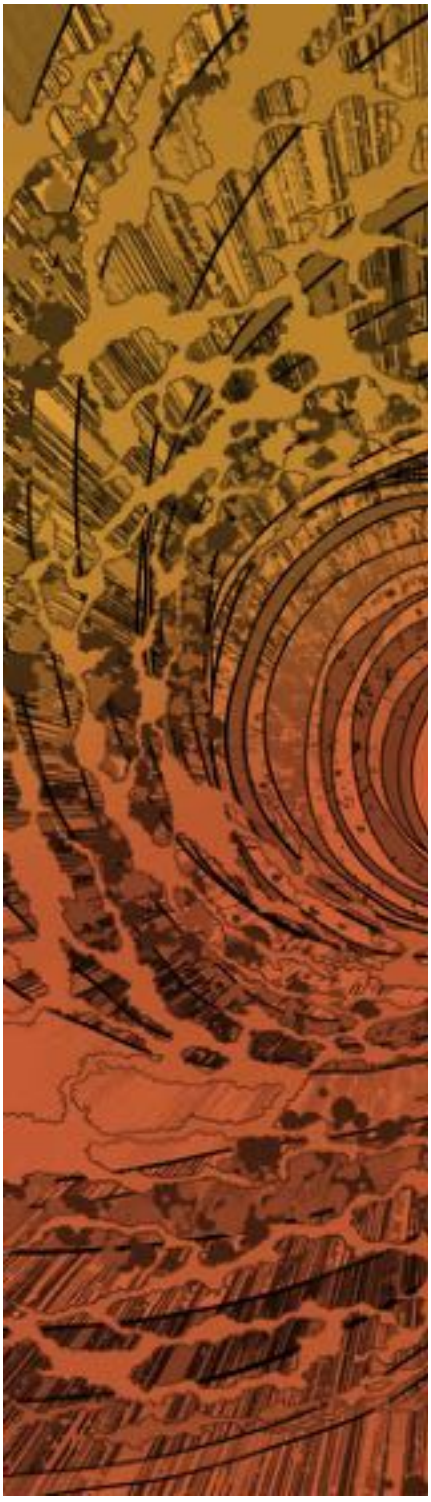


paloalto NETWORKS®

PAN-OS® Administrator's Guide

Version 7.1

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 how to configure other components in the Palo Alto Networks Next-Generation Security Platform, go to the Technical Documentation portal: https://www.paloaltonetworks.com/documentation or search the 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 7.1 release notes, go to https://www.paloaltonetworks.com/documentation/71/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: October 21, 2016

Table of Contents

Getting Started 17

Integrate the Firewall into Your Management Network. 18
- Determine Your Management Strategy 18
- Perform Initial Configuration 19
- Set Up Network Access for External Services. 23

Register the Firewall 28

Activate Licenses and Subscriptions 29

Install Content and Software Updates. 31

Segment Your Network Using Interfaces and Zones 35
- Network Segmentation for a Reduced Attack Surface 35
- Configure Interfaces and Zones. 36

Set Up a Basic Security Policy. 39

Assess Network Traffic 43

Enable Basic Threat Prevention Features 45
- Enable Basic WildFire Forwarding 45
- Scan Traffic for Threats. 47
- Control Access to Web Content 51
- Enable AutoFocus Threat Intelligence 54

Best Practices for Completing the Firewall Deployment 56

Firewall Administration 57

Management Interfaces 58

Use the Web Interface 59
- Launch the Web Interface 59
- Configure Banners, Message of the Day, and Logos 60
- Use the Administrator Login Activity Indicators to Detect Account Misuse 62
- Manage and Monitor Administrative Tasks 64
- Commit, Validate, and Preview Firewall Configuration Changes. 64
- Use Global Find to Search the Firewall or Panorama Management Server 66
- Manage Locks for Restricting Configuration Changes. 67

Manage Configuration Backups 69
- Back Up a Configuration. 69
- Restore a Configuration 70

Manage Firewall Administrators. 72
- Administrative Roles 72
- Administrative Authentication. 73
- Configure Administrative Accounts and Authentication. 74
- Configure an Administrative Account. 74
- Configure Kerberos SSO and External or Local Authentication for Administrators. 75
- Configure Certificate-Based Administrator Authentication to the Web Interface. 76
- Configure SSH Key-Based Administrator Authentication to the CLI 78
- Configure RADIUS Vendor-Specific Attributes for Administrator Authentication. 78

© Palo Alto Networks, Inc.

Reference: Web Interface Administrator Access 80
Web Interface Access Privileges 80
Panorama Web Interface Access Privileges 122
Reference: Port Number Usage 125
Ports Used for Management Functions 125
Ports Used for HA 126
Ports Used for Panorama 126
Ports Used for GlobalProtect 127
Ports Used for User-ID 128
Reset the Firewall to Factory Default Settings 130
Bootstrap the Firewall 131
USB Flash Drive Support 131
Sample init-cfg.txt Files 132
Prepare a USB Flash Drive for Bootstrapping a Firewall 133
Bootstrap a Firewall Using a USB Flash Drive 136

Authentication 139

Configure an Authentication Profile and Sequence 140
Configure Kerberos Single Sign-On 143
Configure Local Database Authentication 144
Configure External Authentication 145
Configure Authentication Server Profiles 145
Configure a RADIUS Server Profile 145
RADIUS Vendor-Specific Attributes Support 146
Configure a TACACS+ Server Profile 147
Configure an LDAP Server Profile 148
Configure a Kerberos Server Profile 150
Set CHAP or PAP Authentication for RADIUS and TACACS+ Servers 150
Enable External Authentication for Users and Services 151
Test Authentication Server Connectivity 152
Run the Test Authentication Command 152
Test a Local Database Authentication Profile 153
Test a RADIUS Authentication Profile 154
Test a TACACS+ Authentication Profile 156
Test an LDAP Authentication Profile 157
Test a Kerberos Authentication Profile 158
Troubleshoot Authentication Issues 160

Certificate Management 161

Keys and Certificates 162
Certificate Revocation 164
Certificate Revocation List (CRL) 164
Online Certificate Status Protocol (OCSP) 165
Certificate Deployment 166
Set Up Verification for Certificate Revocation Status 167
Configure an OCSP Responder 167
Configure Revocation Status Verification of Certificates 168
Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption . 168

 © Palo Alto Networks, Inc.

Configure the Master Key . . . 170
Obtain Certificates . . . 171
Create a Self-Signed Root CA Certificate . . . 171
Generate a Certificate . . . 172
Import a Certificate and Private Key . . . 173
Obtain a Certificate from an External CA . . . 174
Export a Certificate and Private Key . . . 176
Configure a Certificate Profile . . . 177
Configure an SSL/TLS Service Profile . . . 179
Replace the Certificate for Inbound Management Traffic . . . 180
Configure the Key Size for SSL Forward Proxy Server Certificates . . . 181
Revoke and Renew Certificates . . . 182
Revoke a Certificate . . . 182
Renew a Certificate . . . 182
Secure Keys with a Hardware Security Module . . . 183
Set up Connectivity with an HSM . . . 183
Encrypt a Master Key Using an HSM . . . 188
Store Private Keys on an HSM . . . 189
Manage the HSM Deployment . . . 190

High Availability . . . 191

HA Overview . . . 192
HA Concepts . . . 193
HA Modes . . . 193
HA Links and Backup Links . . . 194
Device Priority and Preemption . . . 197
Failover . . . 197
LACP and LLDP Pre-Negotiation for Active/Passive HA . . . 198
Floating IP Address and Virtual MAC Address . . . 198
ARP Load-Sharing . . . 200
Route-Based Redundancy . . . 202
HA Timers . . . 202
Session Owner . . . 205
Session Setup . . . 205
NAT in Active/Active HA Mode . . . 207
ECMP in Active/Active HA Mode . . . 208
Set Up Active/Passive HA . . . 209
Prerequisites for Active/Passive HA . . . 209
Configuration Guidelines for Active/Passive HA . . . 210
Configure Active/Passive HA . . . 212
Define HA Failover Conditions . . . 217
Verify Failover . . . 218
Set Up Active/Active HA . . . 219
Prerequisites for Active/Active HA . . . 219
Configure Active/Active HA . . . 220
Determine Your Active/Active Use Case . . . 225
Use Case: Configure A/A HA with Route-Based Redundancy . . . 226
Use Case: Configure A/A HA with Floating IP Addresses . . . 227

Use Case: Configure A/A HA with ARP Load-Sharing 228
Use Case: Configure A/A HA with Floating IP Address Bound to A-P Firewall. 229
Use Case: Configure A/A HA with Source DIPP NAT Using Floating IP Addresses 233
Use Case: Configure Separate Source NAT IP Address Pools for A/A HA Firewalls 236
Use Case: Configure A/A HA for ARP Load-Sharing with Destination NAT 237
Use Case: Configure A/A HA for ARP Load-Sharing with Destination NAT in Layer 3 . . 240
HA Firewall States 243
Reference: HA Synchronization. 245
What Settings Don't Sync in Active/Passive HA?. 245
What Settings Don't Sync in Active/Active HA?. 247
Synchronization of System Runtime Information 249

Monitoring 251

Use the Dashboard 252
Use the Application Command Center. 253
ACC—First Look 254
ACC Tabs 255
ACC Widgets. 256
Widget Descriptions. 257
ACC Filters. 261
Interact with the ACC 262
Use Case: ACC—Path of Information Discovery 265
App Scope 272
Summary Report 273
Change Monitor Report. 274
Threat Monitor Report. 275
Threat Map Report. 276
Network Monitor Report. 277
Traffic Map Report 278
Use the Automated Correlation Engine 279
Automated Correlation Engine Concepts. 279
View the Correlated Objects 280
Interpret Correlated Events 281
Use the Compromised Hosts Widget in the ACC 283
Take Packet Captures 284
Types of Packet Captures 284
Disable Hardware Offload. 285
Take a Custom Packet Capture 286
Take a Threat Packet Capture 290
Take an Application Packet Capture 291
Take a Packet Capture on the Management Interface. 294
Monitor Applications and Threats. 296
Monitor and Manage Logs 297
Log Types and Severity Levels 297
Work with Logs. 301
Configure Log Storage Quotas and Expiration Periods 307
Schedule Log Exports to an SCP or FTP Server. 307

 © Palo Alto Networks, Inc.

Manage Reporting . . . 309
Report Types . . . 309
View Reports . . . 310
Configure the Report Expiration Period . . . 310
Disable Predefined Reports . . . 311
Custom Reports . . . 311
Generate Custom Reports . . . 314
Generate Botnet Reports . . . 316
Generate the SaaS Application Usage Report . . . 318
Manage PDF Summary Reports . . . 320
Generate User/Group Activity Reports . . . 321
Manage Report Groups . . . 323
Schedule Reports for Email Delivery . . . 324
Use External Services for Monitoring . . . 325
Configure Log Forwarding . . . 326
Configure Email Alerts . . . 329
Use Syslog for Monitoring . . . 330
Configure Syslog Monitoring . . . 330
Syslog Field Descriptions . . . 332
SNMP Monitoring and Traps . . . 348
SNMP Support . . . 348
Use an SNMP Manager to Explore MIBs and Objects . . . 349
Enable SNMP Services for Firewall-Secured Network Elements . . . 353
Monitor Statistics Using SNMP . . . 353
Forward Traps to an SNMP Manager . . . 355
Supported MIBs . . . 357
NetFlow Monitoring . . . 364
Configure NetFlow Exports . . . 364
NetFlow Templates . . . 365
Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors . . . 369

User-ID . . . 371

User-ID Overview . . . 372
User-ID Concepts . . . 374
Group Mapping . . . 374
User Mapping . . . 374
Enable User-ID . . . 378
Map Users to Groups . . . 382
Map IP Addresses to Users . . . 385
Create a Dedicated Service Account for the User-ID Agent . . . 386
Configure User Mapping Using the Windows User-ID Agent . . . 389
Configure User Mapping Using the PAN-OS Integrated User-ID Agent . . . 396
Configure User-ID to Receive User Mappings from a Syslog Sender . . . 398
Map IP Addresses to Usernames Using Captive Portal . . . 406
Configure User Mapping for Terminal Server Users . . . 412
Send User Mappings to User-ID Using the XML API . . . 419
Enable User- and Group-Based Policy . . . 420

Enable Policy for Users with Multiple Accounts 421
Verify the User-ID Configuration 423
Deploy User-ID in a Large-Scale Network 425
Deploy User-ID for Numerous Mapping Information Sources 425
Configure Firewalls to Redistribute User Mapping Information 429

App-ID 435
App-ID Overview 436
Manage Custom or Unknown Applications 437
Manage New App-IDs Introduced in Content Releases 438
Review New App-IDs 438
Review New App-IDs Since Last Content Version 439
Review New App-ID Impact on Existing Policy Rules 440
Disable or Enable App-IDs 441
Prepare Policy Updates for Pending App-IDs 442
Use Application Objects in Policy 444
Create an Application Group 444
Create an Application Filter 445
Create a Custom Application 446
Applications with Implicit Support 451
Application Level Gateways 454
Disable the SIP Application-level Gateway (ALG) 455

Threat Prevention 457
Set Up Security Profiles and Policies 458
Set Up Antivirus, Anti-Spyware, and Vulnerability Protection 458
Set Up Data Filtering 461
Set Up File Blocking 464
Prevent Brute Force Attacks 466
Customize the Action and Trigger Conditions for a Brute Force Signature 467
Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions 470
Enable DNS Proxy 475
Enable Passive DNS Collection for Improved Threat Intelligence 478
Use DNS Queries to Identify Infected Hosts on the Network 479
DNS Sinkholing 479
Configure DNS Sinkholing for a List of Custom Domains 481
Configure the Sinkhole IP Address to a Local Server on Your Network 483
Identify Infected Hosts 487
DoS Protection Against Flooding of New Sessions 489
Multiple-Session DoS Attack 489
Single-Session DoS Attack 492
Configure DoS Protection Against Flooding of New Sessions 492
Use the CLI to End a Single Attacking Session 495
Identify Sessions That Use an Excessive Percentage of the Packet Buffer 495
Discard a Session Without a Commit 498
Content Delivery Network Infrastructure for Dynamic Updates 499
Threat Prevention Resources 501

 © Palo Alto Networks, Inc.

Decryption 503
Decryption Overview 504
Decryption Concepts 505
Keys and Certificates for Decryption Policies 505
SSL Forward Proxy 506
SSL Inbound Inspection 507
SSH Proxy 508
Decryption Exceptions 509
Decryption Mirroring 510
Define Traffic to Decrypt 511
Create a Decryption Profile 511
Create a Decryption Policy Rule 513
Configure SSL Forward Proxy 515
Configure SSL Inbound Inspection 519
Configure SSH Proxy 521
Configure Decryption Exceptions 522
Exclude Traffic from Decryption 522
Exclude a Server from Decryption 523
Enable Users to Opt Out of SSL Decryption 524
Configure Decryption Port Mirroring 526
Temporarily Disable SSL Decryption 528

URL Filtering 529
URL Filtering Overview 530
URL Filtering Vendors 530
Interaction Between App-ID and URL Categories 531
PAN-DB Private Cloud 531
URL Filtering Concepts 534
URL Categories 534
URL Filtering Profile 536
URL Filtering Profile Actions 536
Block and Allow Lists 537
External Dynamic List for URLs 538
Safe Search Enforcement 538
Container Pages 540
HTTP Header Logging 540
URL Filtering Response Pages 540
URL Category as Policy Match Criteria 542
PAN-DB Categorization 544
PAN-DB URL Categorization Components 544
PAN-DB URL Categorization Workflow 545
Enable a URL Filtering Vendor 547
Enable PAN-DB URL Filtering 547
Enable BrightCloud URL Filtering 548
Determine URL Filtering Policy Requirements 551
Use an External Dynamic List in a URL Filtering Profile 553

© Palo Alto Networks, Inc.

Monitor Web Activity 555
Monitor Web Activity of Network Users 555
View the User Activity Report 557
Configure Custom URL Filtering Reports 559
Configure URL Filtering 560
Customize the URL Filtering Response Pages 562
Configure URL Admin Override 563
Enable Safe Search Enforcement 565
Block Search Results that are not Using Strict Safe Search Settings 565
Enable Transparent Safe Search Enforcement 568
Set Up the PAN-DB Private Cloud 572
Configure the PAN-DB Private Cloud 572
Configure the Firewalls to Access the PAN-DB Private Cloud 577
URL Filtering Use Case Examples 578
Use Case: Control Web Access 578
Use Case: Use URL Categories for Policy Matching 582
Troubleshoot URL Filtering 584
Problems Activating PAN-DB 584
PAN-DB Cloud Connectivity Issues 585
URLs Classified as Not-Resolved 586
Incorrect Categorization 586
URL Database Out of Date 587

Quality of Service 589

QoS Overview 590
QoS Concepts 592
QoS for Applications and Users 592
QoS Policy 592
QoS Profile 593
QoS Classes 593
QoS Priority Queuing 594
QoS Bandwidth Management 594
QoS Egress Interface 595
QoS for Clear Text and Tunneled Traffic 595
Configure QoS 596
Configure QoS for a Virtual System 601
Enforce QoS Based on DSCP Classification 606
QoS Use Cases 609
Use Case: QoS for a Single User 609
Use Case: QoS for Voice and Video Applications 611

VPNs 615

VPN Deployments 616
Site-to-Site VPN Overview 617
Site-to-Site VPN Concepts 618
IKE Gateway 618
Tunnel Interface 618

 © Palo Alto Networks, Inc.

Tunnel Monitoring . . . 619
Internet Key Exchange (IKE) for VPN . . . 620
IKEv2 . . . 622
Set Up Site-to-Site VPN . . . 626
Set Up an IKE Gateway . . . 626
Define Cryptographic Profiles . . . 632
Set Up an IPSec Tunnel . . . 635
Set Up Tunnel Monitoring . . . 638
Enable/Disable, Refresh or Restart an IKE Gateway or IPSec Tunnel . . . 639
Test VPN Connectivity . . . 641
Interpret VPN Error Messages . . . 642
Site-to-Site VPN Quick Configs . . . 643
Site-to-Site VPN with Static Routing . . . 643
Site-to-Site VPN with OSPF . . . 647
Site-to-Site VPN with Static and Dynamic Routing . . . 653

Large Scale VPN (LSVPN) . . . 659
LSVPN Overview . . . 660
Create Interfaces and Zones for the LSVPN . . . 661
Enable SSL Between GlobalProtect LSVPN Components . . . 663
About Certificate Deployment . . . 663
Deploy Server Certificates to the GlobalProtect LSVPN Components . . . 663
Deploy Client Certificates to the GlobalProtect Satellites Using SCEP . . . 666
Configure the Portal to Authenticate Satellites . . . 669
Configure GlobalProtect Gateways for LSVPN . . . 671
Prerequisite Tasks . . . 671
Configure the Gateway . . . 671
Configure the GlobalProtect Portal for LSVPN . . . 674
Prerequisite Tasks . . . 674
Configure the Portal . . . 674
Define the Satellite Configurations . . . 675
Prepare the Satellite to Join the LSVPN . . . 679
Verify the LSVPN Configuration . . . 681
LSVPN Quick Configs . . . 682
Basic LSVPN Configuration with Static Routing . . . 683
Advanced LSVPN Configuration with Dynamic Routing . . . 686
Advanced LSVPN Configuration with iBGP . . . 689

Networking . . . 695
Interface Deployments . . . 696
Virtual Wire Deployments . . . 696
Layer 2 Deployments . . . 699
Layer 3 Deployments . . . 699
Tap Mode Deployments . . . 700
Configure an Aggregate Interface Group . . . 702
Use Interface Management Profiles to Restrict Access . . . 705

© Palo Alto Networks, Inc.

Virtual Routers 707
Static Routes 709
RIP 711
OSPF 713
OSPF Concepts 713
Configure OSPF 715
Configure OSPFv3 719
Configure OSPF Graceful Restart 722
Confirm OSPF Operation 723
BGP 725
Session Settings and Timeouts 730
Transport Layer Sessions 730
TCP 730
UDP 735
ICMP 735
Configure Session Timeouts 736
Configure Session Settings 738
Prevent TCP Split Handshake Session Establishment 740
DHCP 742
DHCP Overview 742
Firewall as a DHCP Server and Client 743
DHCP Messages 743
DHCP Addressing 744
DHCP Options 746
Configure an Interface as a DHCP Server 748
Configure an Interface as a DHCP Client 752
Configure the Management Interface as a DHCP Client 753
Configure an Interface as a DHCP Relay Agent 755
Monitor and Troubleshoot DHCP 755
DNS 757
DNS Overview 757
DNS Proxy Object 758
DNS Server Profile 759
Multi-Tenant DNS Deployments 759
Configure a DNS Proxy Object 760
Configure a DNS Server Profile 762
Use Case 1: Firewall Requires DNS Resolution for Management Purposes 763
Use Case 2: ISP Tenant Uses DNS Proxy to Handle DNS Resolution for Security Policies, Reporting, and Services within its Virtual System764
Use Case 3: Firewall Acts as DNS Proxy Between Client and Server 766
Reference: DNS Proxy Rule and FQDN Matching 768
NAT 772
NAT Policy Rules 772
Source NAT and Destination NAT 775
NAT Rule Capacities 776
Dynamic IP and Port NAT Oversubscription 776
Dataplane NAT Memory Statistics 778
Configure NAT 779
NAT Configuration Examples 786

 © Palo Alto Networks, Inc.

NPTv6 . . . 794
NPTv6 Overview . . . 794
How NPTv6 Works . . . 796
NDP Proxy . . . 797
NPTv6 and NDP Proxy Example . . . 799
Create an NPTv6 Policy . . . 800
ECMP . . . 803
ECMP Load-Balancing Algorithms . . . 803
ECMP Platform, Interface, and IP Routing Support . . . 804
Configure ECMP on a Virtual Router . . . 805
Enable ECMP for Multiple BGP Autonomous Systems . . . 806
Verify ECMP . . . 808
LLDP . . . 809
LLDP Overview . . . 809
Supported TLVs in LLDP . . . 810
LLDP Syslog Messages and SNMP Traps . . . 811
Configure LLDP . . . 812
View LLDP Settings and Status . . . 814
Clear LLDP Statistics . . . 815
BFD . . . 816
BFD Overview . . . 816
Configure BFD . . . 819
Reference: BFD Details . . . 826

Policy . . . 829

Policy Types . . . 830
Security Policy . . . 831
Components of a Security Policy Rule . . . 831
Security Policy Actions . . . 834
Create a Security Policy Rule . . . 834
Policy Objects . . . 837
Security Profiles . . . 838
Antivirus Profiles . . . 839
Anti-Spyware Profiles . . . 839
Vulnerability Protection Profiles . . . 840
URL Filtering Profiles . . . 840
Data Filtering Profiles . . . 841
File Blocking Profiles . . . 842
WildFire Analysis Profiles . . . 842
DoS Protection Profiles . . . 842
Zone Protection Profiles . . . 843
Security Profile Group . . . 843
Best Practice Internet Gateway Security Policy . . . 847
What Is a Best Practice Internet Gateway Security Policy? . . . 847
Why Do I Need a Best Practice Internet Gateway Security Policy? . . . 849
How Do I Deploy a Best Practice Internet Gateway Security Policy? . . . 850
Identify Whitelist Applications . . . 851
Create User Groups for Access to Whitelist Applications . . . 854
Decrypt Traffic for Full Visibility and Threat Inspection . . . 854

Create Best Practice Security Profiles . . . 856
Define the Initial Internet Gateway Security Policy . . . 860
Monitor and Fine Tune the Policy Rulebase . . . 868
Remove the Temporary Rules . . . 869
Maintain the Rulebase . . . 870
Enumeration of Rules Within a Rulebase . . . 871
Move or Clone a Policy Rule or Object to a Different Virtual System . . . 872
Use Tags to Group and Visually Distinguish Objects . . . 873
Create and Apply Tags . . . 873
Modify Tags . . . 874
Use the Tag Browser . . . 874
Use an External Dynamic List in Policy . . . 879
External Dynamic List . . . 879
Formatting Guidelines for an External Dynamic List . . . 880
Enforce Policy on Entries in an External Dynamic List . . . 881
View the List of Entries in an External Dynamic List . . . 884
Retrieve an External Dynamic List from the Web Server . . . 885
Register IP Addresses and Tags Dynamically . . . 886
Monitor Changes in the Virtual Environment . . . 887
Enable VM Monitoring to Track Changes on the Virtual Network . . . 887
Attributes Monitored in the AWS and VMware Environments . . . 889
Use Dynamic Address Groups in Policy . . . 890
CLI Commands for Dynamic IP Addresses and Tags . . . 893
Identify Users Connected through a Proxy Server . . . 895
Use XFF Values for Policies and Logging Source Users . . . 895
Add XFF Values to URL Filtering Logs . . . 896
Policy-Based Forwarding . . . 897
PBF . . . 897
Create a Policy-Based Forwarding Rule . . . 900
Use Case: PBF for Outbound Access with Dual ISPs . . . 902

Virtual Systems . . . 909

Virtual Systems Overview . . . 910
Virtual System Components and Segmentation . . . 910
Benefits of Virtual Systems . . . 911
Use Cases for Virtual Systems . . . 911
Platform Support and Licensing for Virtual Systems . . . 912
Administrative Roles for Virtual Systems . . . 912
Shared Objects for Virtual Systems . . . 912
Communication Between Virtual Systems . . . 913
Inter-VSYS Traffic That Must Leave the Firewall . . . 913
Inter-VSYS Traffic That Remains Within the Firewall . . . 914
Inter-VSYS Communication Uses Two Sessions . . . 916
Shared Gateway . . . 917
External Zones and Shared Gateway . . . 917
Networking Considerations for a Shared Gateway . . . 918

 © Palo Alto Networks, Inc.

Service Routes for Virtual Systems 919
Use Cases for Service Routes for a Virtual System 919
PA-7000 Series Firewall LPC Support for Per-Virtual System Paths to Logging Servers 920
Configure Virtual Systems 921
Configure Inter-Virtual System Communication within the Firewall 924
Configure a Shared Gateway 925
Customize Service Routes for a Virtual System 926
Customize Service Routes to Services for Virtual Systems 926
Configure a PA-7000 Series Firewall for Logging Per Virtual System 928
Configure Administrative Access Per Virtual System or Firewall 929
Virtual System Functionality with Other Features 931

Certifications 933
Enable FIPS and Common Criteria Support 934
FIPS-CC Security Functions 935

© 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, registering the firewall, activating licenses and subscriptions, 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 security platform features as necessary to address your network security needs.

- Integrate the Firewall into Your Management Network
- Register the Firewall
- Activate Licenses and Subscriptions
- Install Content and Software Updates
- Segment Your Network Using Interfaces and Zones
- Set Up a Basic Security Policy
- Assess Network Traffic
- 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 data port for managing your firewall 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 an in-band data port to provide access to required external services (using service routes) 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



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 firewall-specific configuration locally on a firewall and enforce shared policies across all firewalls 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.

 © Palo Alto Networks, Inc.

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 firewall.

Set Up Network Access to the Firewall		
Step 1	Gather the required information from your network administrator.	• IP address for MGT port • Netmask • Default gateway • DNS server address
Step 2	Connect your computer to the firewall.	You can connect to the firewall in one of the following ways: • 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 firewall is ready, the prompt changes to the name of the firewall, for example `PA-500 login`. • 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 192.168.1.0/24 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 edit the Management Interface Settings. 2. Configure the address settings for the MGT interface using one of the following methods: • To configure static IP address settings for the MGT interface, set the **IP Type** to **Static** and enter the **IP Address**, **Netmask**, and **Default Gateway**. • To dynamically configure the MGT interface address settings, set the **IP Type** to **DHCP**. To use this method, you must Configure the Management Interface as a DHCP Client. 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. 3. Set the **Speed** to **auto-negotiate**. 4. Select which management services to allow on the interface. 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. 5. Click **OK**.

Set Up Network Access to the Firewall (Continued)		
Step 5	Configure DNS, update server, and proxy server settings.  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**. • For multi-virtual system platforms, select **Global** and edit the Services section. • For single virtual system platforms, edit the Services section. 2. On the **Services** tab, for **DNS**, click one of the following: • **Servers**—Enter the **Primary DNS Server** address and **Secondary DNS Server** address. • **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. 3. Click **OK**.
Step 6	Configure date and time (NTP) settings.	1. Select **Device > Setup > Services**. • For multi-virtual system platforms, select **Global** and edit the Services section. • For single virtual system platforms, edit the Services section. 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. 3. (Optional) Enter a **Secondary NTP Server** address. 4. (Optional) To authenticate time updates from the NTP server(s), for **Authentication Type**, select one of the following for each server: • **None**—(Default) Disables NTP authentication. • **Symmetric Key**—Firewall uses symmetric key exchange (shared secrets) to authenticate time updates. – **Key ID**—Enter the Key ID (1-65534). – **Algorithm**—Select the algorithm to use in NTP authentication (**MD5** or **SHA1**). • **Autokey**—Firewall uses autokey (public key cryptography) to authenticate time updates. 5. Click **OK**.

 © Palo Alto Networks, Inc.

Set Up Network Access to the Firewall (Continued)		
Step 7	(Optional) Configure general firewall settings as needed.	1. Select **Device > Setup > Management** and edit the General Settings. 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. 3. Enter **Login Banner** text that informs users who are about to log in that they require authorization to access the firewall management functions.  As a best practice, avoid using welcoming verbiage. Additionally, you should ask your legal department to review the banner message to ensure it adequately warns that unauthorized access is prohibited. 4. Enter the **Latitude** and **Longitude** to enable accurate placement of the firewall on the world map. 5. Click **OK**.
Step 8	Set a secure password for the admin account.	1. Select **Device > Administrators**. 2. Select the **admin** role. 3. Enter the current default password and the new password. 4. Click **OK** to save your settings.
Step 9	Commit your changes.  When the configuration changes are saved, you lose connectivity to the web interface because the IP address has changed.	Click **Commit** at the top right of the web interface. The firewall can take up to 90 seconds to save your changes.
Step 10	Connect the firewall to your network.	1. Disconnect the firewall from your computer. 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.

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.

You can do this in one of the following ways:

- 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.
- 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.

1. Use the ping utility to verify network connectivity to the Palo Alto Networks Update server as shown in the following example. Verify that DNS resolution occurs and the response includes the IP address for the Update server; the update server does not respond to a ping request.

```
admin@PA-200 > ping host
updates.paloaltonetworks.com
PING updates.paloaltonetworks.com (10.101.16.13)
56(84) bytes of data.
From 192.168.1.1 icmp_seq=1 Destination Host
Unreachable
From 192.168.1.1 icmp_seq=2 Destination Host
Unreachable
From 192.168.1.1 icmp_seq=3 Destination Host
Unreachable
From 192.168.1.1 icmp_seq=4 Destination Host
Unreachable
```



After verifying DNS resolution, press Ctrl+C to stop the ping request.

2. Use the following CLI command to retrieve information on the support entitlement for the firewall from the Palo Alto Networks update server:

```
request support check
```

If you have connectivity, the update server will respond with the support status for your firewall. Because your firewall is not registered, the update server will return the following message:

```
Contact Us
https://www.paloaltonetworks.com/company/contact-u
s.html
Support Home
https://www.paloaltonetworks.com/support/tabs/over
view.html
Device not found on this update server
```

 © Palo Alto Networks, Inc.

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 an in-band data port to provide access to required external services and set up service routes to instruct the firewall what port to use to access the external services.



This task requires familiarity with firewall interfaces, zones, and policies. For more information on these topics, see Configure Interfaces and Zones and Set Up a Basic Security Policy.

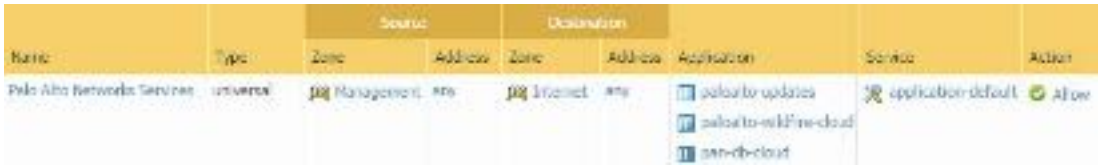
Set Up a Data Port for Access to External Services		
Step 1	Decide which port you want to use for access to external services and connect it to your switch or router port.	The interface you use must have a static IP address.
Step 2	Log in to the web interface.	Using a secure connection (https) from your 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: 1. To delete the default security policy, select **Policies > Security**, select the rule, and click **Delete**. 2. To delete the default virtual wire, select **Network > Virtual Wires**, select the virtual wire and click **Delete**. 3. To delete the default trust and untrust zones, select **Network > Zones**, select each zone and click **Delete**. 4. To delete the interface configurations, select **Network > Interfaces** and then select each interface (ethernet1/1 and ethernet1/2) and click **Delete**. 5. **Commit** the changes.

Set Up a Data Port for Access to External Services (Continued)	
Step 4 Configure the interface you plan to use for external access to management services.	1. Select **Network > Interfaces** and select the interface that corresponds to the port you cabled in Step 1. 2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**. 3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. 4. In the Zone dialog, enter a **Name** for new zone, for example Management, and then click **OK**. 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. You must use a static IP address on this interface. 6. Select **Advanced > Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**. 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**.  These services provide management access to the firewall, 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 firewall 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**). For details, see Use Interface Management Profiles to Restrict Access. 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	Configure the service routes. By default, the firewall uses the MGT interface to access the external services it requires. To change the interface the firewall uses to send requests to external services, you must edit the service routes. This example shows how to set up global service routes. 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.	1. Select **Device > Setup > Services > Global** and click **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**, **WildFire**, and **AutoFocus**. 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.
Step 6	Configure an external-facing interface and an associated zone and then create a security policy rule to allow the firewall to send service requests from the internal zone to the external zone.	1. Select **Network > Interfaces** and then select the 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 Internet. This interface must have a static IP address; 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, select **Policies > Security** and click **Add**. As a best practice when creating Security policy rules, use application-based rules instead of port-based rules to ensure that you are accurately identifying the underlying application regardless of the port, protocol, evasive tactics, or encryption in use. Always leave the **Service** set to **application-default**. In this case, create a security policy rule that allows access to the update server (and other Palo Alto Networks services).

Set Up a Data Port for Access to External Services (Continued)



Step 7 Create a NAT policy rule.	1. 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, Management to Internet in this case (**Original Packet** tab), and define the source address translation settings (**Translated Packet** tab) and then click **OK**. 2. **Commit** your changes.



 © Palo Alto Networks, Inc.

Set Up a Data Port for Access to External Services (Continued)	

Step 8 Verify that you have connectivity from the data port to the external services, including the default gateway, and the Palo Alto Networks Update Server.
After you verify you have the required network connectivity, continue to Register the Firewall and Activate Licenses and Subscriptions.
	1. Use the ping utility to verify network connectivity to the Palo Alto Networks Update server as shown in the following example. Verify that DNS resolution occurs and the response includes the IP address for the Update server; the update server does not respond to a ping request.
	`admin@PA-200 > ping host updates.paloaltonetworks.com`
	`PING updates.paloaltonetworks.com (10.101.16.13) 56(84) bytes of data.`
	`From 192.168.1.1 icmp_seq=1 Destination Host Unreachable`
	`From 192.168.1.1 icmp_seq=2 Destination Host Unreachable`
	`From 192.168.1.1 icmp_seq=3 Destination Host Unreachable`
	`From 192.168.1.1 icmp_seq=4 Destination Host Unreachable`
	After verifying DNS resolution, press Ctrl+C to stop the ping request.
	2. Use the following CLI command to retrieve information on the support entitlement for the firewall from the Palo Alto Networks update server:
	`request support check`
	If you have connectivity, the update server will respond with the support status for your firewall. Because your firewall is not registered, the update server will return the following message:
	`Contact Us https://www.paloaltonetworks.com/company/contact-us.html`
	`Support Home https://www.paloaltonetworks.com/support/tabs/overview.html`
	`Device not found on this update server`

Register the Firewall

Before you can activate support and other licenses and subscriptions, you must first register the firewall.



If you are registering a VM-Series firewall, refer to the VM-Series Deployment Guide.

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>).
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 Customer Support portal and log in.	In a new browser tab or window, go to https://www.paloaltonetworks.com/support/tabs/overview.html.
Step 4	Register the firewall.  You must have a support account to register a firewall. If you do not yet have a support account, click the **Register** link on the support login page and follow the instructions to get your account set up and register the firewall.	If you already have a support account, log in and register the hardware-based firewall as follows: 1. Select **Assets > Devices**. 2. Click **Register New Device**. 3. Select **Register device using Serial Number or Authorization Code** and click **Submit**. 4. Enter the firewall **Serial Number** (you can copy and paste it from the firewall Dashboard). 5. (Optional) Enter the **Device Name** and **Device Tag**. 6. Provide information about where you plan to deploy the firewall including the **City**, **Postal Code**, and **Country**. 7. Read the end-user license agreement (EULA) and then click **Agree and Submit**.

© 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.
- AutoFocus—Provides a graphical analysis of firewall traffic logs and identifies potential risks to your network using threat intelligence from the AutoFocus portal. With an active license, you can also open an AutoFocus search based on logs recorded on the firewall.

Activate Licenses and Subscriptions		
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	Activate your Support license. You will not be able to update your PAN-OS software if you do not have a valid Support license.	1. Log in to the web interface and then select **Device > Support**. 2. Click **Activate support using authorization code**. 3. Enter your **Authorization Code** and then click **OK**.

Activate Licenses and Subscriptions (Continued)		
Step 3	Activate each license you purchased.	Select **Device > Licenses** and then activate your licenses and subscriptions in one of the following ways: • **Retrieve license keys from license server**—Use this option if you activated your license on the Customer 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. When prompted, enter the **Authorization Code** and then click **OK**. • **Manually upload license key**—Use this option if your firewall does not have connectivity to the Palo Alto Networks Customer Support web 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 firewall.
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: 
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. You should either: • Commit any pending changes. • 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.

Install Content and Software 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 Palo Alto Networks firewalls. The firewalls access the web resources in the CDN to perform various App-ID and Content-ID functions. By default, the firewalls 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 firewalls up-to-date with the latest content and software 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 content 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. To review the policy impact of new application updates, see Manage New App-IDs Introduced in Content Releases.
- **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 firewalls 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 Threats update.

Install Content and Software Updates	
Step 1 Ensure that the firewall has access to the update server.	1. By default, the firewall accesses the **Update Server** at `updates.paloaltonetworks.com` so that the firewall receives content updates from the server to which it is closest in the CDN infrastructure. 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` instead of dynamically selecting a server from the CDN infrastructure. 2. (Optional) Click **Verify Update Server Identity** for an extra level of validation to enable the firewall to check that the server's SSL certificate is signed by a trusted authority. 3. (Optional) If the firewall needs to use a proxy server to reach Palo Alto Networks update services, in the **Proxy Server** window, enter: • **Server**—IP address or host name of the proxy server. • **Port**—Port for the proxy server. Range: 1-65535. • **User**—Username to access the server. • **Password**—Password for the user to access the proxy server. Re-enter the password at **Confirm Password**.

 © Palo Alto Networks, Inc.

Install Content and Software Updates (Continued)	
Step 2 Check for the latest content updates.	Select **Device > Dynamic Updates** and 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 update until you have installed the Application and Threats update. • **Upgrade**—Indicates that a new version of the BrightCloud database is 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 on the firewall automatically synchronizes with the PAN-DB cloud.  To check the status of an action, click **Tasks** (on the lower right-hand corner of the window). • **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.
Step 3 Install the content updates. Installation can take up to 20 minutes on a PA-200, PA-500, or PA-2000 Series firewall 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.
	

Install Content and Software Updates (Continued)		
Step 4	Schedule each content update. Repeat this step for each update you want to schedule. 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 Minute**, **Every 15 Minutes**, **Every 30 minutes**, or **Every Hour** whereas Applications and Threats updates can be scheduled for **Daily** or **Weekly** update and Antivirus updates can be scheduled for **Hourly**, **Daily**, or **Weekly**). As new WildFire signatures are made available every five minutes, set the firewall to retrieve WildFire updates **Every Minute** to get the latest signatures within a minute of availability. 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 Only** or, as a best practice, **Download And Install** the update. 5. Enter how long after a release to wait before performing a content update in the **Threshold (Hours)** field. 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. 6. Click **OK** to save the schedule settings. 7. Click **Commit** to save the settings to the running configuration.
Step 5	Update PAN-OS. Always update content before updating PAN-OS. Every PAN-OS version has a minimum supported content release version.	1. Review the Release Notes. 2. Update the PAN-OS software.

 © Palo Alto Networks, Inc.

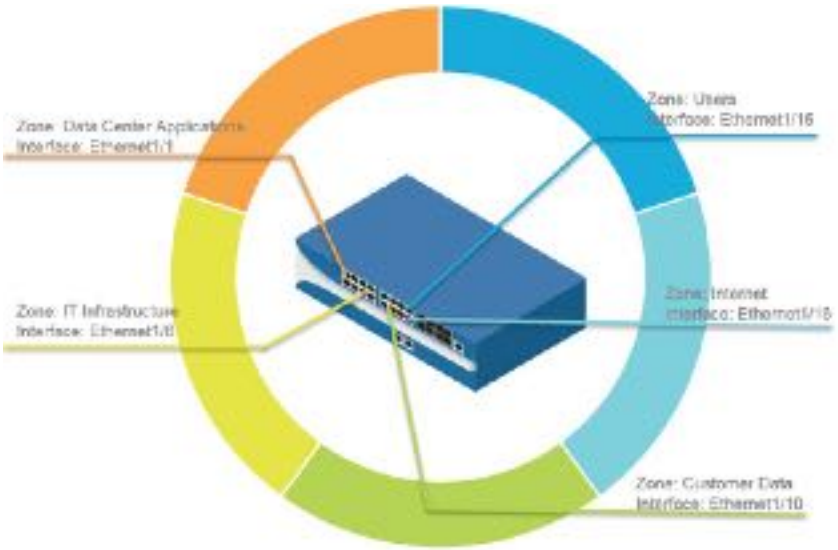
Segment Your Network Using Interfaces and Zones

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 represents a segment of your network that is connected to, and controlled by, the firewall. Because traffic can only flow between zones if there is a Security policy rule to allow it, this is your first line of defense. The more granular the zones you create, the greater control you have over access to sensitive applications and data and the more protection you have against malware moving laterally throughout your network. For example, you might want to segment access to the database servers that store your customer data into a zone called Customer Data. You can then define security policies that only permit certain users or groups of users to access the Customer Data zone, thereby preventing unauthorized internal or external access to the data stored in that segment.

- Network Segmentation for a Reduced Attack Surface
- Configure Interfaces and Zones

Network Segmentation for a Reduced Attack Surface

The following diagram shows a very basic example of how you can create zones to segment your network. The more granular you make your zones (and the corresponding security policy rules that allows traffic between zones), the more you reduce the attack surface on your network. This is because traffic can flow freely within a zone (intra-zone traffic), but traffic cannot flow between zones (inter-zone traffic) until you define a Security policy rule that allows it. Additionally, an interface cannot process traffic until you have assigned it to a zone. Therefore, by segmenting your network into granular zones you have more control over access to sensitive applications or data and you can prevent malicious traffic from establishing a communication channel within your network, thereby reducing the likelihood of a successful attack on your network.



Configure Interfaces and Zones

After you identify how you want to segment your network and the zones you will need to create to achieve the segmentation (as well as the interfaces to map to each zone), you can begin configuring the interfaces and zones on the firewall. Each interface on the firewall supports all Interface Deployments and the deployment you will use depends on the topology of each part of the network you are connecting to. The following workflow shows how to configure Layer 3 interfaces and assign them to zones. For details on integrating the firewall using a different type of interface deployments (for example Virtual Wire Deployments or Layer 2 Deployments), see Networking.



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. 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). 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). 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 Ethernet1/16 as the external interface. 2. Select the **Interface Type**. Although your choice here depends on interface topology, this example shows the steps for **Layer3**. 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 Internet, and then click **OK**. 4. In the **Virtual Router** drop-down, select **default**. 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 203.0.113.23/24. 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**. 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. 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 Ethernet1/15 as the internal interface our users connect to. 2. Select **Layer3** as the **Interface Type**. 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 Users, and then click **OK**. 4. Select the same Virtual Router you used previously, default in this example. 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. 6. To enable you to ping the interface, select the management profile that you just created. 7. To save the interface configuration, click **OK**.
Step 4	Configure the interface that connects to your data center applications. Although this basic security policy example configuration depicts using a single zone for all of your data center applications, as a best practice you would want to define more granular zones to prevent unauthorized access to sensitive applications or data and eliminate the possibility of malware moving laterally within your data center.	1. Select the interface you want to configure. 2. Select **Layer3** from the **Interface Type** drop-down. In this example, we are configuring Ethernet1/1 as the interface that provides access to your data center applications. 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 Data Center Applications, and then click **OK**. 4. Select the same Virtual Router you used previously, default in this example. 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. 6. To enable you to ping the interface, select the management profile that you created. 7. To save the interface configuration, click **OK**.
Step 5	(Optional) Create tags for each zone.	Tags allow you to visually scan policy rules. 1. Select **Objects > Tags** and **Add**. 2. Select a zone **Name**. 3. Select a tag **Color** and click **OK**. 
Step 6	Save the interface configuration.	Click **Commit**.

Set Up Interfaces and Zones (Continued)		
Step 7	Cable the firewall.	Attach straight through cables from the interfaces you configured to the corresponding switch or router on each network segment.
Step 8	Verify that the interfaces are active.	Select **Dashboard** and verify that the interfaces you configured show as green in the Interfaces widget. 

 © Palo Alto Networks, Inc.

Set Up a Basic Security Policy

Now that you have defined some zones and attached them to interfaces, you are ready to begin creating your Security Policy. The firewall will not allow any traffic to flow from one zone to another unless there is a Security policy rule to allow it. When a packet enters a firewall interface, the firewall matches the attributes in the packet against the Security policy rules to determine whether to block or allow the session based on attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service. The firewall evaluates incoming traffic against the security policy rulebase from left to right and from top to bottom and then takes the action specified in the first security rule that matches (for example, whether to allow, deny, or drop the packet). This means that you must order the rules in your security policy rulebase so that more specific rules are at the top of the rulebase and more general rules are at the bottom to ensure that the firewall is enforcing policy as expected.

The following workflow shows how to set up a very basic Internet gateway security policy that enables access to the network infrastructure, to data center applications, and to the Internet. This will enable you to get the firewall up and running so that you can verify that you have successfully configured the firewall. This policy is not comprehensive enough to protect your network. After you verify that you have successfully configured the firewall and integrated it into your network, proceed to Policy to learn how to create a Best Practice Internet Gateway Security Policy that will safely enable application access while protecting your network from attack.

Define Basic Security Policy Rules		
Step 1	(Optional) Delete the default security policy rule.	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 conventions.
Step 2	Create the File Blocking profiles you will need to prevent upload/download of malicious files and for drive-by download protection.	1. Configure a File Blocking profile for general use. You will attach this profile to most of your security profiles to block files known to carry threats or that have no real business use for upload/download. 2. Configure a File Blocking profile for risky traffic. You will attach this profile to security policy rules that allow general web access to prevent users from unknowingly downloading malicious files from the Internet.

Define Basic Security Policy Rules (Continued)	
Step 3 Allow access to your network infrastructure resources.	1. Select **Policies > Security** and click **Add**. 2. Enter a descriptive **Name** for the rule in the **General** tab. 3. In the **Source** tab, set the **Source Zone** to Users. 4. In the **Destination** tab, set the **Destination Zone** to IT Infrastructure. As a best practice, consider using address objects in the **Destination Address** field to enable access to specific servers or groups of servers only, particularly for services such as DNS and SMTP that are commonly exploited. By restricting users to specific destination server addresses you can prevent data exfiltration and command and control traffic from establishing communication through techniques such as DNS tunneling. 5. In the **Applications** tab, **Add** the applications that correspond to the network services you want to safely enable. For example, select **dns**, **ntp**, **ocsp**, **ping**, **smtp**. 6. In the **Service/URL Category** tab, keep the **Service** set to **application-default**. 7. In the **Actions** tab, set the **Action Setting** to **Allow**. 8. Select **Profiles** as the **Profile Type**. Select the **default** profiles for **Antivirus** and **URL Filtering** and the **strict** profiles for **Vulnerability Protection** and **Anti-Spyware** and select the **File Blocking** profile you configured for general traffic. 9. Verify that **Log at Session End** is enabled. Only traffic that matches a security rule will be logged. 10. Click **OK**.



 © Palo Alto Networks, Inc.

Define Basic Security Policy Rules (Continued)

Step 4 Enable access to general Internet applications.

This is a temporary rule that allows you to gather information about the traffic on your network. After you have more insight into what applications your users need access to, you can make informed decisions about what applications to allow and create more granular application-based rules for each user group.

1. Select **Policies > Security** and click **Add**.
2. Enter a descriptive **Name** for the rule in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Users.
4. In the **Destination** tab, set the **Destination Zone** to Internet.
5. In the **Applications** tab, **Add** an **Application Filter** and enter a Name. To safely enable access to legitimate web-based applications, set the Category in the application filter to **general-internet** and then click **OK**. To enable access to encrypted sites, **Add** the **ssl** application.
6. In the **Service/URL Category** tab, keep the **Service** set to **application-default**.
7. In the **Actions** tab, set the **Action Setting** to **Allow**.
8. Select **Profiles** as the **Profile Type**. Select the **default** profiles for **Antivirus** and **URL Filtering** and the **strict** profiles for **Vulnerability Protection** and **Anti-Spyware** and select the **File Blocking** strict profile you configured for risky traffic.
9. Verify that **Log at Session End** is enabled. Only traffic that matches a security rule will be logged.
10. Click **OK**.



Step 5 Enable access to data center applications.

1. Select **Policies > Security** and click **Add**.
2. Enter a descriptive **Name** for the rule in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Users.
4. In the **Destination** tab, set the **Destination Zone** to Data Center Applications.
5. In the **Applications** tab, **Add** the applications that correspond to the network services you want to safely enable. For example, select **activesync**, **imap**, **kerberos**, **ldap**, **ms-exchange**, and **ms-lync**.
6. In the **Service/URL Category** tab, keep the **Service** set to **application-default**.
7. In the **Actions** tab, set the **Action Setting** to **Allow**.
8. Select **Profiles** as the **Profile Type**. Select the **default** profiles for **Antivirus** and **URL Filtering** and the **strict** profiles for **Vulnerability Protection** and **Anti-Spyware** and select the **File Blocking** profile you configured for general traffic.
9. Verify that **Log at Session End** is enabled. Only traffic that matches a security rule will be logged.
10. Click **OK**.

Define Basic Security Policy Rules (Continued)



Step 6	Save your policies to the running configuration on the firewall.	Click **Commit**.

Step 7	To verify that you have set up your basic policies effectively, test whether your security policy rules are being evaluated and determine which security policy rule applies to a traffic flow.	To verify the policy rule that matches a flow, use the following CLI command:
		`test security-policy-match source <IP_address> destination <IP_address> destination port <port_number> application <application_name> protocol <protocol_number>`
		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 client in the user zone with the IP address 10.35.14.150 when it sends a DNS query to the DNS server in the data center:
		`test security-policy-match source 10.35.14.150 destination 10.43.2.2 application dns protocol 53`
		
		`"Network Infrastructure" {`
		`from Users;`
		`source any;`
		`source-region none;`
		`to Data_Center;`
		`destination any;`
		`destination-region none;`
		`user any;`
		`category any;`
		`application/service dns/any/any/any;`
		`action allow;`
		`icmp-unreachable: no`
		`terminal yes;`
		`}`

 © Palo Alto Networks, Inc.

Assess 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. Use this information to identify where you need to create more granular security policy rules.

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. 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.
• Determine what updates/modifications are required for your network security policy rules and implement the changes.	For example: • Evaluate whether to allow web 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.
• Work with Logs.	Specifically, view the traffic and threat logs (**Monitor > Logs**).  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 Network Traffic	
• View AutoFocus Threat Data for Logs.	Review the AutoFocus intelligence summary for artifacts in your logs. An *artifact* is an item, property, activity, or behavior associated with logged events on the firewall. The intelligence summary reveals the number of sessions and samples in which WildFire detected the artifact. Use WildFire verdict information (benign, grayware, malware) and AutoFocus matching tags to look for potential risks in your network.  AutoFocus tags created by Unit 42, the Palo Alto Networks threat intelligence team, call attention to advanced, targeted campaigns and threats in your network. From the AutoFocus intelligence summary, you can start an AutoFocus search for artifacts and assess their pervasiveness within global, industry, and network contexts.
• 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.

 © 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 the use of evasion, tunneling, or circumvention techniques. The threat prevention features on the firewall include the WildFire service, Security Profiles that support Antivirus, Anti-Spyware, Vulnerability Protection, URL Filtering, File Blocking and Data Filtering capabilities, the Denial of Service (DoS) and Zone protection functionality, and AutoFocus threat intelligence.

Threat Prevention contains more in-depth information on how to protect your network from threats. For details on how to scan encrypted (SSH or SSL) traffic for threats, see Decryption. Visit Applipedia and Threat Vault to learn more about the applications and threats that Palo Alto Networks products can identify, respectively.



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 a Basic Security Policy.

To begin protecting your network from threats, start here:

- Enable Basic WildFire Forwarding
- Scan Traffic for Threats
- Control Access to Web Content
- Enable AutoFocus Threat Intelligence

Enable Basic WildFire Forwarding

WildFire is a cloud-based virtual environment that analyzes and executes unknown samples (files and email links) and determines the samples to be malicious, grayware, or benign. With WildFire enabled, a Palo Alto Networks firewall can forward unknown samples to WildFire for analysis. For newly-discovered malware, WildFire generates a signature to detect the malware and distributes it to all firewalls with active WildFire licenses. This enables global firewalls to detect and prevent malware found by a single firewall.

A basic WildFire service is included as part of the Palo Alto Networks next generation firewall and does not require a WildFire subscription. With the basic WildFire service, you can enable the firewall to forward portable executable (PE) files. Additionally, if do not have a WildFire subscription, but you do have a Threat Prevention subscription, you can receive signatures for malware WildFire identifies every 24- 48 hours (as part of the antivirus updates).

Beyond the basic WildFire service, a WildFire subscription is required for the firewall to:

- Get the latest WildFire signatures every five minutes.
- Forward advanced file types and email links for analysis.
- Use the WildFire API.
- Use a WF-500 appliance to host a WildFire private cloud or a WildFire hybrid cloud.

If you have a WildFire subscription, go ahead and get started with WildFire to get the most out of your subscription. Otherwise, take the following steps to enable basic WildFire forwarding:

Enable Basic WildFire Forwarding		
Step 1	Confirm that your firewall is registered and that you have a valid support account as well as any subscriptions you require.	1. Go to the Palo Alto Networks Customer Support web site, log in, and select **My Devices**. 2. Verify that the firewall is listed. If it is not listed, see Register the Firewall. 3. (Optional) If you have a Threat Prevention subscription, be sure to Activate Licenses and Subscriptions.
Step 2	Configure WildFire forwarding settings.	1. Select **Device > Setup > WildFire** and edit the General Settings. 2. Set the **WildFire Public Cloud** field to: `wildfire.paloaltonetworks.com`. 3. Review the **File Size Limits** for PEs the firewall forwards for WildFire analysis. set the **Size Limit** for PEs that the firewall can forward to the maximum available limit of 10 MB.  As a WildFire best practice, set the **Size Limit** for PEs to the maximum available limit of 10 MB. 4. Click **OK** to save your changes.
Step 3	Enable the firewall to forward PEs for analysis.	1. Select **Objects > Security Profiles > WildFire Analysis** and **Add** a new profile rule. 2. **Name** the new profile rule. 3. Click **Add** to create a forwarding rule and enter a name. 4. In the **File Types** column, add **pe** files to the forwarding rule. 5. In the **Analysis** column, select **public-cloud** to forward PEs to the WildFire public cloud. 6. Click **OK**.
Step 4	Apply the new WildFire Analysis profile to traffic that the firewall allows.	1. Select **Policies > Security** and either select an existing policy or create a new policy as described in Set Up a Basic Security Policy. 2. Select **Actions** and in the Profile Settings section, set the **Profile Type** to **Profiles**. 3. Select the **WildFire Analysis** profile you just created to apply that profile rule to all traffic this policy allows. 4. Click **OK**.
Step 5	Enable the firewall to forward decrypted SSL traffic for WildFire analysis.	
Step 6	Review and implement WildFire best practices to ensure that you are getting the most of WildFire detection and prevention capabilities.	
Step 7	Click **Commit** to save your configuration updates.	
Step 8	Verify that the firewall is forwarding PE files to the WildFire public cloud.	Select **Monitor > Logs > WildFire Submissions** to view log entries for PEs the firewall successfully submitted for WildFire analysis. The **Verdict** column displays whether WildFire found the PE to be malicious, grayware, or benign.

 © Palo Alto Networks, Inc.

Enable Basic WildFire Forwarding		
Step 9	(Threat Prevention subscription only) If you have a Threat Prevention subscription, but do not have a WildFire subscription, you can still receive WildFire signature updates every 24- 48 hours.	1. Select **Device > Dynamic Updates**. 2. Check that the firewall is set to retrieve, download, and install Antivirus updates.

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 redefined 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 Vulnerability 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.

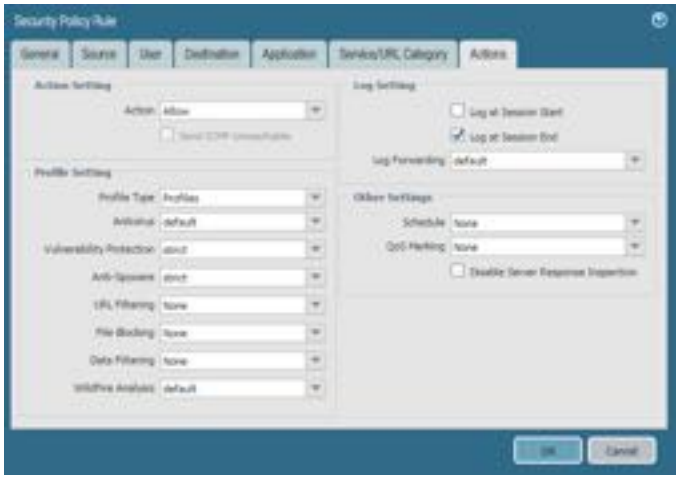
Set up Antivirus/Anti-Spyware/Vulnerability Protection		
Step 1	Verify that you have a Threat Prevention license.	• The Threat Prevention license bundles the Antivirus, Anti-Spyware, and the Vulnerability Protection features in one license. • Select **Device > Licenses** to verify that the **Threat Prevention** 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. 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 firewall; you need to configure the schedule on each firewall.

Recommendations for HA Configurations:

- **Active/Passive HA**—If the firewalls use the MGT port for content updates, configure a schedule on each firewall so that each firewall downloads and installs content independently. If the firewalls are using a data port for content updates, the passive firewall will not perform downloads while it is in the passive state. In this case set a schedule on each peer and enable **Sync To Peer** to ensure that content updates on the active peer sync to the passive peer.
- **Active/Active HA**—If the firewalls use the MGT port for content updates, configure a schedule on each firewall, but do not enable **Sync To Peer**. If the firewalls are using a data port for content updates, schedule content updates on each firewall and select **Sync To Peer** to enable the active-primary firewall to download and install the content updates and then push the content update to the active-secondary peer.

 © Palo Alto Networks, Inc.

Set up Antivirus/Anti-Spyware/Vulnerability Protection (Continued)		
Step 4	Attach the security profiles to a security policy. Attach a clone of a predefined security profile to your basic Security policy rules. 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. 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** and **WildFire Analysis**, and **strict** for **Vulnerability Protection** and **Anti-Spyware**. If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down. 
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. For most traffic (including traffic on your internal network) you will want to block files that are known to carry threats or that have no real use case for upload/download. Currently, these include batch files, DLLs, Java class files, help files, Windows shortcuts (.lnk), and BitTorrent files. Additionally, to provide drive-by download protection, allow download/upload of executables and archive files (.zip and .rar), but force users to acknowledge that they are transferring a file so that they will notice that the browser is attempting to download something they were not aware of. For policy rules that allow general web browsing, be more strict with your file blocking because the risk of users unknowingly downloading malicious files is much higher. For this type of traffic you will want to attach a more strict file blocking profile that also blocks portable executable (PE) files.

Configure File Blocking		
Step 6	Configure a File Blocking profile for general use.	1. Select **Objects > Security Profiles > File Blocking** and click **Add**. 2. Enter a **Name** for the file blocking profile, for example general-file-blocking. 3. Optionally enter a **Description**, such as block-risky-apps. Click **Add** to define the profile settings. 4. Enter a **Name**, such as block-risky. 5. Set **File Types** to block. For example, **Add** the following: **bat**, **dll**, **jar**, **hlp**, **lnk**, and **torrent**. 6. Leave the **Direction** set to **both**. 7. Set the **Action** to **block**. 8. **Add** a second rule and enter a **Name**, for example continue exe and archive. 9. Set **File Types** to continue. For example, **Add** the following: **PE**, **zip** and **rar**. 10. Leave the **Direction** set to **both**. 11. Set the **Action** to **block**. 12. Click **OK** to save the profile.
Step 7	Configure a File Blocking profile for risky traffic.  When users are web browsing it is much more likely that they will download a malicious file unintentionally. Therefore, it is important to attach a stricter file blocking policy than you would attach to Security policy rules that allow access to less risk-prone application traffic.	1. On the **Objects > Security Profiles > File Blocking** page, select the file blocking profile you just created for general traffic and click **Clone**. Select the profile to clone and click **OK**. 2. Select the cloned profile and give it a new **Name**, such as strict-block-risky-apps. 3. Click in the File Types section of the block rule and **Add** the **PE** file type. 4. Click in the File Types section of the continue rule, select **PE** and click **Delete**. 5. Click **OK** to save the profile.
Step 8	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 Set Up a Basic Security Policy. 2. Click the **Actions** tab within the security policy. 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.

 © Palo Alto Networks, Inc.

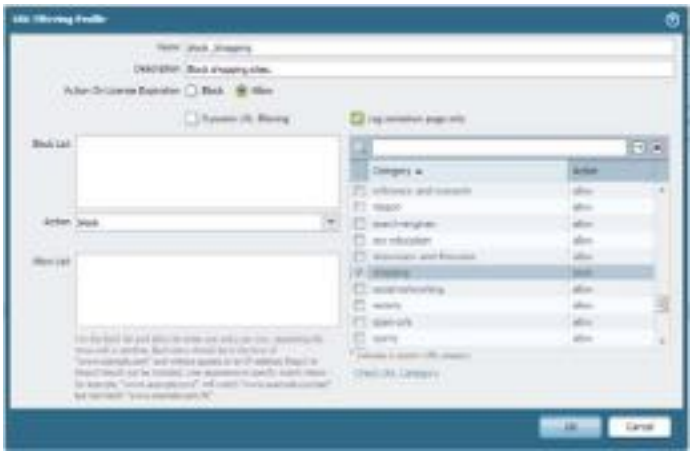
Configure File Blocking (Continued)	
Step 9 Enable response pages in the management profile for each interface on which you are attaching file blocking profile with a **continue** action.	1. Select **Network > Network Profiles > Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile. 2. Select **Response Pages**, as well as any other management services required on the interface. 3. Click **OK** to save the interface management profile. 4. Select **Network > Interfaces** and select the interface to which to attach the profile. 5. On the **Advanced > Other Info** tab, select the interface management profile you just created. 6. Click **OK** to save the interface settings.
Step 10 Save the configuration.	1. Click **Commit**.
Step 11 Test the file blocking configuration.	From a client PC in the trust zone of the firewall, attempt to download an.exe file from a website in the Internet zone. Make sure the file is blocked as expected based on the action you defined in the file blocking profile: • If you selected **alert** as the action, check the data filtering log to make sure you see a log entry for the request. • If you selected **block** as the action, the File Blocking Block Page response page should display. • 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. 

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. 2. Select **Device > Licenses** and verify that the URL Filtering license is valid. 
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. 2. Choose a region (North America, Europe, APAC, Japan) and then click **OK** to start the download. 3. After the download completes, click **Activate**. 
Step 3	Create a URL filtering profile. 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**. 2. Select the default profile and then click **Clone**. The new profile will be named default-1. 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. 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. 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: • If you do not care about traffic to a particular category (that is you neither want to block it nor log it), select **allow**. • For visibility into traffic to sites in a category, select **alert**. • 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**. • To prevent access to traffic that matches the associated policy, select **block** (this also generates a log entry).  2. Click **OK** to save the URL filtering profile.
Step 5 Attach the URL filtering profile to a security policy.	1. Select **Policies > Security**. 2. Select the desired policy to modify it and then click the **Actions** tab. 3. If this is the first time you are defining a security profile, select **Profiles** from the **Profile Type** drop-down. 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.) 5. Click **OK** to save the profile. 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. 2. Select **Response Pages**, as well as any other management services required on the interface. 3. Click **OK** to save the interface management profile. 4. Select **Network > Interfaces** and select the interface to which to attach the profile. 5. On the **Advanced > Other Info** tab, select the interface management profile you just created. 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: • If you selected **alert** as the action, check the data filtering log to make sure you see a log entry for the request. • If you selected the **continue** action, the URL Filtering Continue and Override Page response page should display. **Continue** to the site. • If you selected **block** as the action, the URL Filtering and Category Match Block Page response page should display as follows: 

Enable AutoFocus Threat Intelligence

With a valid AutoFocus subscription, you can compare the activity on your network with the latest threat data available on the AutoFocus portal. Connecting your firewall and AutoFocus unlocks the following features:

- Ability to view an AutoFocus intelligence summary for session artifacts recorded in the firewall logs.
- Ability to open an AutoFocus search for log artifacts from the firewall.

The AutoFocus intelligence summary reveals the prevalence of an artifact on your network and on a global scale. The WildFire verdicts and AutoFocus tags listed for the artifact indicate whether the artifact poses a security risk.

 © Palo Alto Networks, Inc.

Enable AutoFocus Threat Intelligence on the Firewall	
Step 1 Verify that the AutoFocus license is activated on the firewall.	1. Select **Device > Licenses** to verify that the AutoFocus Device License is installed and valid (check the expiration date). 2. If the firewall doesn't detect the license, see Activate Licenses and Subscriptions.
Step 2 Connect the firewall to AutoFocus.	1. Select **Device > Setup > Management** and edit the AutoFocus settings. 2. Enter the **AutoFocus URL**: `https://autofocus.paloaltonetworks.com:10443` 3. Use the **Query Timeout** field to set the duration of time for the firewall to attempt to query AutoFocus for threat intelligence data. If the AutoFocus portal does not respond before the end of the specified period, the firewall closes the connection. As a best practice, set the query timeout to the default value of 15 seconds. AutoFocus queries are optimized to complete within this duration. 4. Select **Enabled** to allow the firewall to connect to AutoFocus. 5. Click **OK**. 6. **Commit** your changes to retain the AutoFocus settings upon reboot.
Step 3 Connect AutoFocus to the firewall.	1. Log in to the AutoFocus portal: https://autofocus.paloaltonetworks.com 2. Select **Settings**. 3. **Add new** remote systems. 4. Enter a descriptive **Name** to identify the firewall. 5. Select **PanOS** as the System Type. 6. Enter the firewall IP **Address**. 7. Click **Save changes** to add the remote system. 8. Click **Save changes** again on the Settings page to ensure the firewall is successfully added.
Step 4 Test the connection between the firewall and AutoFocus.	1. On the firewall, select **Monitor > Logs > Traffic**. 2. Verify that you can View AutoFocus Threat Data for Logs.

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.
- Replace the Certificate for Inbound Management Traffic. By default, the firewall ships with a default certificate that enables HTTPS access to the web interface over the management (MGT) interface or any other interface that supports HTTPS management traffic. To improve the security of inbound management traffic, replace the default certificate with a new certificate issued specifically for your organization.
- Configure a best-practice security policy rulebase to safely enable applications and protect your network from attack. See Best Practice Internet Gateway Security Policy for details.
- Set up High Availability—High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration and session tables are 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 a two-firewall 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 all private keys on the firewall used for cryptographic protocols. As a best practice to safeguard the keys, configure the master key on each firewall to be unique. However, if you use Panorama, you must use the same master key on Panorama and all managed firewalls. Otherwise, Panorama cannot push configurations to the firewalls.
- 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 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 allows you to better protect the firewall 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.

Firewall Administration

Administrators can configure, manage, and monitor Palo Alto Networks firewalls using the web interface, CLI, and API management interface. You can customize role-based administrative access to the management interfaces to delegate specific tasks or permissions to certain administrators.

- Management Interfaces
- Use the Web Interface
- Manage Configuration Backups
- Manage Firewall Administrators
- Reference: Web Interface Administrator Access
- Reference: Port Number Usage
- Reset the Firewall to Factory Default Settings
- Bootstrap the Firewall

Management Interfaces

You can use the following user interfaces to manage the Palo Alto Networks firewall and Panorama:

- 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 enter 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 become familiar with the nesting structure and syntax of the commands, the CLI provides quick response times and administrative efficiency.
- Use the XML API to streamline your operations and integrate with existing, internally developed applications and repositories. The XML API is a web service implemented using HTTP/HTTPS requests and responses.

 © Palo Alto Networks, Inc.

Use the Web Interface

The following topics describe how to use the firewall web interface. For detailed information about specific tabs and fields in the web interface, refer to the Web Interface Reference Guide.

- Launch the Web Interface
- Configure Banners, Message of the Day, and Logos
- Use the Administrator Login Activity Indicators to Detect Account Misuse
- Manage and Monitor Administrative Tasks
- Commit, Validate, and Preview Firewall Configuration Changes
- Use Global Find to Search the Firewall or Panorama Management Server
- Manage Locks for Restricting Configuration Changes

Launch the Web Interface

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

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

Launch the Web Interface	
Step 1	Launch an Internet browser and enter the IP address of the firewall in the URL field (https://<IP address>). By default, the management (MGT) interface allows only HTTPS access to the web interface. To enable other protocols, select **Device > Setup > Management** and edit the Management Interface Settings.
Step 2	Enter your user **Name** and **Password**. If this is your first login session, enter the default `admin` for both fields.
Step 3	If the login dialog has a banner, read it. If the dialog requires you to acknowledge reading the banner, select **I Accept and Acknowledge the Statement Below**.
Step 4	**Login** to the web interface.
Step 5	Read and **Close** the messages of the day. You can select **Do not show again** for messages you don't want to see in future login sessions. If you want to change the language that the web interface uses, click **Language** at the bottom of the web interface, select a **Language** from the drop-down, and click **OK**.

Configure Banners, Message of the Day, and Logos

A *login banner* is optional text that you can add to the login page so that administrators will see information they must know before they log in. For example, you could add a message to notify users of restrictions on unauthorized use of the firewall.

You can add colored bands that highlight overlaid text across the top (*header banner*) and bottom (*footer banner*) of the web interface to ensure administrators see critical information, such as the classification level for firewall administration.

A *message of the day* dialog automatically displays after you log in. The dialog displays messages that Palo Alto Networks embeds to highlight important information associated with a software or content release. You can also add one custom message to ensure administrators see information, such as an impending system restart, that might affect their tasks.

You can replace the default logos that appear on the login page and in the header of the web interface with the logos of your organization.

Configure Banners, Message of the Day, and Logos	
Step 1 Configure the login banner.	1. Select **Device > Setup > Management** and edit the General Settings. 2. Enter the **Login Banner** (up to 3,200 characters). 3. (Optional) Select **Force Admins to Acknowledge Login Banner** to force administrators to select an **I Accept and Acknowledge the Statement Below** check box above the banner text to activate the **Login** button. 4. Click **OK**.
Step 2 Set the message of the day.	1. Select **Device > Setup > Management** and edit the Banners and Messages settings. 2. Enable the **Message of the Day**. 3. Enter the **Message of the Day** (up to 3,200 characters).  After you enter the message and click **OK**, administrators who subsequently log in, and active administrators who refresh their browsers, see the new or updated message immediately; a commit isn't necessary. This enables you to inform other administrators of an impending commit that might affect their configuration changes. Based on the commit time that your message specifies, the administrators can then decide whether to complete, save, or undo their changes. 4. (Optional) Select **Allow Do Not Display Again** (default is disabled) to give administrators the option to suppress a message of the day after the first login session. Each administrator can suppress messages only for his or her own login sessions. In the message of the day dialog, each message will have its own suppression option. 5. (Optional) Enter a header **Title** for the message of the day dialog (default is `Message of the Day`).

 © Palo Alto Networks, Inc.

Configure Banners, Message of the Day, and Logos (Continued)		
Step 3	Configure the header and footer banners.  A bright background color and contrasting text color can increase the likelihood that administrators will notice and read a banner. You can also use colors that correspond to classification levels in your organization.	1. Enter the **Header Banner** (up to 3,200 characters). 2. (Optional) Clear **Same Banner Header and Footer** (enabled by default) to use different header and footer banners. 3. Enter the **Footer Banner** (up to 3,200 characters) if the header and footer banners differ. 4. Click **OK**.
Step 4	Replace the logos on the login page and in the header.  The maximum size for any logo image is 128KB.	1. Select **Device > Setup > Operations** and click **Custom Logos** in the Miscellaneous section. 2. Perform the following steps for both the **Login Screen** logo and the **Main UI** (header) logo: a. Click upload  . b. Select a logo image and click **Open**.  You can preview  the image to see how PAN-OS will crop it to fit. c. Click **Close**. 3. **Commit** your changes.
Step 5	Verify that the banners, message of the day, and logos display as expected.	1. Log out to return to the login page, which displays the new logos you selected. 2. Enter your login credentials, review the banner, select **I Accept and Acknowledge the Statement Below** to enable the **Login** button, and then **Login**. A dialog displays the message of the day. Messages that Palo Alto Networks embedded display on separate pages in the same dialog. To navigate the pages, click the right  or left  arrows along the sides of the dialog or click a page selector at the bottom of the dialog. 3. (Optional) You can select **Do not show again** for the message you configured and for any messages that Palo Alto Networks embedded. 4. **Close** the message of the day dialog to access the web interface. Header and footer banners display in every web interface page with the text and colors that you configured. The new logo you selected for the web interface displays below the header banner.

Use the Administrator Login Activity Indicators to Detect Account Misuse

The last login time and failed login attempts indicators provide a visual way to detect misuse of your administrator account on a Palo Alto Networks firewall or Panorama management server. Use the last login information to determine if someone else logged in using your credentials and use the failed login attempts indicator to determine if your account is being targeted in a brute-force attack.

Use the Login Activity Indicators to Detect Account Misuse	
Step 1 View the login activity indicators to monitor recent activity on your account.	1. Log in to the web interface on your firewall or Panorama management server. 2. View the last login details located at the bottom left of the window and verify that the timestamp corresponds to your last login.  3. Look for a caution symbol to the right of the last login time information for failed login attempts. The failed login indicator appears if one or more failed login attempts occurred using your account since the last successful login. a. If you see the caution symbol, hover over it to display the number of failed login attempts.  b. Click the caution symbol to view the failed login attempts summary. Details include the admin account name, the reason for the login failure, the source IP address, and the date and time. After you successfully log in and then log out, the failed login counter resets to zero so you will see new failed login details, if any, the next time you log in.

 © Palo Alto Networks, Inc.

Use the Login Activity Indicators to Detect Account Misuse (Continued)	
	4. Locate hosts that are continually attempting to login to your firewall or Panorama management server. a. Click the failed login caution symbol to view the failed login attempts summary. b. Locate and record the source IP address of the host that attempted to log in. For example, the following figure shows multiple failed login attempts from the IP address 192.168.2.10.  c. Work with your network administrator to locate the user and host that is using the IP address that you identified. If you cannot locate the system that is performing the brute-force attack, consider renaming the account to prevent future attacks.
Step 2 Take the following actions if you detect an account compromise.	1. Select **Monitor > Logs > Configuration** and view the configuration changes and commit history to determine if your account was used to make changes without your knowledge. 2. Select **Device > Config Audit** to compare the current configuration and the configuration that was running just prior to the configuration you suspect was changed using your credentials. You can also do this using Panorama. If your administrator account was used to create a new account, performing a configuration audit helps you detect changes that are associated with any unauthorized accounts, as well. 3. Revert the configuration to a known good configuration if you see that logs were deleted or if you have difficulty determining if improper changes were made using your account. Before you commit to a previous configuration, review it to ensure that it contains the correct settings. For example, the configuration that you revert to may not contain recent changes, so apply those changes after you commit the backup configuration.

Use the following best practices to help prevent brute-force attacks on privileged accounts.

- Limit the number of failed attempts allowed before the firewall locks a privileged account by setting the number of Failed Attempts and the Lockout Time (min) in the authentication profile or in the Authentication Settings for the Management interface (**Device > Setup > Management > Authentication Settings**).
- Use Interface Management Profiles to Restrict Access.
- Enforce complex passwords for privileged accounts.

Manage and Monitor Administrative Tasks

The Task Manager displays details about all the operations that you and other administrators initiated (such as manual commits) or that the firewall initiated (such as scheduled report generation) since the last firewall reboot. You can use the Task Manager to troubleshoot failed operations, investigate warnings associated with completed commits, view details about queued commits, or cancel pending commits.



You can also view System Logs to monitor system events on the firewall or view Config Logs to monitor firewall configuration changes.

Manage and Monitor Administrative Tasks	
Step 1	Click **Tasks** at the bottom of the web interface.
Step 2	**Show** only **Running** tasks (in progress) or **All** tasks (default). Optionally, filter the tasks by type: • **Jobs**—Administrator-initiated commits, firewall-initiated commits, and software or content downloads and installations. • **Reports**—Scheduled reports. • **Log Requests**—Log queries that you trigger by accessing the **Dashboard** or a **Monitor** page.
Step 3	Perform any of the following actions: • **Display or hide task details**—By default, the Task Manager displays the Type, Status, Start Time, and Messages for each task. To see the End Time and Job ID for a task, you must manually configure the display to expose those columns. To display or hide a column, open the drop-down in any column header, select **Columns**, and select or deselect the column names as needed. • **Investigate warnings or failures**—Read the entries in the Messages column for task details. If the column says `Too many messages`, click the corresponding entry in the Type column to see more information. • **Display a commit description**—If an administrator entered a description when configuring a commit, you can click **Commit Description** in the Messages column to display the description. • **Check the position of a commit in the queue**—The Messages column indicates the queue position of commits that are in progress. • **Cancel pending commits**—Click **Clear Commit Queue** to cancel all pending commits (available only to predefined administrative roles). To cancel an individual commit, click **x** in the Action column for that commit (the commit remains in the queue until the firewall dequeues it). You cannot cancel commits that are in progress.

Commit, Validate, and Preview Firewall Configuration Changes

A commit is the process of activating changes that you made to the firewall configuration. The firewall queues commit operations in the order you and other administrators initiate them. If the queue already has the maximum number of commits (which varies by platform), you must wait for the firewall to process a pending commit before initiating a new commit. To cancel pending commits or view details about commits of any status, see Manage and Monitor Administrative Tasks. To check which changes a commit will activate, you can run a commit preview.



For details on candidate and running configurations, see Manage Configuration Backups.
To prevent multiple administrators from making configuration changes during concurrent sessions, see Manage Locks for Restricting Configuration Changes.

 © Palo Alto Networks, Inc.

When you initiate a commit, the firewall checks the validity of the changes before activating them. The validation output displays conditions that either block the commit (errors) or that are important to know but that do not block the commit (warnings). For example, validation could indicate an invalid route destination that you need to fix for the commit to succeed. To identify and fix configuration errors before initiating a commit, you can validate changes without committing. A pre-commit validation displays the same errors and warnings as a commit, including reference errors, rule shadowing, and application dependency warnings. Pre-commit validations are useful if your organization allows commits only within certain time windows; you can find and fix errors to avoid failures that could cause you to miss a commit window.

Preview, Validate, or Commit Firewall Configuration Changes		
Step 1	Configure the commit, validation, or preview options.	1. Click **Commit** at the top of the web interface. 2. (Optional) Exclude certain types of configuration changes. These options are included (enabled) by default.  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 the policy and object configuration and the device and network configuration. • **Include Device and Network configuration** • **Include Policy and Object configuration**—This is available only on firewalls for which multiple virtual systems capability is disabled. • **Include Shared Object configuration**—This is available only on firewalls with multiple virtual systems. • **Include Virtual System configuration**—This is available only on firewalls with multiple virtual systems. **Select All virtual systems** (default) or **Select one or more virtual systems** in the list. 3. (Optional) Enter a **Description** for the commit. A brief summary of what changed in the configuration is useful to other administrators who want to know what changes were made without performing a configuration audit.
Step 2	(Optional) Preview the changes that the commit will activate. This can be useful if, for example, you don't remember all your changes and you're not sure you want to activate all of them. The firewall displays the changes in a new window that shows the running and candidate configurations side by side using colors to highlight the differences line by line.	1. Click **Preview Changes**. 2. Select the **Lines of Context**, which is the number of lines from the compared configuration files to display before and after each highlighted difference. These additional lines help you correlate the preview output to settings in the web interface.  Because the preview results display in a new window, your browser must allow pop-up windows. If the preview window does not open, refer to your browser documentation for the steps to unblock pop-up windows. 3. Close the preview window when you finish reviewing the changes.
Step 3	(Optional) Validate the changes before you commit to ensure the commit will succeed.	1. Click **Validate Changes**. The results display all the errors and warnings that an actual commit would display. 2. Resolve any errors that the validation results identify.

Preview, Validate, or Commit Firewall Configuration Changes (Continued)	
Step 4 Commit your configuration changes.	Click **Commit**. To view details about commits that are pending (which you can still cancel), in progress, completed, or failed, see Manage and Monitor Administrative Tasks.

Use Global Find to Search the Firewall or Panorama Management Server

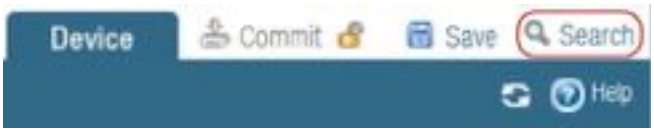
Global Find enables you to search the candidate configuration on a firewall or on Panorama for a particular string, such as an IP address, object name, policy rule name, threat ID, or application name. The search results are grouped by category and provide links to the configuration location in the web 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.

Watch the video.

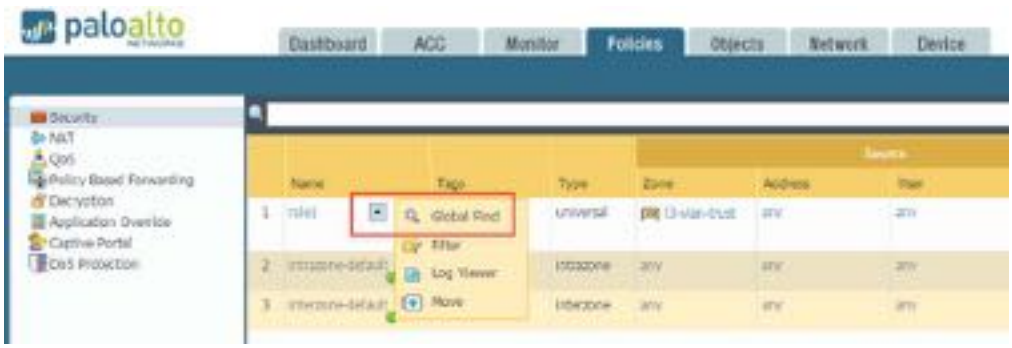
Global Find will not search dynamic content (such as logs, address ranges, or allocated DHCP addresses). In the case of DHCP, you can search on a DHCP server attribute, such as the DNS entry, but you cannot search for individual addresses allocated to users. Global Find also does not search for individual 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.

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 select **Global Find**:



 © Palo Alto Networks, Inc.

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 Search (located on the upper right of the web interface) to see a list of the last 20 searches. Click an item in the list to rerun that search. Search history is unique to each administrator account.

Manage Locks for Restricting Configuration Changes

Locking the candidate or running configuration prevents other administrators from changing the configuration until you manually remove the lock, a superuser removes the lock, or the firewall automatically removes it (after a commit). Locks ensure that administrators don't make conflicting changes to the same settings or interdependent settings during concurrent login sessions.

The firewall queues commit requests and performs them in the order that administrators initiate the commits. For details, see Commit, Validate, and Preview Firewall Configuration Changes. To view the status of queued commits, see Manage and Monitor Administrative Tasks.

Manage Locks for Restricting Configuration Changes	
• View details about current locks. For example, you can check whether other administrators have set locks and read comments they entered to explain the locks.	Click the lock at the top of the web interface. An adjacent number indicates the number of current locks.

Manage Locks for Restricting Configuration Changes (Continued)	
• Lock a configuration.	1. Click the lock at the top of the web interface. The lock image varies based on whether existing locks are or are not set. 2. **Take a Lock** and select the lock **Type**: • **Config**—Blocks other administrators from changing the candidate configuration. • **Commit**—Blocks other administrators from changing the running configuration. 3. (Firewall with multiple virtual systems only) Select a **Location** to lock the configuration for a specific virtual system or the **Shared** location. 4. (Optional) As a best practice, enter a **Comment** so that other administrators will understand the reason for the lock. 5. Click **OK** and **Close**.
• Unlock a configuration. Only a superuser or the administrator who locked the configuration can manually unlock it. However, the firewall automatically removes a lock after completing the commit operation.	1. Click the lock at the top of the web interface. 2. Select the lock entry in the list. 3. Click **Remove Lock**, **OK**, and **Close**.
• Configure the firewall to automatically lock the running configuration when you change the candidate configuration. This setting applies to all administrators.	1. Select **Device > Setup > Management** and edit the General Settings. 2. Select **Automatically Acquire Commit Lock** and then click **OK** and **Commit**.

 © Palo Alto Networks, Inc.

Manage Configuration Backups

The running configuration comprises all settings you have committed and that are therefore active, such as policy rules that currently block or allow various types of traffic in your network. The candidate configuration is a copy of the running configuration plus any inactive changes that you made after the last commit. Backing up versions of the running or candidate configuration enables you to later restore those versions on the firewall. For example, if a commit validation shows that the current candidate configuration has more errors than you are able or have time to fix, then you can restore a previous candidate configuration or revert to the running configuration.



See Commit, Validate, and Preview Firewall Configuration Changes for related information.

- Back Up a Configuration
- Restore a Configuration

Back Up a Configuration

Creating configuration backups enables you to later Restore a Configuration. This is useful when you want to revert the firewall to all the settings of an earlier configuration because you can perform the restoration as a single operation instead of manually reconfiguring each setting in the current configuration. You can either save backups locally on the firewall or export backups to an external host.

When you commit changes, the firewall automatically saves a new version of the running configuration. If a system event or administrator action causes the firewall to reboot, it automatically reverts to the current version of the running configuration, which the firewall stores in a file named running-config.xml. However, the firewall does not automatically save a backup of the candidate configuration; you must manually save a backup of the candidate configuration as a snapshot file using either the default name (.snapshot.xml) or a custom name.



When you edit a setting and click **OK**, the firewall updates the candidate configuration but does not save a backup snapshot.

Additionally, saving changes does not activate them. To activate changes, perform a commit (see Commit, Validate, and Preview Firewall Configuration Changes).

As a best practice, back up any important configuration to a host external to the firewall.

Back Up a Configuration	
Step 1 Save a local backup snapshot of the candidate configuration if it contains changes that you want to preserve in the event the firewall reboots. These are changes you are not ready to commit—for example, changes you cannot finish in the current login session.	Perform one of the following tasks based on whether you want to overwrite the default snapshot (.snapshot.xml) or create a snapshot with a custom name: • Overwrite the default snapshot—Click **Save** at the top of the web interface. • Create a custom-named snapshot: a. Select **Device > Setup > Operations** and **Save named configuration snapshot**. b. Enter a **Name** for the snapshot or select an existing snapshot to overwrite. c. Click **OK** and **Close**.
Step 2 Export a candidate configuration, a running configuration, or the firewall state information to a host external to the firewall.	Select **Device > Setup > Operations** and click an export option: • **Export named configuration snapshot**—Export the current running configuration, a named candidate configuration snapshot, or a previously imported configuration (candidate or running). The firewall exports the configuration as an XML file with the **Name** you specify. • **Export configuration version**—Select a **Version** of the running configuration to export as an XML file. The firewall creates a version whenever you commit configuration changes. • **Export device state**—Export the firewall state information as a bundle. Besides the running configuration, the state information includes device group and template settings pushed from Panorama. If the firewall is a GlobalProtect portal, the information also includes certificate information, a list of satellites, and satellite authentication information. If you replace a firewall or portal, you can restore the exported information on the replacement by importing the state bundle.

Restore a Configuration

Restoring a firewall configuration overwrites the current candidate configuration with another configuration. This is useful when you want to revert all firewall settings used in an earlier configuration; you can perform this restoration as a single operation instead of manually reconfiguring each setting in the current configuration.

The firewall automatically saves a new version of the running configuration whenever you commit changes and you can restore any of those versions. However, you must manually save a candidate configuration to later restore it (see Back Up a Configuration).

Restore a Configuration	
• Restore the current running configuration. This operation undoes all the changes you made to the candidate configuration since the last commit.	1. Select **Device > Setup > Operations** and **Revert to running configuration**. 2. Click **Yes** to confirm the operation.

 © Palo Alto Networks, Inc.

Restore a Configuration (Continued)	
• Restore the default snapshot of the candidate configuration. This is the snapshot that you create or overwrite when you click **Save** at the top right of the web interface.	1. Select **Device > Setup > Operations** and **Revert to last saved configuration**. 2. Click **Yes** to confirm the operation. 3. (Optional) Click **Commit** to overwrite the running configuration with the snapshot.
• Restore a previous version of the running configuration that is stored on the firewall. The firewall creates a version whenever you commit configuration changes.	1. Select **Device > Setup > Operations** and **Load configuration version**. 2. Select a configuration **Version** and click **OK**. 3. (Optional) Click **Commit** to overwrite the running configuration with the version you just restored.

• Restore one of the following:
  • Current running configuration (named running-config.xml)
  • Custom-named version of the running configuration that you previously imported
  • Custom-named candidate configuration snapshot (instead of the default snapshot)
	1. Select **Device > Setup > Operations** and click **Load named configuration snapshot**.
	2. Select the snapshot **Name** and click **OK**.
	3. (Optional) Click **Commit** to overwrite the running configuration with the snapshot.

• Restore a running or candidate configuration that you previously exported to an external host.	1. Select **Device > Setup > Operations**, click **Import named configuration snapshot**, **Browse** to the configuration file on the external host, and click **OK**. 2. Click **Load named configuration snapshot**, select the **Name** of the configuration file you just imported, and click **OK**. 3. (Optional) Click **Commit** to overwrite the running configuration with the snapshot you just imported.
• Restore state information that you exported from a firewall. Besides the running configuration, the state information includes device group and template settings pushed from Panorama. If the firewall is a GlobalProtect portal, the information also includes certificate information, a list of satellites, and satellite authentication information. If you replace a firewall or portal, can you can restore the information on the replacement by importing the state bundle.	Import state information: 1. Select **Device > Setup > Operations**, click **Import device state**, **Browse** to the state bundle, and click **OK**. 2. (Optional) Click **Commit** to apply the imported state information to the running configuration.

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

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).

 © Palo Alto Networks, Inc.

- **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 need to access to perform their jobs.

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: • **Device** role—**superuser**, **superreader**, **deviceadmin**, **devicereader**, or **None** • **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.

Administrative Authentication

You can configure the following types of administrator authentication:

Account Type	Authentication Method	Description
Local	Local (no database)	The administrator account credentials and the authentication mechanisms are local to the firewall. You can further secure local accounts by setting global password complexity and expiration settings for all accounts or by creating a password profile that defines password expiration settings for specific accounts. For details, see Configure an Administrative Account.

Account Type	Authentication Method	Description
Local	Local database	The firewall uses a local database to store the administrator account credentials and to perform authentication. 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

If you have already configured Administrative Roles and external authentication services (if applicable), you can Configure an Administrative Account. Otherwise, perform one of the other procedures listed below to configure administrative accounts for specific types of authentication.



Administrative accounts specify how administrators authenticate to the firewall. To configure how the firewall authenticates to administrators, see Replace the Certificate for Inbound Management Traffic.

- 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

Configure an Administrative Account

Administrative accounts specify roles and authentication methods for the administrators of Palo Alto Networks firewalls.

 © Palo Alto Networks, Inc.

Configure an Administrative Account	
Step 1 (Optional) Define password complexity and expiration settings for administrator accounts that are local to the firewall. These settings can help protect the firewall against unauthorized access by making it harder for attackers to guess passwords. You cannot configure these settings for local accounts that use a local database or external service for authentication.	1. Define global password complexity and expiration settings for all local administrators. a. Select **Device > Setup > Management** and edit the Minimum Password Complexity settings. b. Select **Enabled**. c. Define the password settings and click **OK**. 2. Define a Password Profile if you want certain local administrators to have password expiration settings that override the global settings. a. Select **Device > Password Profiles** and **Add** a profile. b. Enter a **Name** to identify the profile. c. Define the password expiration settings and click **OK**.
Step 2 Add an administrative account.	1. Select **Device > Administrators** and **Add** an administrator. 2. Enter a user **Name**. 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. In this case, you must enter and confirm a **Password**. 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** (default) and select a dynamic role. If the dynamic role is **virtual system administrator**, add one or more virtual systems that the virtual system administrator is allowed to manage. 5. (Optional) Select a **Password Profile** for local administrators. This option is available only if you set the **Authentication Profile** to **None**. 6. 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. 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.

Configure Kerberos SSO and External or Local Authentication for Administrators (Continued)		
Step 2	Configure a local database or external server profile. Required for local database or external authentication.	• Local database authentication—Perform the following tasks: a. Configure the user account. b. (Optional) Configure a user group. • External authentication—Perform one of the following tasks: • Configure a RADIUS Server Profile. • Configure a TACACS+ Server Profile. • Configure an LDAP Server Profile. • Configure a Kerberos Server Profile.
Step 3	Configure an authentication profile. 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 4).	Configure an Authentication Profile and Sequence.
Step 4	Configure an administrator account.	Configure an Administrative Account. • For local database authentication, specify the **Name** of the user you defined in Step 2. • 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. You will use this CA certificate to sign the client certificate of each administrator.	Create a Self-Signed Root CA Certificate. Alternatively, Import a Certificate and Private Key from your enterprise CA.

 © Palo Alto Networks, Inc.

Configure Certificate-Based Administrator Authentication to the Web Interface (Continued)		
Step 2	Configure a certificate profile for securing access to the web interface.	Configure a Certificate Profile. • Set the **Username Field** to **Subject**. • In the CA Certificates section, **Add** 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. 2. Select the **Certificate Profile** you created for authenticating administrators 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 and select **Use only client certificate authentication**. 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. In the **Signed By** drop-down, select a self-signed root CA certificate.
Step 6	Export the client certificate.	1. Export a Certificate and Private Key. 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.
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. 2. When prompted, select the certificate you imported and click **OK**. The browser displays a certificate warning. 3. Add the certificate to the browser exception list. 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 (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. 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. 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. • 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**. • Select **Use Public Key Authentication (SSH)**, then **Import Key**, **Browse** to the public key you just generated, and click **OK**. 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.
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. 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): `Authenticating with public key "dsa-key-20130415"` 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

 © Palo Alto Networks, Inc.

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.

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. 2. Configure an access domain if the firewall has more than one virtual system (vsys): a. Select **Device > Access Domain**, **Add** an access domain, and enter a **Name** to identify the access domain. b. **Add** each vsys that the administrator will access, and then click **OK**. 3. Configure a RADIUS Server Profile. 4. Configure an authentication profile. Set the authentication **Type** to **RADIUS** and assign the RADIUS **Server Profile**. 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**. 6. Click **OK** and **Commit**.
Step 2 Configure the RADIUS server.	1. Add the firewall IP address or hostname as the RADIUS client. 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 Support.

Reference: Web Interface Administrator Access

You can configure privileges for an entire firewall 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). You apply the admin role to a custom role-based administrator account where you can assign one or more virtual systems. The following topics describe the privileges you can configure for custom administrator roles.

- Web Interface Access Privileges
- Panorama Web Interface Access Privileges

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 as defined by the **Device** or **Virtual System** radio button. If you select **Virtual System**, 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 topics describe how to set admin role privileges to the different parts of the web interface:

- Define Access to the Web Interface Tabs
- Provide Granular Access to the Monitor Tab
- Provide Granular Access to the Policy Tab
- Provide Granular Access to the Objects Tab
- Provide Granular Access to the Network Tab
- Provide Granular Access to the Device Tab
- Define User Privacy Settings in the Admin Role Profile
- Restrict Administrator Access to Commit and Validate Functions
- Provide Granular Access to Global Settings

 © Palo Alto Networks, Inc.

Define Access to the Web Interface Tabs

The following table describes the top-level access privileges you can assign to an admin role profile (**Device > Admin Roles**). You can enable, disable, or define read-only access privileges at the top-level tabs in 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 administrator 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 administrator 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 administrator 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 administrator 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 firewall-wide configuration information, such as User-ID, high availability, server profile or certificate configuration information. For more granular control over what objects the administrator 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. 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 administrator 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	Administrator Role Availability	Enable	Read Only	Disable
Monitor	Enables or disables access to the **Monitor** tab. If disabled, the administrator will not see this tab or any of the associated logs or reports.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Administrator 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 administrator 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 Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Traffic	Specifies whether the administrator can see the traffic logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Threat	Specifies whether the administrator can see the threat logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
URL Filtering	Specifies whether the administrator can see the URL filtering logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
WildFire Submissions	Specifies whether the administrator can see the WildFire logs. These logs are only available if you have a WildFire subscription.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Data Filtering	Specifies whether the administrator can see the data filtering logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
HIP Match	Specifies whether the administrator can see the HIP Match logs. HIP Match logs are only available if you have a GlobalProtect portal license and gateway subscription.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Configuration	Specifies whether the administrator can see the configuration logs.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
System	Specifies whether the administrator can see the system logs.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
Alarms	Specifies whether the administrator can see system-generated alarms.	Firewall: Yes Panorama: Yes 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 Panorama: Yes Device Group/Template: Yes	Yes	No	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Correlation Objects	Specifies whether the administrator can view and enable/disable the correlation objects.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Correlated Events	Specifies whether the administrator	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Packet Capture	Specifies whether the administrator 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 Panorama: No Device Group/Template: No	Yes	Yes	Yes
App Scope	Specifies whether the administrator can see the App Scope visibility and analysis tools. Enabling App Scope enables access to all of the **App Scope** charts.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Session Browser	Specifies whether the administrator 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 Panorama: No Device Group/Template: No	Yes	No	Yes
Botnet	Specifies whether the administrator 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 Panorama: No Device Group/Template: No	Yes	Yes	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
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 administrator 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 Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Manage PDF Summary	Specifies whether the administrator can view, add or delete PDF summary report definitions. With read-only access, the administrator can see PDF summary report definitions, but not add or delete them. If you disable this option, the administrator can neither view the report definitions nor add/delete them.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes
PDF Summary Reports	Specifies whether the administrator 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 Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
User Activity Report	Specifies whether the administrator can view, add or delete User Activity report definitions and download the reports. With read-only access, the administrator can see User Activity report definitions, but not add, delete, or download them. If you disable this option, the administrator cannot see this category of PDF report.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes
SaaS Application Usage Report	Specifies whether the administrator can view, add or delete a SaaS application usage report. With read-only access, the administrator can see the SaaS application usage report definitions, but cannot add or delete them. If you disable this option, the administrator can neither view the report definitions nor add or delete them.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes
Report Groups	Specifies whether the administrator can view, add or delete report group definitions. With read-only access, the administrator can see report group definitions, but not add or delete them. If you disable this option, the administrator cannot see this category of PDF report.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Email Scheduler	Specifies whether the administrator 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 administrator does not have access, you should disable the **Email Scheduler** option if you have user privacy requirements.	Firewall: Yes Panorama: Yes 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 administrator 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.  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 administrator role.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Application Statistics	Specifies whether the administrator can create a custom report that includes data from the application statistics database.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Data Filtering Log	Specifies whether the administrator can create a custom report that includes data from the Data Filtering logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Threat Log	Specifies whether the administrator can create a custom report that includes data from the Threat logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Threat Summary	Specifies whether the administrator can create a custom report that includes data from the Threat Summary database.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Traffic Log	Specifies whether the administrator can create a custom report that includes data from the Traffic logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Traffic Summary	Specifies whether the administrator can create a custom report that includes data from the Traffic Summary database.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
URL Log	Specifies whether the administrator can create a custom report that includes data from the URL Filtering logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Hipmatch	Specifies whether the administrator can create a custom report that includes data from the HIP Match logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
WildFire Log	Specifies whether the administrator can create a custom report that includes data from the WildFire logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
View Scheduled Custom Reports	Specifies whether the administrator can view a custom report that has been scheduled to generate.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
View Predefined Application Reports	Specifies whether the administrator 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 Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
View Predefined Threat Reports	Specifies whether the administrator 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 Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
View Predefined URL Filtering Reports	Specifies whether the administrator 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 Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
View Predefined Traffic Reports	Specifies whether the administrator 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 Panorama: Yes 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 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 administrator to view the corresponding policy rule base, but not add or delete rules. Disabling access to a specific type of policy prevents the administrator 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 administrator to view, add, and/or delete security rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the security rulebase, disable this privilege.	Yes	Yes	Yes
NAT	Enable this privilege to allow the administrator to view, add, and/or delete NAT rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the NAT rulebase, disable this privilege.	Yes	Yes	Yes
QoS	Enable this privilege to allow the administrator to view, add, and/or delete QoS rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the QoS rulebase, disable this privilege.	Yes	Yes	Yes
Policy Based Forwarding	Enable this privilege to allow the administrator to view, add, and/or delete Policy-Based Forwarding (PBF) rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the PBF rulebase, disable this privilege.	Yes	Yes	Yes
Decryption	Enable this privilege to allow the administrator to view, add, and/or delete decryption rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the decryption rulebase, disable this privilege.	Yes	Yes	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
Application Override	Enable this privilege to allow the administrator to view, add, and/or delete application override policy rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the application override rulebase, disable this privilege.	Yes	Yes	Yes
Captive Portal	Enable this privilege to allow the administrator to view, add, and/or delete Captive Portal rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator from seeing the Captive Portal rulebase, disable this privilege.	Yes	Yes	Yes
DoS Protection	Enable this privilege to allow the administrator to view, add, and/or delete DoS protection rules. Set the privilege to read-only if you want the administrator to be able to see the rules, but not modify them. To prevent the administrator 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 administrator will have policy definition responsibilities. If not, the administrator probably does not need access to the tab. If, however, the administrator 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 administrator the privilege to view, add, and delete the corresponding object type. Giving read-only access allows the administrator to view the already defined objects, but not create or delete any. Disabling a node prevents the administrator from seeing the node in the web interface.

Access Level	Description	Enable	Read Only	Disable
Addresses	Specifies whether the administrator can view, add, or delete address objects for use in security policy.	Yes	Yes	Yes
Address Groups	Specifies whether the administrator can view, add, or delete address group objects for use in security policy.	Yes	Yes	Yes
Regions	Specifies whether the administrator can view, add, or delete regions objects for use in security, decryption, or DoS policy.	Yes	Yes	Yes
Applications	Specifies whether the administrator can view, add, or delete application objects for use in policy.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
Application Groups	Specifies whether the administrator can view, add, or delete application group objects for use in policy.	Yes	Yes	Yes
Application Filters	Specifies whether the administrator can view, add, or delete application filters for simplification of repeated searches.	Yes	Yes	Yes
Services	Specifies whether the administrator 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 administrator can view, add, or delete service group objects for use in security policy.	Yes	Yes	Yes
Tags	Specifies whether the administrator can view, add, or delete tags that have been defined on the firewall.	Yes	Yes	Yes
GlobalProtect	Specifies whether the administrator 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 administrator 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 administrator 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 administrator can view, add, or delete dynamic block lists for use in security policy.	Yes	Yes	Yes
Custom Objects	Specifies whether the administrator 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 administrator can view, add, or delete custom data pattern signatures for use in creating custom Vulnerability Protection profiles.	Yes	Yes	Yes
Spyware	Specifies whether the administrator can view, add, or delete custom spyware signatures for use in creating custom Vulnerability Protection profiles.	Yes	Yes	Yes
Vulnerability	Specifies whether the administrator can view, add, or delete custom vulnerability signatures for use in creating custom Vulnerability Protection profiles.	Yes	Yes	Yes
URL Category	Specifies whether the administrator can view, add, or delete custom URL categories for use in policy.	Yes	Yes	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
Security Profiles	Specifies whether the administrator 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 administrator can view, add, or delete antivirus profiles.	Yes	Yes	Yes
Anti-Spyware	Specifies whether the administrator can view, add, or delete Anti-Spyware profiles.	Yes	Yes	Yes
Vulnerability Protection	Specifies whether the administrator can view, add, or delete Vulnerability Protection profiles.	Yes	Yes	Yes
URL Filtering	Specifies whether the administrator can view, add, or delete URL filtering profiles.	Yes	Yes	Yes
File Blocking	Specifies whether the administrator can view, add, or delete file blocking profiles.	Yes	Yes	Yes
Data Filtering	Specifies whether the administrator can view, add, or delete data filtering profiles.	Yes	Yes	Yes
DoS Protection	Specifies whether the administrator can view, add, or delete DoS protection profiles.	Yes	Yes	Yes
Security Profile Groups	Specifies whether the administrator can view, add, or delete security profile groups.	Yes	Yes	Yes
Log Forwarding	Specifies whether the administrator can view, add, or delete log forwarding profiles.	Yes	Yes	Yes
Decryption Profile	Specifies whether the administrator can view, add, or delete decryption profiles.	Yes	Yes	Yes
Schedules	Specifies whether the administrator 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 administrator will have network administration responsibilities, including GlobalProtect administration. If not, the administrator 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 administrator the privilege to view, add, and delete the corresponding network configurations. Giving read-only access allows the administrator to view the already-defined configuration, but not create or delete any. Disabling a node prevents the administrator from seeing the node in the web interface.

Access Level	Description	Enable	Read Only	Disable
Interfaces	Specifies whether the administrator can view, add, or delete interface configurations.	Yes	Yes	Yes
Zones	Specifies whether the administrator can view, add, or delete zones.	Yes	Yes	Yes
VLANs	Specifies whether the administrator can view, add, or delete VLANs.	Yes	Yes	Yes
Virtual Wires	Specifies whether the administrator can view, add, or delete virtual wires.	Yes	Yes	Yes
Virtual Routers	Specifies whether the administrator can view, add, modify or delete virtual routers.	Yes	Yes	Yes
IPSec Tunnels	Specifies whether the administrator can view, add, modify, or delete IPSec Tunnel configurations.	Yes	Yes	Yes
DHCP	Specifies whether the administrator can view, add, modify, or delete DHCP server and DHCP relay configurations.	Yes	Yes	Yes
DNS Proxy	Specifies whether the administrator can view, add, modify, or delete DNS proxy configurations.	Yes	Yes	Yes
GlobalProtect	Specifies whether the administrator 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 administrator can view, add, modify, or delete GlobalProtect portal configurations.	Yes	Yes	Yes
Gateways	Specifies whether the administrator can view, add, modify, or delete GlobalProtect gateway configurations.	Yes	Yes	Yes
MDM	Specifies whether the administrator can view, add, modify, or delete GlobalProtect MDM server configurations.	Yes	Yes	Yes
Device Block List	Specifies whether the administrator can view, add, modify, or delete device block lists.	Yes	Yes	Yes
QoS	Specifies whether the administrator can view, add, modify, or delete QoS configurations.	Yes	Yes	Yes
LLDP	Specifies whether the administrator 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

 © Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
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. 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. 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. 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. 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. 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

Access Level	Description	Enable	Read Only	Disable
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. 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. 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. 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. 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
BFD Profile	Controls access to the **Network Profiles > BFD Profile** node. If you disable this privilege, the administrator will not see the **Network Profiles > BFD Profile** node or be able to configure a BFD profile. A Bidirectional Forwarding Detection (BFD) profile allows you to configure BFD settings to apply to one or more static routes or routing protocols. Thus, BFD detects a failed link or BFD peer and allows an extremely fast failover. If the privilege state is set to read-only, you can view the currently configured BFD profile but cannot add or edit a BFD profile.	Yes	Yes	Yes

 © Palo Alto Networks, Inc.

Provide Granular Access to the Device Tab

To define granular access privileges for the **Device** tab, when creating or editing an admin role profile (**Device > Admin Roles**), scroll down to the **Device** node on the **WebUI** 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 firewall-wide setup configuration information, such as Management, Operations, Service, Content-ID, Wildfire or Session setup information. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
Management	Controls access to the **Management** node. If you disable this privilege, the administrator will not be able to configure settings such as the hostname, domain, timezone, authentication, logging and reporting, Panorama, management interface, banner, message, and password complexity settings, and more. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
Operations	Controls access to the **Operations** node. If you disable this privilege, the administrator will not be able to manage configuration files, or reboot or shut down the firewall, among other things.	Yes	Yes	Yes
Services	Controls access to the **Services** node. If you disable this privilege, the administrator will not be able to configure services for DNS servers, an update server, proxy server, or NTP servers, or set up service routes. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
Content-ID	Controls access to the **Content-ID** node. If you disable this privilege, the administrator will not be able to configure URL filtering or Content-ID. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
WildFire	Controls access to the **WildFire** node. If you disable this privilege, the administrator will not be able to configure WildFire settings. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
Session	Controls access to the **Session** node. If you disable this privilege, the administrator will not be able to configure session settings or timeouts for TCP, UDP or ICMP, or configure decryption or VPN session settings. If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.	Yes	Yes	Yes
HSM	Controls access to the **HSM** node. If you disable this privilege, the administrator will not be able to configure a Hardware Security Module. 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 firewall-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. If you disable this privilege, the administrator will not see the **Admin Roles** node or have access to any firewall-wide information concerning Admin Role profiles configuration. If you set this privilege to read-only, you can view the configuration information for all administrator roles configured on the firewall.	No	Yes	Yes
Administrators	Controls access to the **Administrators** node. This function can only be allowed for read-only access. If you disable this privilege, the administrator will not see the **Administrators** node or have access to information about their own administrator account. If you set this privilege to read-only, the administrator can view the configuration information for their own administrator account. They will not see any information about other administrator accounts configured on the firewall.	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. 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. If you disable this privilege, the administrator will not see or be able to configure shared gateways. 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 firewall-wide User Identification configuration information, such as User Mapping, User-ID Agents, Service, Terminal Services Agents, Group Mappings Settings or Captive Portal Settings. If you set this privilege to read-only, the administrator can view configuration information for the firewall 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. 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.  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 firewall-wide high availability configuration information such as General setup information or Link and Path Monitoring. If you set this privilege to read-only, the administrator can view High Availability configuration information for the firewall 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	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. If you set this privilege to read-only, the administrator can view Certificate configuration information for the firewall 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. If you set this privilege to read-only, the administrator can view Certificate Profiles that are currently configured for the firewall 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 firewall. If you set this privilege to read-only, the administrator can view the **OCSP Responder** configuration for the firewall 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. 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 firewall services that use SSL/TLS. 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
SCEP	Controls access to the **SCEP** node. If you disable this privilege, the administrator will not see the node or be able to define a profile that specifies simple certificate enrollment protocol (SCEP) settings for issuing unique device certificates. If you set this privilege to read-only, the administrator can view existing SCEP profiles but cannot create or edit them.	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. If you set this privilege to read-only, the administrator can view the **Response Page** configuration for the firewall 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. If you set this privilege to read-only, the administrator can view the **Log Settings > System** configuration for the firewall 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. If you set this privilege to read-only, the administrator can view the **Log Settings > Config** configuration for the firewall 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 If you set this privilege to read-only, the administrator can view the **Log Settings > HIP** configuration for the firewall 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. If you set this privilege to read-only, the administrator can view the **Log Settings > Alarms** configuration for the firewall 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. 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. 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. 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 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. 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. 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. If you disable this privilege, the administrator will not see the node or configure settings for the TACACS+ servers that authentication profiles reference. 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. 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. 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	Description	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, firewall administrators, and captive portal users. 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. 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. 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. 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. 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

 © Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
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 firewall and a remote host. 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. 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. 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. 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. If you set this privilege to read-only, the administrator can view the installed **Licenses**, but cannot perform license management functions.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
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. 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. 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

To define what private end user data an administrator has access to, when creating or editing an admin role profile (**Device > Admin Roles**), scroll down to the **Privacy** option on the **WebUI** tab.

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 disabled, 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.  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

 © Palo Alto Networks, Inc.

Access Level	Description	Enable	Read Only	Disable
Show User Names in Logs and Reports	When disabled, 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. 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 disabled, packet capture files that are normally available within the Traffic, Threat and Data Filtering logs are not displayed.	Yes	N/A	Yes

Restrict Administrator Access to Commit and Validate Functions

To restrict access to commit and validate functions when creating or editing an admin role profile (**Device > Admin Roles**), scroll down to the **Commit** and **Validate** options on the **WebUI** tab.

Access Level	Description	Enable	Read Only	Disable
Commit	When disabled, an administrator cannot commit any changes to a configuration.	Yes	N/A	Yes
Validate	When disabled, an administrator cannot validate a configuration.	Yes	N/A	Yes

Provide Granular Access to Global Settings

To define what global settings and administrator has access to, when creating or editing an admin role profile (**Device > Admin Roles**), scroll down to the **Global** option on the **WebUI** tab.

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 disabled, an administrator 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	Administrator 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**. If you set the privilege to: • read-only, the administrator can see the information but cannot edit it. • disable this privilege, the administrator cannot see or edit the information.	Panorama: Yes 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. 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. If you disable this privilege, the administrator can't see or manage HA configuration settings for the Panorama management server.	Panorama: Yes 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 Device Group/Template: No	Yes	No	Yes
Administrators	Specifies whether the administrator can view Panorama administrator account details. 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. If you disable this privilege, the administrator can't see information about any Panorama administrator account, including his or her own.	Panorama: Yes Device Group/Template: No	No	Yes	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Admin Roles	Specifies whether the administrator can view Panorama administrator roles. 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. If you disable this privilege, the administrator can't see or manage Panorama administrator roles.	Panorama: Yes 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.) If you set this privilege to read-only, the administrator can view Panorama access domain configurations but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama access domain configurations.	Panorama: Yes Device Group/Template: No  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. If you set this privilege to read-only, the administrator can view Panorama authentication profiles but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama authentication profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Authentication Sequence	Specifies whether the administrator can view, add, edit, delete, or clone authentication sequences for Panorama administrators. If you set this privilege to read-only, the administrator can view Panorama authentication sequences but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama authentication sequences.	Panorama: Yes 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. 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. If you disable this privilege, the administrator can't view, add, edit, tag, delete, or install updates on managed firewalls.  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 Device Group/Template: Yes	Yes (No for Device Group and Template roles)	Yes	Yes
Templates	Specifies whether the administrator can view, edit, add, or delete templates and template stacks. If you set the privilege to read-only, the administrator can see template and stack configurations but can't manage them. If you disable this privilege, the administrator can't see or manage template and stack configurations.	Panorama: Yes Device Group/Template: Yes  Device Group and Template administrators can see only the templates and stacks that are within the access domains assigned to those administrators.	Yes (No for Device Group and Template admins)	Yes	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Device Groups	Specifies whether the administrator can view, edit, add, or delete device groups. If you set this privilege to read-only, the administrator can see device group configurations but can't manage them. If you disable this privilege, the administrator can't see or manage device group configurations.	Panorama: Yes Device Group/Template: Yes 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. If you set this privilege to read-only, the administrator can see managed collector configurations but can't manage them. If you disable this privilege, the administrator can't view, edit, add, or delete managed collector configurations. 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 Device Group/Template: No	Yes	Yes	Yes
Collector Groups	Specifies whether the administrator can view, edit, add, or delete Collector Groups. If you set this privilege to read-only, the administrator can see Collector Groups but can't manage them. If you disable this privilege, the administrator can't see or manage Collector Groups.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
VMware Service Manager	Specifies whether the administrator can view and edit VMware Service Manager settings. If you set this privilege to read-only, the administrator can see the settings but can't perform any related configuration or operational procedures. If you disable this privilege, the administrator can't see the settings or perform any related configuration or operational procedures.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Certificate Management	Sets the default state, enabled or disabled, for all of the Panorama certificate management privileges.	Panorama: Yes 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. If you set this privilege to read-only, the administrator can see Panorama certificates but can't manage the certificates or HA keys. If you disable this privilege, the administrator can't see or manage Panorama certificates or HA keys.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Certificate Profile	Specifies whether the administrator can view, add, edit, delete or clone Panorama certificate profiles. If you set this privilege to read-only, the administrator can see Panorama certificate profiles but can't manage them. If you disable this privilege, the administrator can't see or manage Panorama certificate profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
SSL/TLS Service Profile	Specifies whether the administrator can view, add, edit, delete or clone SSL/TLS Service profiles. If you set this privilege to read-only, the administrator can see SSL/TLS Service profiles but can't manage them. If you disable this privilege, the administrator can't see or manage SSL/TLS Service profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Log Settings	Sets the default state, enabled or disabled, for all the log setting privileges.	Panorama: Yes Device Group/Template: No	Yes	No	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
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). If you set this privilege to read-only, the administrator can see the System log forwarding settings but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  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 Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator 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). If you set this privilege to read-only, the administrator can see the Config log forwarding settings but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  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 Device Group/Template: No	Yes	Yes	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Administrator 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). If you set this privilege to read-only, the administrator can see the forwarding settings of HIP Match logs but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  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 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). If you set this privilege to read-only, the administrator can see the Correlation log forwarding settings but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  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 Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator 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). If you set this privilege to read-only, the administrator can see the forwarding settings of Traffic logs but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  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 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). If you set this privilege to read-only, the administrator can see the forwarding settings of Threat logs but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  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 Device Group/Template: No	Yes	Yes	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Administrator 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). If you set this privilege to read-only, the administrator can see the forwarding settings of WildFire logs but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  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 Device Group/Template: No	Yes	Yes	Yes
Server Profiles	Sets the default state, enabled or disabled, for all the server profile privileges.  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 Device Group/Template: No	Yes	No	Yes
SNMP Trap	Specifies whether the administrator can see and configure SNMP trap server profiles. If you set this privilege to read-only, the administrator can see SNMP trap server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage SNMP trap server profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
Syslog	Specifies whether the administrator can see and configure Syslog server profiles. If you set this privilege to read-only, the administrator can see Syslog server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage Syslog server profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Email	Specifies whether the administrator can see and configure email server profiles. If you set this privilege to read-only, the administrator can see email server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage email server profiles.	Panorama: Yes 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. If you set this privilege to read-only, the administrator can see the RADIUS server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage the RADIUS server profiles.	Panorama: Yes 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. If you disable this privilege, the administrator can't see the node or configure settings for the TACACS+ servers that authentication profiles reference. 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 Device Group/Template: No	Yes	Yes	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Administrator 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. If you set this privilege to read-only, the administrator can see the LDAP server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage the LDAP server profiles.	Panorama: Yes 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. If you set this privilege to read-only, the administrator can see the Kerberos server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage the Kerberos server profiles.	Panorama: Yes 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. If you set this privilege to read-only, the administrator can view the scheduled exports but can't manage them. If you disable this privilege, the administrator can't see or manage the scheduled exports.	Panorama: Yes Device Group/Template: No	Yes	No	Yes

Access Level	Description	Administrator 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. 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. If you disable this privilege, the administrator can't see Panorama software updates, see the associated release notes, or perform any related operations.  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 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. 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. If you disable this privilege, the administrator can't see Panorama content updates, see the associated release notes, or perform any related operations.  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 Device Group/Template: No	Yes	Yes	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Administrator 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 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. 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 Device Group/Template: No	Yes	Yes	Yes
Device Deployment	Sets the default state, enabled or disabled, for all the device deployment privileges.  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 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. 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. 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 Device Group/Template: Yes	Yes	Yes	Yes

Access Level	Description	Administrator 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. 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. 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 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. 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. 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 Device Group/Template: Yes	Yes	Yes	Yes

 © Palo Alto Networks, Inc.

Access Level	Description	Administrator 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. 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. 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 Device Group/Template: Yes	Yes	Yes	Yes
Licenses	Specifies whether the administrator can view, refresh, and activate firewall licenses. If you set this privilege to read-only, the administrator can view firewall licenses but can't refresh or activate those licenses. If you disable this privilege, the administrator can't view, refresh, or activate firewall licenses.	Panorama: Yes 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. If you set this privilege to read-only, the administrator can view the Panorama master key configuration but can't change it. If you disable this privilege, the administrator can't see or edit the Panorama master key configuration.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Panorama Web Interface Access Privileges

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 administrator 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 profile. If you assign superuser privileges for the CLI to a **Panorama** Admin Role profile, 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 administrator 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 administrator 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 administrator 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

 © Palo Alto Networks, Inc.

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 administrator 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 firewall-wide configuration information, such as User-ID, High Availability, server profile or certificate configuration information. For more granular control over what objects the administrator 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. 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. For more granular control over what objects the administrator 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
Privacy	Controls access to the privacy settings described in Define User Privacy Settings in the Admin Role Profile.	Yes	No	Yes
Validate	When disabled, an administrator cannot validate a configuration.	Yes	No	Yes
Commit	Sets the default state (enabled or disabled) for all the commit settings described below (Panorama, Device Groups, Templates, Force Template Values, Collector Groups).	Yes	No	Yes
Panorama	When disabled, an administrator cannot commit changes to the Panorama configuration.	Yes	No	Yes
Device Groups	When disabled, an administrator cannot commit changes to device groups.	Yes	No	Yes

Access Level	Description	Enable	Read Only	Disable
Templates	When disabled, an administrator cannot commit changes to templates.	Yes	No	Yes
Force Template Values	This privilege controls access to the **Force Template Values** option in the Commit dialog. When disabled, an administrator cannot replace overridden settings in local firewall configurations with settings that Panorama pushes from a template.	Yes	No	Yes
Collector Groups	When disabled, an administrator cannot commit changes to Collector Groups.	Yes	No	Yes
Global	Controls access to the global settings (system alarms) described in Provide Granular Access to Global Settings.	Yes	No	Yes

 © Palo Alto Networks, Inc.

Reference: Port Number Usage

The following tables list the ports that firewalls and Panorama 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 GlobalProtect
- Ports Used for User-ID

Ports Used for Management Functions

The firewall and Panorama use the following ports 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. 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 a Log Collector uses to Forward Traps to an SNMP Manager. This port doesn't need to be open on the Palo Alto Networks firewall. 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.
514 514 6514	TCP UDP 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. 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 6081 6082	TCP TCP 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 28260	TCP 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 29281	IP 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 firewall (Active/Passive) or active-primary (Active/Active) to the passive firewall (Active/Passive) or active-secondary (Active/Active). The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default. 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

Panorama uses the following ports.

Destination Port	Protocol	Description
22	TCP	Used for communication from a client system to the Panorama CLI interface.

 © Palo Alto Networks, Inc.

Destination Port	Protocol	Description
443	TCP	Used for communication from a client system to the Panorama web interface.
3978	TCP	Used for communication between Panorama and managed firewalls or managed collectors, as well as for communication among managed collectors in a Collector Group: • 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. • Log Collectors use this destination port to forward logs to Panorama. • 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) 28260 (5.0 and later) 49160 (5.0 and earlier)	TCP TCP 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) 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 GlobalProtect

GlobalProtect uses the following ports.

Destination Port	Protocol	Description
443	TCP	Used for communication between GlobalProtect agents and portals, or GlobalProtect agents and gateways and for SSL tunnel connections. GlobalProtect gateways also use this port to collect host information from GlobalProtect agents and perform host information profile (HIP) checks.
4501	UDP	Used for IPSec tunnel connections between GlobalProtect agents and gateways.

For tips on how to use a loopback interface to provide access to GlobalProtect on different ports and addresses, refer to Can GlobalProtect Portal Page be Configured to be Accessed on any Port?.

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.
514 514 6514	TCP UDP 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 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 firewall 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.
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. The User-ID agent also uses this port to connect to client systems to perform Windows Management Instrumentation (WMI) probing.

 © Palo Alto Networks, Inc.

Destination Port	Protocol	Description
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. 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). If your computer does not have a 9-pin serial port, use a USB-to-serial port connector. 2. Enter your login credentials. 3. Enter the following CLI command: `debug system maintenance-mode` The firewall will reboot in the maintenance mode.
Step 2 Reset the system to factory default settings.	1. When the firewall reboots, press `Enter` to continue to the maintenance mode menu. 2. Select `Factory Reset` and press `Enter`. 3. Select `Factory Reset` and press `Enter` again. The firewall will reboot without any configuration settings. The default username and password to log in to the firewall is admin/admin. 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.

Bootstrap the Firewall

Bootstrapping speeds up the process of configuring and licensing the firewall to make it operational on the network with or without Internet access. Bootstrapping allows you to choose whether to configure the firewall with a basic configuration file (init-cfg.txt) so that it can connect to Panorama and obtain the complete configuration or to fully configure the firewall with the basic configuration and the optional bootstrap.xml file.

▲ USB Flash Drive Support

▲ Sample init-cfg.txt Files

▲ Prepare a USB Flash Drive for Bootstrapping a Firewall

▲ Bootstrap a Firewall Using a USB Flash Drive

USB Flash Drive Support

The USB flash drive that bootstraps a hardware-based Palo Alto Networks firewall must support one of the following:

- File Allocation Table 32 (FAT32)
- Third Extended File System (ext3)

The firewall can bootstrap from the following flash drives with USB2.0 or USB3.0 connectivity:

USB Flash Drives Supported	
Kingston	• Kingston SE9 8GB (2.0) • Kingston SE9 16GB (3.0) • Kingston SE9 32GB (3.0)
SanDisk	• SanDisk Cruzer Fit CZ33 8GB (2.0) • SanDisk Cruzer Fit CZ33 16GB (2.0) • SanDisk Cruzer CZ36 16GB (2.0) • SanDisk Cruzer CZ36 32GB (2.0) • SanDisk Extreme CZ80 32GB (3.0)
Silicon Power	• Silicon Power Jewel 32GB (3.0) • Silicon Power Blaze 16GB (3.0)
PNY	• PNY Attache 16GB (2.0) • PNY Turbo 32GB (3.0)

Sample init-cfg.txt Files

An init-cfg.txt file is required for the bootstrap process; this file is a basic configuration file that you create using a text editor. You create this file is Step 5 in Prepare a USB Flash Drive for Bootstrapping a Firewall. The following sample init-cfg.txt files show the parameters that are supported in the file; the parameters that you must provide are in bold.

Sample init-cfg.txt (Static IP Address)	Sample init-cfg.txt (DHCP Client)
type=static **ip-address=10.5.107.19** **default-gateway=10.5.107.1** **netmask=255.255.255.0** **ipv6-address=2001:400:f00::1/64** **ipv6-default-gateway=2001:400:f00::2** hostname=Ca-FW-DC1 panorama-server=10.5.107.20 panorama-server-2=10.5.107.21 tplname=FINANCE_TG4 dgname=finance_dg dns-primary=10.5.6.6 dns-secondary=10.5.6.7 op-command-modes=multi-vsys,jumbo-frame dhcp-send-hostname=no dhcp-send-client-id=no dhcp-accept-server-hostname=no dhcp-accept-server-domain=no	**type=dhcp-client** ip-address= default-gateway= netmask= ipv6-address= ipv6-default-gateway= hostname=Ca-FW-DC1 panorama-server=10.5.107.20 panorama-server-2=10.5.107.21 tplname=FINANCE_TG4 dgname=finance_dg dns-primary=10.5.6.6 dns-secondary=10.5.6.7 op-command-modes=multi-vsys,jumbo-frame dhcp-send-hostname=yes dhcp-send-client-id=yes dhcp-accept-server-hostname=yes dhcp-accept-server-domain=yes

The following table describes the fields in the init-cfg.txt file. The type is required; if the type is static, the IP address, default gateway and netmask are required, or the IPv6 address and IPv6 default gateway are required.

Fields in the init-cfg.txt File	
Field	**Description**
type	(Required) Type of management IP address: static or dhcp-client.
ip-address	(Required for IPv4 static management address) IPv4 address. The firewall ignores this field if the type is dhcp-client.
default-gateway	(Required for IPv4 static management address) IPv4 default gateway for the management interface. The firewall ignores this field if the type is dhcp-client.
netmask	(Required for IPv4 static management address) IPv4 netmask. The firewall ignores this field if the type is dhcp-client.
ipv6-address	(Required for IPv6 static management address) IPv6 address and /prefix length of the management interface. The firewall ignores this field if the type is dhcp-client.
ipv6-default-gateway	(Required for IPv6 static management address) IPv6 default gateway for the management interface. The firewall ignores this field if the type is dhcp-client.
hostname	(Optional) Host name for the firewall.

 © Palo Alto Networks, Inc.

Fields in the init-cfg.txt File	
Field	**Description**
panorama-server	(Recommended) IPv4 or IPv6 address of the primary Panorama server.
panorama-server-2	(Optional) IPv4 or IPv6 address of the secondary Panorama server.
tplname	(Recommended) Panorama template name.
dgname	(Recommended) Panorama device group name.
dns-primary	(Optional) IPv4 or IPv6 address of the primary DNS server.
dns-secondary	(Optional) IPv4 or IPv6 address of the secondary DNS server.
vm-auth-key	(VM-Series firewalls only) Virtual machine authentication key.
op-command-modes	(Optional) Enter multi-vsys, jumbo-frame, or both separated by a comma only. Enables multiple virtual systems and jumbo frames while bootstrapping.
dhcp-send-hostname	(DHCP client type only) The DHCP server determines a value of yes or no. If yes, the firewall sends its hostname to the DHCP server.
dhcp-send-client-id	(DHCP client type only) The DHCP server determines a value of yes or no. If yes, the firewall sends its client ID to the DHCP server.
dhcp-accept-server-hostname	(DHCP client type only) The DHCP server determines a value of yes or no. If yes, the firewall accepts its hostname from the DHCP server.
dhcp-accept-server-domain	(DHCP client type only) The DHCP server determines a value of yes or no. If yes, the firewall accepts its DNS server from the DHCP server.

Prepare a USB Flash Drive for Bootstrapping a Firewall

You can use a USB flash drive to bootstrap a physical firewall. However, to do so you must upgrade to PAN-OS 7.1 and Reset the Firewall to Factory Default Settings. For security reasons, you can bootstrap a firewall only when it is in factory default state or has all private data deleted.

Prepare a USB Flash Drive for Bootstrapping a Firewall		
Step 1	Obtain serial numbers (S/Ns) and auth codes for support subscriptions from your order fulfillment email.	
Step 2	Register S/Ns of new firewalls on the Customer Support portal.	1. Go to support.paloaltonetworks.com, log in, and select **Assets > Register New Device > Register device using Serial Number or Authorization Code**. 2. Follow the steps to Register the Firewall. 3. Click **Submit**.

Prepare a USB Flash Drive for Bootstrapping a Firewall (Continued)		
Step 3	Activate authorization codes on the Customer Support portal, which creates license keys.	1. Go to support.paloaltonetworks.com, log in, and select the **Assets** tab. 2. For each S/N you just registered, click the **Action** link. 3. Select **Activate Auth-Code**. 4. Enter the **Authorization code** and click **Agree** and **Submit**.
Step 4	Add the S/Ns in Panorama.	Complete Step 1 in Add a Firewall as a Managed Device in the Panorama Administrator's Guide.
Step 5	Create the init-cfg.txt file.	Create the init-cfg.txt file, a mandatory file that provides bootstrap parameters. The fields are described in Sample init-cfg.txt Files. If the init-cfg.txt file is missing, the bootstrap process will fail and the firewall will boot up with the default configuration in the normal boot-up sequence. There are no spaces between the key and value in each field; do not add spaces because they cause failures during parsing on the management server side. You can have multiple init-cfg.txt files—one each for different remote sites—by prepending the S/N to the file name. For example: 0008C200105-init-cfg.txt 0008C200107-init-cfg.txt If no prepended filename is present, the firewall uses the init-cfg.txt file and proceeds with bootstrapping.
Step 6	(Optional) Create the bootstrap.xml file.	The optional bootstrap.xml file is a complete firewall configuration that you can export from an existing production firewall. 1. Select **Device > Setup > Operations > Export named configuration snapshot**. 2. Select the **Name** of the saved or the running configuration. 3. Click **OK**. 4. Rename the file as `bootstrap.xml`.

 © Palo Alto Networks, Inc.

Prepare a USB Flash Drive for Bootstrapping a Firewall (Continued)	
Step 7 Create and download the bootstrap bundle from the Customer Support portal. For a physical firewall, the bootstrap bundle requires only the /license and /config directories.	Use one of the following methods to create and download the bootstrap bundle: • Use Method 1 to create a bootstrap bundle specific to a remote site (you have only one init-cfg.txt file). • Use Method 2 to create one bootstrap bundle for multiple sites. **Method 1** 1. On your local system, go to support.paloaltonetworks.com and log in. 2. Select **Assets**. 3. Select the S/N of the firewall you want to bootstrap. 4. Select **Bootstrap Container**. 5. Click **Select**. 6. Upload and **Open** the init-cfg.txt file you created. 7. (Optional) Select the bootstrap.xml file you created and **Upload Files**. You must use a bootstrap.xml file from a firewall of the same model and PAN-OS version. 8. Select **Bootstrap Container Download** to download a tar.gz file named `bootstrap_<S/N>_<date>.tar.gz` to your local system. This bootstrap container includes the license keys associated with the S/N of the firewall. **Method 2** Create a tar.gz file on your local system with two top-level directories: /license and /config. Include all licenses and all init-cfg.txt files with S/Ns prepended to the filenames. The license key files you download from the Customer Support portal have the S/N in the license file name. PAN-OS checks the S/N in the file name against the firewall S/N while executing the bootstrap process.

Step 8 Import the tar.gz file you created to a PAN-OS 7.1 firewall using Secure Copy (SCP) or TFTP.	Access the CLI and enter one of the following commands:
	• `tftp import bootstrap-bundle file <path and filename> from <host IP address>`
	For example:
	`tftp import bootstrap-bundle file /home/userx/bootstrap/devices/pa5000.tar.gz from 10.1.2.3`
	• `scp import bootstrap-bundle from <<user>@<host>:<path to file>>`
	For example:
	`scp import bootstrap-bundle from userx@10.1.2.3:/home/userx/bootstrap/devices/pa200_bootstrap_bundle.tar.gz`

Prepare a USB Flash Drive for Bootstrapping a Firewall (Continued)	

Step 9 Prepare the USB flash drive.	1. Insert the USB flash drive into the firewall that you used in Step 8.
	2. Enter the following CLI operational command, using your tar.gz filename in place of "`pa5000.tar.gz`". This command formats the USB flash drive, unzips the file, and validates the USB flash drive:
	`request system bootstrap-usb prepare from pa5000.tar.gz`
	3. Press `y` to continue. The following message displays when the USB drive is ready:
	`USB prepare completed successfully.`
	4. Remove the USB flash drive from the firewall.
	5. You can prepare as many USB flash drives as needed.

Step 10 Deliver the USB flash drive to your remote site.	If you used Method 2 to create the bootstrap bundle, you can use the same USB flash drive content for bootstrapping firewalls at multiple remote sites. You can translate the content into multiple USB flash drives or a single USB flash drive used multiple times.

Bootstrap a Firewall Using a USB Flash Drive

After you receive a new Palo Alto Networks firewall and a USB flash drive loaded with bootstrap files, you can bootstrap the firewall.

Microsoft Windows and Apple Mac operating systems are unable to read the bootstrap USB flash drive because the drive is formatted using an ext4 file system. You must install third-party software or use a Linux system to read the USB drive.

Bootstrap a Firewall Using a USB Flash Drive	
Step 1	The firewall must be in a factory default state or must have all private data deleted.
Step 2	To ensure connectivity with your corporate headquarters, cable the firewall by connecting the management interface (MGT) using an Ethernet cable to one of the following: • An upstream modem • A port on the switch or router • An Ethernet jack in the wall
Step 3	Insert the USB flash drive into the USB port on the firewall and power on the firewall. The factory default firewall bootstraps itself from the USB flash drive. The firewall Status light turns from yellow to green when the firewall is configured; autocommit is successful.

 © Palo Alto Networks, Inc.

Bootstrap a Firewall Using a USB Flash Drive

Step 4 Verify bootstrap completion. You can see basic status logs on the console during the bootstrap and you can verify that the process is complete.

1. If you included Panorama values (panorama-server, tplname, and dgname) in your init-cfg.txt file, check Panorama managed devices, device group, and template name.
2. Verify the general system settings and configuration by accessing the web interface and selecting **Dashboard > Widgets > System** or by using the CLI operational commands `show system info` and `show config running`.
3. Verify the license installation by selecting **Device > Licenses** or by using the CLI operational command `request license info`.
4. If you have Panorama configured, manage the content versions and software versions from Panorama. If you do not have Panorama configured, use the web interface to manage content versions and software versions.

 © Palo Alto Networks, Inc.

Authentication

Many of the services that Palo Alto Networks firewalls and Panorama 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 firewall and Panorama services, procedures to configure them, how to test authentication profiles, and how to troubleshoot authentication issues:

- Configure an Authentication Profile and Sequence
- Configure Kerberos Single Sign-On
- Configure Local Database Authentication
- Configure External Authentication
- Test Authentication Server Connectivity
- Troubleshoot Authentication Issues

Configure an Authentication Profile and Sequence

An authentication profile defines the authentication service that validates the login credentials of firewall or Panorama administrators and Captive Portal or GlobalProtect end users. The authentication service can be a local database (firewalls only), an external service (RADIUS, TACACS+, LDAP, or Kerberos server), or Kerberos single sign-on (SSO).

Some networks have multiple databases for different users and user groups (for example, TACACS+ and LDAP). To authenticate users in such cases, configure an authentication sequence, which is a ranked order of authentication profiles that the firewall or Panorama matches a user against during login. The firewall or Panorama checks against each profile in sequence until one successfully authenticates the user (the firewall always checks the local database first if the sequence includes one). A user is denied access only if authentication fails for all the profiles in the authentication sequence.

Configure an Authentication Profile and Sequence		
Step 1	Create a Kerberos keytab. Required if the firewall or Panorama 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 firewall or Panorama.
Step 2	Configure a local database (firewall only) or external server profile (firewall or Panorama). Required for local database or external authentication.	• Local database authentication—Perform the following tasks: a. Configure the user account. b. (Optional) Configure a user group. • External authentication—Perform one of the following tasks: • Configure a RADIUS Server Profile. • Configure a TACACS+ Server Profile. • Configure an LDAP Server Profile. • Configure a Kerberos Server Profile.

 © Palo Alto Networks, Inc.

Configure an Authentication Profile and Sequence (Continued)	
Step 3 Configure an authentication profile. Define one or both of the following: • Kerberos SSO—The firewall or Panorama first tries SSO authentication. If that fails, it falls back to the specified authentication **Type**. • Local database or external authentication—The firewall or Panorama prompts the user to enter login credentials, and uses its local database (firewalls only) or an external service to authenticate the user.	1. Select **Device > Authentication Profile** and **Add** the authentication profile. 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. 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) Select the **User Domain** and **Username Modifier** options as follows 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. • 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 firewall or Panorama. 7. Select **Advanced** and **Add** the users and groups that can authenticate with this profile. You can select users and 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 firewall or Panorama 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 firewall or Panorama 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.

Configure an Authentication Profile and Sequence (Continued)		
Step 4	Configure an authentication sequence. Required if you want the firewall or Panorama to try multiple authentication profiles to authenticate users. The firewall or Panorama evaluates the profiles in top-to-bottom order until one profile successfully authenticates the user.	1. Select **Device > Authentication Sequence** and **Add** the authentication sequence. 2. Enter a **Name** to identify the authentication sequence. 3. If the firewall has more than one virtual system (vsys), select a **Location** (a vsys or **Shared**) where the sequence is available.  To expedite the authentication process, the best practice is to **Use domain to determine authentication profile**: the firewall or Panorama 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 firewall or Panorama doesn't find a match, or if you clear the check box, it tries the profiles in the top-to-bottom sequence. 4. **Add** each authentication profile. To change the evaluation order of the profiles, select a profile and **Move Up** or **Move Down**. 5. Click **OK** to save the authentication sequence.
Step 5	Assign the authentication profile or sequence.	Assign the authentication profile or sequence to an administrator account or to a firewall service for end users.  Test Authentication Server Connectivity to verify that an authentication profile can communicate with the back-end authentication server and that the authentication request succeeded.

 © Palo Alto Networks, Inc.

Configure Kerberos Single Sign-On

Palo Alto Networks firewalls and Panorama 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 firewall or Panorama 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 the firewall or Panorama 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 firewall or Panorama. 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. 2. Enter the following command, where <*principal_name*>, <*password*>, and <*algorithm*> are variables. The Kerberos principal name and password are of the firewall or Panorama, not the user. `ktpass /princ <principal_name> /pass <password> /crypto <algorithm> /ptype KRB5_NT_PRINCIPAL /out <file_name>.keytab` If the firewall is in FIPS/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 firewall or Panorama account. 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 and Sequence: 1. Enter the **Kerberos Realm** (usually the DNS domain of the users, except that the realm is uppercase). 2. **Import** the **Kerberos Keytab** that you created for the firewall or Panorama.
Step 3	Assign the authentication profile to the administrator account or to the Captive Portal settings.	• Configure an administrator account. • Configure Captive Portal.

Configure Local Database Authentication

You can use a local firewall database instead of an external service to manage user account credentials and authentication. For example, you might create a local database of users and user groups for specialized purposes if you don't have permission to add them to the directory servers that your organization uses to manage regular accounts and groups. Local database authentication is available for firewall administrators and for Captive Portal and GlobalProtect end users.



If your network supports Kerberos single sign-on (SSO), you can configure local authentication as a fall-back in case SSO fails. For details, see Configure Kerberos SSO and External or Local Authentication for Administrators.

You can also Configure an Administrative Account to use local account management and authentication without a local database, but only for firewall administrators.

Configure Local Database Authentication		
Step 1	Configure the user account.	1. Select **Device > Local User Database > Users** and click **Add**. 2. Enter a user **Name** for the administrator. 3. Enter a **Password** and **Confirm Password** or enter a **Password Hash**. 4. **Enable** the account (enabled by default) and click **OK**.
Step 2	Configure a user group. Required if your users require group membership.	1. Select **Device > Local User Database > User Groups** and click **Add**. 2. Enter a **Name** to identify the group. 3. **Add** each user who is a member of the group and click **OK**.
Step 3	Configure an authentication profile.	Set the authentication **Type** to **Local Database**.

Step 4	Assign the authentication profile to an administrator account or firewall service.	• Administrators—Configure an Administrative Account:
		  • Specify the **Name** of a user you defined in Step 1.
		  • Assign the **Authentication Profile** that you configured for the account.
		• End users—For all services, you must assign the **Authentication Profile** that you configured for the accounts:
		  • Configure Captive Portal.
		  • Configure the GlobalProtect portal.
		  • Configure the GlobalProtect gateway.

Step 5	Verify that the firewall can communicate with the authentication server.	Test a Local Database Authentication Profile.

 © Palo Alto Networks, Inc.

Configure External Authentication

Palo Alto Networks firewalls and Panorama 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 firewalls and Panorama 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 firewall or Panorama 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 firewall/Panorama service by assigning the authentication profile to it.

- Configure Authentication Server Profiles
- Enable External Authentication for Users and Services

Configure Authentication Server Profiles

- Configure a RADIUS Server Profile
- RADIUS Vendor-Specific Attributes Support
- Configure a TACACS+ Server Profile
- Configure an LDAP Server Profile
- Configure a Kerberos Server Profile
- Set CHAP or PAP Authentication for RADIUS and TACACS+ Servers

Configure a RADIUS Server Profile

You can configure the firewall or Panorama to use a RADIUS server for managing administrator accounts (if they are not local). You can also configure the firewall to use a RADIUS server for authenticating end users 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 Support.



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 or PAP Authentication for RADIUS and TACACS+ Servers.

When sending authentication requests to a RADIUS server, the firewall and Panorama use the authentication profile name as the network access server (NAS) identifier, even if the profile is assigned to an authentication sequence for the service that initiates the authentication process.

Configure a RADIUS Server Profile	
Step 1 Add a RADIUS server profile.	1. Select **Device > Server Profiles > RADIUS** and click **Add**. 2. Enter a **Profile Name** to identify the server profile. 3. For a firewall with more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available. 4. For the **Timeout**, enter an interval in seconds after which an authentication request times out (range is 1-30, default is 3). 5. Enter the number of automatic **Retries** following a **Timeout** before the request fails (range is 1-5, default is 3). 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). 7. Click **OK**.
Step 2 Implement the RADIUS server profile.	1. Assign the RADIUS server profile to an authentication profile or sequence. 2. Test a RADIUS Authentication Profile to verify that the firewall or Panorama can connect to the RADIUS server. 3. Assign the authentication profile or sequence to an administrator account or to a firewall service for end users. 4. **Commit** your changes.

RADIUS Vendor-Specific Attributes Support

Palo Alto Networks firewalls and Panorama 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 firewalls or Panorama) 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.

 © Palo Alto Networks, Inc.

Name	Number	Value
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	

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 or PAP Authentication for RADIUS and TACACS+ Servers.

Configure a TACACS+ Server Profile	
Step 1 Add a TACACS+ server profile.	1. Select **Device > Server Profiles > TACACS+** and click **Add**. 2. Enter a **Profile Name** to identify the server profile. 3. For a firewall with more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available. 4. For the **Timeout**, enter an interval in seconds after which an authentication request times out (range is 1-20, default is 3). 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. 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). 7. Click **OK**.

Configure a TACACS+ Server Profile (Continued)	
Step 2 Implement the TACACS+ server profile.	1. Assign the TACACS+ server profile to an authentication profile or sequence. 2. Test a TACACS+ Authentication Profile to verify that the firewall or Panorama can connect to the TACACS+ server. 3. Assign the authentication profile or sequence to an administrator account or to a firewall service for end users. 4. **Commit** your changes.

Configure an LDAP Server Profile

An LDAP server profile enables you to:

- Authenticate administrators and end users of Palo Alto Networks firewalls and Panorama.
- Define security rules based on user or user 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.

 © Palo Alto Networks, Inc.

Configure an LDAP Server Profile	
Step 1 Add an LDAP server profile.	1. Select **Device > Server Profiles > LDAP** and click **Add**. 2. Enter a **Profile Name** to identify the server profile. 3. For a firewall with more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available. 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). 5. Select the server **Type** from the drop-down: **active-directory**, **e-directory**, **sun**, or **other**. 6. If you want the firewall or Panorama 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 firewall or Panorama uses depends on the server **Port**: • 389 (default)—TLS (Specifically, the firewall or Panorama uses the Start TLS operation, which upgrades the initial plaintext connection to TLS.) • 636—SSL • Any other port—The firewall or Panorama first tries to use TLS. If the directory server doesn't support TLS, the firewall or Panorama falls back to SSL. 7. To improve security, you can select the **Verify Server Certificate for SSL sessions** check box (it is cleared by default) so that the firewall or Panorama 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 firewall or Panorama verifies the certificate in two respects: • The certificate is trusted and valid. For the firewall or Panorama 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. • The certificate name must match the host **Name** of the LDAP server. The firewall or Panorama 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. 8. Click **OK**.

Configure an LDAP Server Profile (Continued)	
Step 2 Implement the LDAP server profile.	1. Assign the LDAP server profile to an authentication profile or sequence. 2. Test an LDAP Authentication Profile to verify that the firewall or Panorama can connect to the LDAP server. 3. Assign the authentication profile or sequence to an administrator account or to a firewall service for end users. 4. **Commit** your changes.

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 Add a Kerberos server profile.	1. Select **Device > Server Profiles > Kerberos** and click **Add**. 2. Enter a **Profile Name** to identify the server profile. 3. For a firewall with more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available. 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). 5. Click **OK**.
Step 2 Implement the Kerberos server profile.	1. Assign the Kerberos server profile to an authentication profile or sequence. 2. Test a Kerberos Authentication Profile to verify that the firewall or Panorama can connect to the Kerberos server. 3. Assign the authentication profile or sequence to an administrator account or to a firewall service for end users. 4. **Commit** your changes.

Set CHAP or PAP Authentication for RADIUS and TACACS+ Servers

When you configure a Palo Alto Networks firewall or Panorama to use RADIUS or TACACS+ server authentication for a particular service (such as Captive Portal), it first tries to authenticate to the server using Challenge-Handshake Authentication Protocol (CHAP). The firewall or Panorama falls back to Password

 © Palo Alto Networks, Inc.

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 firewall or Panorama falls back to PAP for a particular RADIUS or TACACS+ server, it 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 firewall or Panorama reverts to first trying CHAP for that server.

If you want the firewall or Panorama to always use a specific protocol for authenticating to the RADIUS or TACACS+ server, enter 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.

Enable External Authentication for Users and Services

Palo Alto Networks firewalls and Panorama can use external services to authenticate administrators and end users.

Enable External Authentication		
Step 1	Configure an external server profile.	• Configure a RADIUS Server Profile. • Configure a TACACS+ Server Profile. • Configure an LDAP Server Profile. • Configure a Kerberos Server Profile.
Step 2	Assign the server profile to an authentication profile. Optionally, you can assign multiple authentication profiles to an authentication sequence.	1. Configure an Authentication Profile and Sequence. 2. Test Authentication Server Connectivity.

Step 3	Assign the authentication profile or sequence to an administrator account or to a firewall service for end users.	• Administrators: Configure an Administrative Account.
		• End user services:
		  • Configure Captive Portal.
		  • Configure the GlobalProtect portal.
		  • Configure the GlobalProtect gateway.

Test Authentication Server Connectivity

After you configure an authentication profile on a Palo Alto Networks firewall or Panorama, you can use the test authentication feature to determine if it can communicate with the back-end authentication server and if the authentication request succeeded. 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 examples:

- Run the Test Authentication Command
- Test a Local Database Authentication Profile
- Test a RADIUS Authentication Profile
- Test a TACACS+ Authentication Profile
- Test an LDAP Authentication Profile
- Test a Kerberos Authentication Profile

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. 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. To define the target vsys: `admin@PA-3060>` **set system setting target-vsys <vsys-name>** For example, if the user is defined in vsys2, run the following command: `admin@PA-3060>` **set system setting target-vsys vsys2** The `target-vsys` command is per-login session, so the system clears the option when you log off.

 © Palo Alto Networks, Inc.

Run the Test Authentication Command	
Step 4	Test an authentication profile by entering the following command: `admin@PA-3060>` **test authentication authentication-profile <authentication-profile-name> username <username> password** For example, to test an authentication profile named my-profile for a user named bsimpson, run the following command: `admin@PA-3060>` **test authentication authentication-profile my-profile username bsimpson password**  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. 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. For example use cases on the supported authentication profile types, see Test Authentication Server Connectivity.  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.

Test a Local Database Authentication Profile

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.

Local Database Authentication Profile Test Example	
Step 3	(Firewalls with virtual systems configured) Define the target virtual system that the test command will access. 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. To define the target vsys: `admin@PA-3060>` **set system setting target-vsys <vsys-name>** For example, if the user is defined in vsys2, run the following command: `admin@PA-3060>` **set system setting target-vsys vsys2** 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: `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:
	`Allow list check error:`
	`Do allow list check before sending out authentication request...`
	`User User1-LocalDB is not allowed with authentication profile LocalDB-Profile`
	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. 1. On the firewall, select **Device > Authentication Profile** and modify the profile named LocalDB-Profile. 2. Click the **Advanced** tab and add User1-LocalDB to the Allow List. 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 "User1-LocalDB" has an exact match in allow list`
	`Authentication by Local User Database for user "User1-LocalDB"`
	`Authentication succeeded for Local User Database user "User1-LocalDB"`

Test a RADIUS Authentication Profile

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.

 © Palo Alto Networks, Inc.

RADIUS Authentication Profile Test Example

Step 3 (Firewalls with virtual systems configured) Define the target virtual system that the test command will access.

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.

To define the target vsys:

```
admin@PA-3060> set system setting target-vsys <vsys-name>
```

For example, if the user is defined in vsys2, run the following command:

```
admin@PA-3060> set system setting target-vsys vsys2
```



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:

`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:

```
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 error: Invalid RADIUS response received - Bad MD5
Authentication failed against RADIUS server at 10.5.104.99:1812 for user "User2-RADIUS"
```

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.

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"
```

Test a TACACS+ Authentication Profile

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. 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. To define the target vsys: `admin@PA-3060>` **set system setting target-vsys <*vsys-name*>** For example, if the user is defined in vsys2, run the following command: `admin@PA-3060>` **set system setting target-vsys vsys2**  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: `admin@PA-3060>` **test authentication authentication-profile TACACS-Profile username User3-TACACS password**

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.

 © Palo Alto Networks, Inc.

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"`

Test an LDAP Authentication Profile

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. 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. To define the target vsys: `admin@PA-3060>` **set system setting target-vsys <*vsys-name*>** For example, if the user is defined in vsys2, run the following command: `admin@PA-3060>` **set system setting target-vsys vsys2** 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: `admin@PA-3060>` **test authentication authentication-profile LDAP-Profile username User4-LDAP password**

LDAP Authentication Profile Test Example	

Step 5	When prompted, enter the password for the User4-LDAP account. The following output shows that the test failed:
	`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`
	`parse error of dn and attributes for user "User4-LDAP"`
	`Authentication failed against LDAP server at 10.5.104.99:389 for user "User4-LDAP"`
	`Authentication failed for user "User4-LDAP"`
	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.

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"`

Test a Kerberos Authentication Profile

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.

 © Palo Alto Networks, Inc.

Kerberos Authentication Profile Test Example

Step 3 (Firewalls with virtual systems configured) Define the target virtual system that the test command will access.

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.

To define the target vsys:

`admin@PA-3060>` **set system setting target-vsys <*vsys-name*>**

For example, if the user is defined in vsys2, run the following command:

`admin@PA-3060>` **set system setting target-vsys vsys2**



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:

`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:

```
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: 'Bad-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 failure: Wrong realm: 'Bad-MGMT-GROUP.LOCAL' (code: -1765328316)
Authentication failed against KERBEROS server at 10.5.104.99:88 for user "User5-Kerberos"
Authentication failed for user "User5-Kerberos"
```

In this case, the output shows `Wrong realm,` which indicates that the Kerberos realm has an incorrect value.

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"
```

Troubleshoot Authentication Issues

When users fail to authenticate to a Palo Alto Networks firewall or Panorama, 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

Display the number of locked user accounts associated with the authentication profile (`auth-profile`), authentication sequence (`is-seq`), or virtual system (`vsys`).
  To unlock users, use the following operational command:
`request authentication [unlock-admin | unlock-user]`
	`show authentication locked-users`
	`{`
	`vsys <value> |`
	`auth-profile <value> |`
	`is-seq`
	`{yes | no}`
	`{auth-profile | vsys} <value>`
	`}`

Use the `debug authentication` command to troubleshoot authentication events.
Use the `show` options to display authentication request statistics and the current debugging level:
• `show` displays the current debugging level for the authentication service (authd).
• `show-active-requests` displays the number of active checks for authentication requests, allow lists, and locked user accounts.
• `show-pending-requests` displays the number of pending checks for authentication requests, allow lists, and locked user accounts.
• `connection-show` displays authentication request and response statistics for all authentication servers or for a specific protocol type.
Use the `connection-debug` options to enable or disable authentication debugging:
• Use the `on` option to enable or the `off` option to disable debugging for authd.
• 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`
	`{`
	`on {debug | dump | error | info | warn} |`
	`show |`
	`show-active-requests |`
	`show-pending-requests |`
	`connection-show |`
	`{`
	`connection-id |`
	`protocol-type`
	`{`
	`Kerberos connection-id <value> |`
	`LDAP connection-id <value> |`
	`RADIUS connection-id <value> |`
	`TACACS+ connection-id <value> |`
	`}`
	`connection-debug-on |`
	`{`
	`connection-id |`
	`debug-prefix |`
	`protocol-type`
	`{`
	`Kerberos connection-id <value> |`
	`LDAP connection-id <value> |`
	`RADIUS connection-id <value> |`
	`TACACS+ connection-id <value> |`
	`}`
	`connection-debug-off |`
	`{`
	`connection-id |`
	`protocol-type`
	`{`
	`Kerberos connection-id <value> |`
	`LDAP connection-id <value> |`
	`RADIUS connection-id <value> |`
	`TACACS+ connection-id <value> |`
	`}`
	`connection-debug-on`
	`}`

 © Palo Alto Networks, Inc.

Certificate Management

The following topics describe the different keys and certificates that Palo Alto Networks® firewalls and Panorama 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
- Replace the Certificate for Inbound Management Traffic
- 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 firewalls and Panorama 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 firewalls and Panorama use certificates in the following applications:

- User authentication for Captive Portal, GlobalProtect™, Mobile Security Manager, and web interface access to a firewall or Panorama.
- 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 firewalls and Panorama 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 firewall or Panorama 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 firewall or Panorama 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 private 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.
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.

 © Palo Alto Networks, Inc.

Key/Certificate Usage	Description
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. 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). 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 firewalls and Panorama use digital certificates to ensure trust between parties in a secure communication session. Configuring a firewall or Panorama 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 firewall or Panorama checks the status of every certificate in the chain except the root CA certificate, for which it 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.

The firewall and Panorama support the following methods for verifying certificate revocation status. If you configure both methods, the firewall or Panorama first tries the OCSP method; if the OCSP server is unavailable, it uses 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 the firewall or Panorama.

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 firewalls or Panorama. For details, see Configure Revocation Status Verification of Certificates.

 © 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 firewalls or Panorama. 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.

Certificate Deployment

The basic approaches to deploy certificates for Palo Alto Networks firewalls or Panorama 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 the firewall or Panorama, 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
- 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**. 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. 3. If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate. 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. 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. 5. Click **OK**.
Step 2	Enable OCSP communication on the firewall.	1. Select **Device > Setup > Management**. 2. In the Management Interface Settings section, edit to select the **HTTP OCSP** check box, then click **OK**.

Configure an OCSP Responder	
Step 3 (Optional) 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**. 2. Click **Add** to create a new profile or click the name of an existing profile. 3. Select the **HTTP OCSP** check box and click **OK**. 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. 5. Select **Advanced > Other info** and select the Interface Management Profile you configured. 6. Click **OK** and **Commit**.

Configure Revocation Status Verification of Certificates

The firewall and Panorama use certificates to authenticate users and devices for such applications as Captive Portal, GlobalProtect, site-to-site IPSec VPN, and web interface access to the firewall/Panorama. To improve security, it is a best practice to configure the firewall or Panorama to verify the revocation status of certificates that it uses for device/user authentication.

Configure Revocation Status Verification of Certificates	
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. 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**. 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. • 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. • 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. 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: • 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. • 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. • 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. 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 and Panorama management server has a default master key that encrypts all the private keys and passwords in the configuration to secure them (such as the private key used for SSL Forward Proxy Decryption). For the best security posture, configure a new master key and change it periodically.

In a high availability (HA) configuration, ensure both firewalls or Panorama management servers in the pair use the same master key. If the master keys differ, HA configuration synchronization will not work properly.

Additionally, if you are using Panorama to manage your firewalls, you must use the same master key on Panorama and all managed firewalls so that Panorama can push configurations to the firewalls.



For added security, Encrypt a Master Key Using an HSM.

Be sure to store the master key in a safe location. You cannot recover the master key and the only way to restore the default master key is to Reset the Firewall to Factory Default Settings.

Configure a Master Key	
Step 1	Select **Device > Master Key and Diagnostics** and edit the Master Key section.
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	To specify the master key **Life Time**, enter the number of **Days** and/or **Hours** after which the key will expire.  You must configure a new master key before the current key expires. If the master key expires, the firewall or Panorama automatically reboots in Maintenance mode. You must then Reset the Firewall to Factory Default Settings.
Step 5	Enter a **Time for Reminder** that specifies the number of **Days** and **Hours** before the master key expires when the firewall generates an expiration alarm. The firewall automatically opens the System Alarms dialog to display the alarm.  To ensure the expiration alarm displays, select **Device > Log Settings**, edit the Alarm Settings, and **Enable Alarms**.
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
- 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 satellites 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 firewall or Panorama, 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 firewall 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 firewall 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**.

Generate a Certificate

Palo Alto Networks firewalls and Panorama 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 web interface access to the firewall/Panorama. 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	
Step 1	Select **Device > Certificate Management > Certificates > Device Certificates**.
Step 2	If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate.
Step 3	Click **Generate**.
Step 4	Select **Local** (default) as the **Certificate Type** unless you want to deploy SCEP certificates to GlobalProtect clients.
Step 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.
Step 6	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 7	If the firewall has more than one vsys and you want the certificate to be available to every vsys, select the **Shared** check box.
Step 8	In the **Signed By** field, select the root CA certificate that will issue the certificate.
Step 9	(Optional) Select an **OCSP Responder**.
Step 10	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. 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 11	Select the **Number of Bits** to define the certificate key length. Higher numbers are more secure but require more processing time.
Step 12	Select the **Digest** algorithm. From most to least secure, the options are: **sha512**, **sha384**, **sha256** (default), **sha1**, and **md5**.
Step 13	For the **Expiration**, enter the number of days (default is 365) for which the certificate is valid.
Step 14	(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.
Step 15	Click **Generate** and, in the Device Certificates page, 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.

 © Palo Alto Networks, Inc.

Generate a Certificate (Continued)	
Step 16	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.
Step 17	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.



On a Palo Alto Networks firewall or Panorama, 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. 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 firewall 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 firewall 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**: • **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 the next step. Otherwise, select the **Import Private Key** check box, enter the **Key File** or **Browse** to it, then continue to the next step.
Step 8	Enter and re-enter (confirm) the **Passphrase** used to encrypt the private key.

Import a Certificate and 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**. 2. If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for the certificate. 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. If the firewall has more than one vsys and you want the certificate to be available to every vsys, select the **Shared** check box. 7. In the **Signed By** field, select **External Authority (CSR)**. 8. If applicable, select an **OCSP Responder**. 9. (Optional) **Add** the **Certificate Attributes** to uniquely identify the firewall and the service that will use the certificate.  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. 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. 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**. 2. Enter the **Certificate Name** used to generate the CSR. 3. Enter the path and name of the PEM **Certificate File** that the CA sent, or **Browse** to it. 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**. 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. 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:

- Configure Certificate-Based Administrator Authentication to the Web Interface
- GlobalProtect agent/app authentication to portals and gateways
- SSL Forward Proxy decryption
- Obtain a Certificate from an External CA

Export a Certificate and Private Key	
Step 1	Select **Device > Certificate Management > Certificates > Device Certificates**.
Step 2	If the firewall 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**: • **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. • **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. • **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 firewalls or Panorama. 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. • Generate a Certificate. • 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**. 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. 3. If the firewall 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: 1. In the CA Certificates table, click **Add**. 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**. 3. (Optional) 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. • 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://*). • 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. 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. 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. 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: • 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. • 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. • 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. 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. 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

Palo Alto Networks firewalls and Panorama use SSL/TLS service profiles to specify a certificate and the allowed protocol versions for SSL/TLS services. The firewall and Panorama 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 the firewall or Panorama to avoid SSL/TLS versions that have known weaknesses. If a service request involves a protocol version that is outside the specified range, the firewall or Panorama downgrades or upgrades the connection to a supported version.



In the client systems that request firewall services, the certificate trust list (CTL) must include the certificate authority (CA) certificate that issued the certificate specified in the SSL/TLS service profile. Otherwise, users will see a certificate error when requesting firewall services. Most third-party CA certificates are present by default in client browsers. If an enterprise or firewall-generated CA certificate is the issuer, you must deploy that CA certificate to the CTL in client browsers.

Configure an SSL/TLS Service Profile	
Step 1	For each desired service, generate or import a certificate on the firewall (see Obtain Certificates).  Use only signed certificates, not CA certificates, in SSL/TLS service profiles.
Step 2	Select **Device > Certificate Management > SSL/TLS Service Profile**.
Step 3	If the firewall 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: • For the **Min Version**, select the earliest allowed TLS version: **TLSv1.0** (default), **TLSv1.1**, or **TLSv1.2**. • 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**.

Replace the Certificate for Inbound Management Traffic

When you first boot up the firewall or Panorama, it automatically generates a default certificate that enables HTTPS access to the web interface and XML API over the management (MGT) interface and (on the firewall only) over any other interface that supports HTTPS management traffic (for details, see Use Interface Management Profiles to Restrict Access). To improve the security of inbound management traffic, replace the default certificate with a new certificate issued specifically for your organization.



You cannot view, modify, or delete the default certificate.
Securing management traffic also involves configuring how administrators authenticate to the firewall or to Panorama.

Replace the Certificate for Inbound Management Traffic		
Step 1	Obtain the certificate that will authenticate the firewall or Panorama to the client systems of administrators.	You can simplify your Certificate Deployment by using a certificate that the client systems already trust. Therefore, we recommend that you Import a Certificate and Private Key from your enterprise certificate authority (CA) or Obtain a Certificate from an External CA; the trusted root certificate store of the client systems is likely to already have the associated root CA certificate that ensures trust.  If you Generate a Certificate on the firewall or Panorama, administrators will see a certificate error because the root CA certificate is not in the trusted root certificate store of client systems. To prevent this, deploy the self-signed root CA certificate to all client systems.  Regardless of how you obtain the certificate, we recommend a **Digest** algorithm of **sha256** or higher for enhanced security.
Step 2	Configure an SSL/TLS Service Profile.	Select the **Certificate** you just obtained.  For enhanced security, we recommend that you set the **Min Version** (earliest allowed TLS version) to **TLSv1.1** for inbound management traffic. We also recommend that you use a different SSL/TLS Service Profile for each firewall or Panorama service instead of reusing this profile for all services.
Step 3	Apply the SSL/TLS Service Profile to inbound management traffic.	1. Select **Device > Setup > Management** and edit the General Settings. 2. Select the **SSL/TLS Service Profile** you just configured. 3. 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 **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 1024-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**.

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 firewall 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 firewall 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 firewalls or Panorama 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 master keys.

The following topics describe how to integrate an HSM with your firewall or Panorama:

- 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-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 IP address on the HSM client firewall must be a static IP address, not a dynamic address assigned by DHCP. HSM authenticates the firewall using the IP address before the HSM connection comes up. Operations on HSM would stop working if the IP address were to change during runtime.

The following topics describe how to set up connectivity to 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 firewall 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. Before starting the configuration, make sure you have created a partition for the Palo Alto Networks firewalls 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**. 2. Edit the Hardware Security Module Provider section and select **Safenet Luna SA** (SafeNet Network) as the **Provider Configured**. 3. Click **Add** and enter a **Module Name**. This can be any ASCII string up to 31 characters in length. 4. Enter the IPv4 address of the HSM module as the **Server Address**. If you are configuring a high availability HSM configuration, enter module names and IP addresses for the additional HSM devices. 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**. If two HSM servers are configured, you should configure high availability. Otherwise the second HSM server is not used. 6. Click **OK** and **Commit.**
Step 2	(Optional) Configure a service route to enable the firewall to connect to the HSM. 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**. 2. Select **Service Route Configuration** from the Services Features area. 3. Select **Customize** from the Service Route Configuration area. 4. Select the **IPv4** tab. 5. Select **HSM** from the **Service** column. 6. Select an interface to use for HSM from the **Source Interface** drop-down. 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. 7. Click **OK** and **Commit.**
Step 3	Configure the firewall to authenticate to the HSM.	1. Select **Device > Setup > HSM**. 2. Select **Setup Hardware Security Module** in the Hardware Security Operations area. 3. Select the HSM **Server Name** from the drop-down. 4. Enter the **Administrator Password** to authenticate the firewall to the HSM. 5. Click **OK.** The firewall attempts to perform an authentication with the HSM and displays a status message. 6. Click **OK.**

 © Palo Alto Networks, Inc.

Set up a Connectivity with a SafeNet Network HSM (Continued)		

Step 4	Register the firewall (the HSM client) with the HSM and assign it to a partition on the HSM.
	 
	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:
	`client delete -client <cl-name>`
	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.
		2. Register the firewall using the following command:
		`client register -c <cl-name> -ip <fw-ip-addr>`
		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. It must be a static IP address, not an address assigned by DHCP.
		3. Assign a partition to the firewall using the following command:
		`client assignpartition -c <cl-name> -p <partition-name>`
		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.

Step 5	Configure the firewall to connect to the HSM partition.	1. Select **Device > Setup > HSM**. 2. Click the **Refresh** icon. 3. Select the **Setup HSM Partition** in the Hardware Security Operations area. 4. Enter the **Partition Password** to authenticate the firewall to the partition on the HSM. 5. Click **OK**.
Step 6	(Optional) Configure an additional HSM for high availability (HA).	1. Repeat the previous steps to add an additional HSM for high availability (HA). This process adds a new HSM to the existing HA group. 2. If you remove an HSM from your configuration, repeat the previous step. This will remove the deleted HSM from the HA group.
Step 7	Verify connectivity with the HSM.	1. Select **Device > Setup > HSM**. 2. Check the **Status** of the HSM connection: Green—HSM is authenticated and connected. Red—HSM was not authenticated or network connectivity to the HSM is down. 3. View the following columns in Hardware Security Module Status area to determine authentication status: **Serial Number**—The serial number of the HSM partition if the HSM was successfully authenticated. **Partition**—The partition name on the HSM that was assigned on the firewall. **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 firewall is in an active/passive high availability configuration, you must manually perform one failover to configure and authenticate each HA peer individually to the HSM. After you perform this initial manual failover, no further user interaction is required for 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. 2. Select **Thales Nshield Connect** as the **Provider Configured**. 3. Click **Add** and enter a **Module Name**. This can be any ASCII string up to 31 characters in length. 4. Enter the IPv4 address as the **Server Address** of the HSM module. If you are configuring a high availability HSM configuration, enter module names and IP addresses for the additional HSM devices. 5. Enter the IPv4 address of the **Remote Filesystem Address**. 6. Click **OK** and **Commit**.
Step 2	(Optional) Configure a service route to enable the firewall to connect to the HSM. 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**. 2. Select **Service Route Configuration** from the Services Features area. 3. Select **Customize** from the Service Route Configuration area. 4. Select the **IPv4** tab. 5. Select **HSM** from the **Service** column. 6. Select an interface to use for HSM from the **Source Interface** drop-down. 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. 7. Click **OK** and **Commit.**
Step 3	Register the firewall (the HSM client) with the HSM server. 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. 2. On the unit front panel, use the right-hand navigation button to select **System > System configuration > Client config > New client**. 3. Enter the IP address of the firewall. It must be a static IP address, not an address assigned by DHCP. 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.
		2. Obtain the electronic serial number (ESN) and the hash of the K_{NETI} key. The K_{NETI} key authenticates the module to clients:
		`anonkneti <ip-address>`
		where `<ip-address>` is the IP address of the HSM.
		The following is an example:
		`anonkneti 192.0.2.1`
		`B1E2-2D4C-E6A2 5a2e5107e70d525615a903f6391ad72b1c03352c`
		In this example, **`B1E2-2D4C-E6A2`** is the ESM and 5a2e5107e70d525615a903f6391ad72b1c03352c is the hash of the K_{NETI} key.
		3. Use the following command from a superuser account to perform the remote filesystem setup:
		`rfs-setup --force <ip-address> <ESN> <hash-Kneti-key>`
		where `<ip-address>` is the IP address of the HSM,
		`<ESN>` is the electronic serial number (ESN) and
		`<hash-Kneti-key>` is the hash of the KNETI key.
		The following example uses the values obtained in this procedure:
		`rfs-setup --force <192.0.2.1> <B1E2-2D4C-E6A2>`
		`<5a2e5107e70d525615a903f6391ad72b1c03352c>`
		4. Use the following command to permit client submit on the Remote Filesystem:
		`rfs-setup --gang-client --write-noauth <FW-IPaddress>`
		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**. 2. Select **Setup Hardware Security Module** in the Hardware Security Operations area. 3. Click **OK.** The firewall attempts to perform an authentication with the HSM and displays a status message. 4. Click **OK.**
Step 6	Synchronize the firewall with the remote filesystem.	1. Select the **Device > Setup > HSM**. 2. Select **Synchronize with Remote Filesystem** in the Hardware Security Operations section.
Step 7	Verify that the firewall can connect to the HSM.	1. Select **Device > Setup > HSM**. 2. Check the Status indicator to verify that the firewall is connected to the HSM: Green—HSM is authenticated and connected. Red—HSM was not authenticated or network connectivity to the HSM is down. 3. View the following columns in Hardware Security Module Status section to determine authentication status. **Name**—The name of the HSM attempting to be authenticated. **IP address—**The IP address of the HSM that was assigned on the firewall. **Module State**—The current operating state of the HSM: **Authenticated** or **Not Authenticated**.

Encrypt a Master Key Using an HSM

A master key encrypts all private keys and passwords on the firewall and Panorama. 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 or Panorama 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 in a highly secure location that is separate from the firewall or Panorama for greater security.

The HSM encrypts the master key using a wrapping key. To maintain security, you must occasionally change (refresh) this wrapping key.



Firewalls configured in FIPS/CC mode do not support master key encryption using an HSM.

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 firewall, 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. **Life Time**—The number of days and hours after which the master key expires (range 1-730 days). **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**.

 © Palo Alto Networks, Inc.

Refresh the Master Key Encryption

As a best practice, periodically refresh the master key encryption by rotating the wrapping key that encrypts it. The frequency of the rotation depends on your application. The wrapping key resides on your HSM. The following command is the same for 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.

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 Forward Trust certificate that signs 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 the certificates 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 certificate and private key used in your decryption deployment.	For instructions on importing or generating a certificate and private key on the HSM, refer to your HSM documentation.
Step 2	(Thales nShield Connect only) Synchronize the key data from the Thales nShield remote file system to the firewall. Synchronization with the SafeNet Network HSM is automatic.	1. Access the firewall web interface and select **Device > Setup > HSM**. 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**. 2. Enter the **Certificate Name**. 3. **Browse** to the **Certificate File** on the HSM. 4. Select a **File Format**. 5. Select **Private Key resides on Hardware Security Module**. 6. Click **OK** and **Commit**.

Store Private Keys on an HSM (Continued)		
Step 4	(Forward Trust certificates only) Enable the certificate for use in SSL/TLS Forward Proxy.	1. Open the certificate you imported in Step 3 for editing. 2. Select **Forward Trust Certificate**. 3. Click **OK** and **Commit**.
Step 5	Verify that you successfully imported the certificate onto the firewall.	Locate the certificate you imported in Step 3 and check the icon in the Key column: • Lock icon—The private key for the certificate is on the HSM. • Error icon—The private key is not on 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. 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. 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. 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 deployment 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.

Palo Alto Networks firewalls support stateful active/passive or active/active high availability with session and configuration synchronization with a few exceptions:

- The PA-200 firewall supports HA Lite only.
- The VM-Series firewall in AWS supports active/passive HA only; if it is deployed with Amazon Elastic Load Balancing (ELB), it does not support HA (in this case ELB provides the failover capabilities).
- The VM-Series firewall in Microsoft Azure does not support 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

▲ Set Up Active/Active HA

▲ HA Firewall States

▲ Reference: HA Synchronization

HA Overview

You can set up two Palo Alto Networks 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)
- A critical chip or software component fails, known as packet path health monitoring.

You can use Panorama to manage HA firewalls. See Context Switch—Firewall or Panorama in the Panorama Administrator's Guide.

After you understand the HA Concepts, proceed to Set Up Active/Passive HA or Set Up Active/Active 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
- LACP and LLDP Pre-Negotiation for Active/Passive HA
- Floating IP Address and Virtual MAC Address
- ARP Load-Sharing
- Route-Based Redundancy
- HA Timers
- Session Owner
- Session Setup
- NAT in Active/Active HA Mode
- ECMP in Active/Active HA Mode

HA Modes

You can set up the firewalls for HA in one of 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 mode, 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.

  The PA-200 firewall supports HA Lite only.
  HA Lite is an active/passive deployment that provides configuration synchronization and some runtime data synchronization such as IPSec security associations. It does not support any session synchronization (HA2), and therefore does not offer stateful failover.

- Active/Active— Both firewalls in the pair are active and processing traffic and work synchronously to handle session setup and session ownership. Both firewalls individually maintain session tables and routing tables and synchronize to each other. Active/active HA is supported in virtual wire and Layer 3 deployments.

  In active/active HA mode, the firewall does not support DHCP client. Furthermore, only the active-primary firewall can function as a DHCP Relay. If the active-secondary firewall receives DHCP broadcast packets, it drops them.



An active/active configuration does not load-balance traffic. Although you can load-share by sending traffic to the peer, no load balancing occurs. Ways to load share sessions to both firewalls include using ECMP, multiple ISPs, and load balancers.

When deciding whether to use active/passive or active/active mode, consider the following differences:

- Active/passive mode has simplicity of design; it is significantly easier to troubleshoot routing and traffic flow issues in active/passive mode. Active/passive mode supports a Layer 2 deployment; active/active mode does not.
- Active/active mode requires advanced design concepts that can result in more complex networks. Depending on how you implement active/active HA, it might require additional configuration such as activating networking protocols on both firewalls, replicating NAT pools, and deploying floating IP addresses to provide proper failover. Because both firewalls are actively processing traffic, the firewalls use additional concepts of session owner and session setup to perform Layer 7 content inspection. Active/active mode is recommended if each firewall needs its own routing instances and you require full, real-time redundancy out of both firewalls all the time. Active/active mode has faster failover and can handle peak traffic flows better than active/passive mode because both firewalls are actively processing traffic.



In active/active mode, the HA pair can be used to temporarily process more traffic than what one firewall can normally handle. However, this should not be the norm because a failure of one firewall causes all traffic to be redirected to the remaining firewall in the HA pair.

Your design must allow the remaining firewall to process the maximum capacity of your traffic loads with content inspection enabled. If the design oversubscribes the capacity of the remaining firewall, high latency and/or application failure can occur.

For information on setting up your firewalls in active/passive mode, see Set Up Active/Passive HA. For information on setting up your firewalls in active/active mode, see Set Up Active/Active HA.

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.

 © Palo Alto Networks, Inc.

HA Links and Backup Links	Description
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.
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 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. PA-7000 Series firewalls synchronize sessions across the NPCs one-for-one. On PA-3000 Series, PA-4000 Series, and PA-5000 Series firewalls, you can configure aggregate interfaces as an HA3 link. The aggregate interfaces can also provide redundancy for the HA3 link; you cannot configure backup links for the HA3 link. On PA-7000 Series firewalls, the dedicated HSCI ports support the HA3 link. The firewall adds a proprietary packet header to packets traversing the HA3 link, so the MTU over this link must be greater than the maximum packet length forwarded.

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). PA-7000 Series firewalls synchronize sessions across the NPCs one-for-one.

The following table describes the SMC ports that are designed for HA connectivity:

HA Links and Backup Links	Ports on the SMC	Description
Control Link	HA1-A 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 port on the second firewall in the pair, or connect them through a switch or router. HA1 cannot be configured on NPC data ports or the MGT port.
Control Link Backup	HA1-B 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 port on the second firewall in the pair, or connect them through a switch or router. HA1 Backup cannot be configured on NPC data ports or the MGT port.
Data Link Data Link Backup	HSCI-A HSCI-B	The High Speed Chassis Interconnect (HSCI) ports are Layer 1 Quad Port SFP+ (QSFP+) interfaces used to connect two PA-7000 Series firewalls in an HA configuration. Each port is comprised of four 10 gigabit channels multiplexed for a combined speed of 40 gigabits. The traffic carried on the HSCI ports is raw layer-1, which is not routable or switchable; therefore the HSCI ports 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 provides full 80 gigabit transfer rates. In software, both ports (HSCI-A and HSCI-B) are treated as one HA interface. Palo Alto Networks recommends using the dedicated HSCI ports for the HA2 link. The HA3 link, required for packet forwarding in an active/active deployment, must use the HSCI port; the HA3 traffic cannot be configured on data ports. If the firewalls are deployed in: • an active/active configuration, the HA3 link must use the HSCI ports. The HA2 link and HA2 backup links can use the HSCI ports or data ports on the NPC. • an active/passive configuration, you can configure a data port on the NPC for the HA2 link or the HA2 backup link, if needed.

 © Palo Alto Networks, Inc.

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

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, for example, 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. A ping is sent every 1000 milliseconds and if there are three consecutive heartbeat losses, a failovers occurs. 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 (or to tentative state in active/active mode) 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 (or to tentative state in active/active mode) to indicate a failure of a monitored object.

In addition to the failover triggers listed above, a failover also occurs when the administrator suspends the firewall or when 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 monitor the critical components, such as the FPGA and CPUs. Additionally, general health checks occur on any platform causing failover.

LACP and LLDP Pre-Negotiation for Active/Passive HA

If a firewall uses LACP or LLDP, negotiation of those protocols upon failover prevents sub-second failover. However, you can enable an interface on a passive firewall to negotiate LACP and LLDP prior to failover. Thus, a firewall in Passive or Non-functional HA state can communicate with neighboring devices using LACP or LLDP. Such pre-negotiation speeds up failover.

The PA-3000 Series, PA-5000 Series, and PA-7000 Series firewalls support a pre-negotiation configuration depending on whether the Ethernet or AE interface is in a Layer 2, Layer 3, or virtual wire deployment. An HA passive firewall handles LACP and LLDP packets in one of two ways:

- **Active**—The firewall has LACP or LLDP configured on the interface and actively participates in LACP or LLDP pre-negotiation, respectively.
- **Passive**—LACP or LLDP is not configured on the interface and the firewall does not participate in the protocol, but allows the peers on either side of the firewall to pre-negotiate LACP or LLDP, respectively.

Pre-negotiation is not supported on subinterfaces or tunnel interfaces.

To configure LACP or LLDP pre-negotiation, see Step 14 of Configure Active/Passive HA.

Floating IP Address and Virtual MAC Address

In a Layer 3 deployment of HA active/active mode, you can assign floating IP addresses, which move from one HA firewall to the other if a link or firewall fails. The interface on the firewall that owns the floating IP address responds to ARP requests with a virtual MAC address.

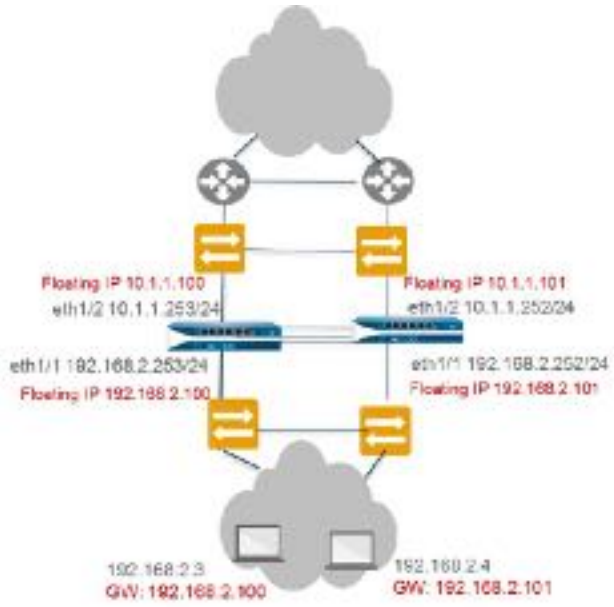
Floating IP addresses are recommended when you need functionality such as Virtual Router Redundancy Protocol (VRRP). Floating IP addresses can also be used to implement VPNs and source NAT, allowing for persistent connections when a firewall offering those services fails.

As shown in the figure below, each HA firewall interface has its own IP address and floating IP address. The interface IP address remains local to the firewall, but the floating IP address moves between the firewalls upon firewall failure. You configure the end hosts to use a floating IP address as its default gateway, allowing you to load balance traffic to the two HA peers. You can also use external load balancers to load balance traffic.

If a link or firewall fails or a path monitoring event causes a failover, the floating IP address and virtual MAC address move over to the functional firewall. (In the figure below, each firewall has two floating IP addresses and virtual MAC addresses; they all move over if the firewall fails.) The functioning firewall sends a gratuitous ARP to update the MAC tables of the connected switches to inform them of the change in floating IP address and MAC address ownership to redirect traffic to itself.

After the failed firewall recovers, by default the floating IP address and virtual MAC address move back to firewall with the Device ID [0 or 1] to which the floating IP address is bound. More specifically, after the failed firewall recovers, it comes on line. The currently active firewall determines that the firewall is back online and checks whether the floating IP address it is handling belongs natively to itself or the other firewall. If the floating IP address was originally bound to the other Device ID, the firewall automatically gives it back. (For an alternative to this default behavior, see Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall.)

 © Palo Alto Networks, Inc.

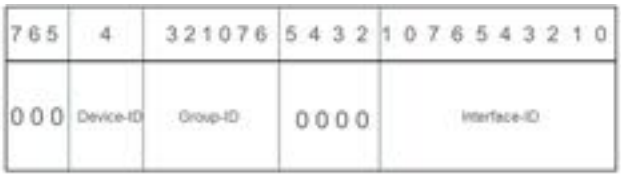


Each firewall in the HA pair creates a virtual MAC address for each of its interfaces that has a floating IP address or ARP Load-Sharing IP address.

The format of the virtual MAC address (on firewalls other than PA-7000 Series firewalls) is 00-1B-17-00-xx-yy, where 00-1B-17 is the vendor ID (of Palo Alto Networks in this case), 00 is fixed, xx indicates the Device ID and Group ID as shown in the following figure, and yy is the Interface ID:



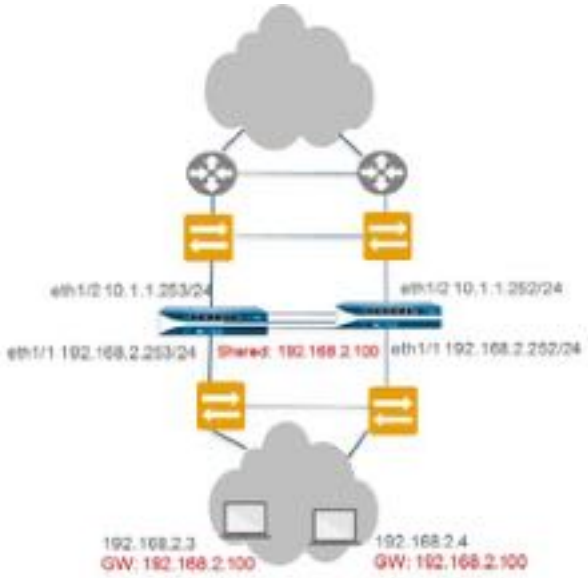
The format of the virtual MAC address on PA-7000 Series firewalls is 00-1B-17-xx-xx-xx, where 00-1B-17 is the vendor ID (of Palo Alto Networks in this case), and the next 24 bits indicate the Device ID, Group ID and Interface ID as follows:



When a new active firewall takes over, it sends gratuitous ARPs from each of its connected interfaces to inform the connected Layer 2 switches of the new location of the virtual MAC address. To configure floating IP addresses, see Use Case: Configure Active/Active HA with Floating IP Addresses.

ARP Load-Sharing

In a Layer 3 interface deployment and active/active HA configuration, ARP load-sharing allows the firewalls to share an IP address and provide gateway services. Use ARP load-sharing only when no Layer 3 device exists between the firewall and end hosts, that is, when end hosts use the firewall as their default gateway.



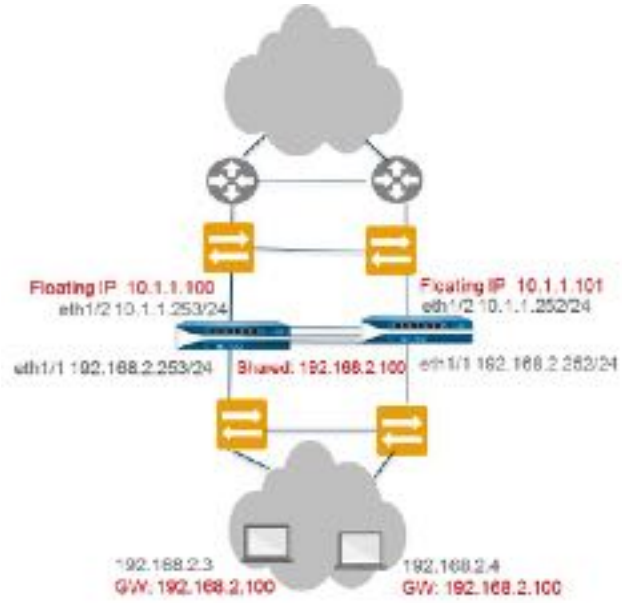
In such a scenario, all hosts are configured with a single gateway IP address. One of the firewalls responds to ARP requests for the gateway IP address with its virtual MAC address. Each firewall has a unique virtual MAC address generated for the shared IP address. The load-sharing algorithm that controls which firewall will respond to the ARP request is configurable; it is determined by computing the hash or modulo of the source IP address of the ARP request.

After the end host receives the ARP response from the gateway, it caches the MAC address and all traffic from the host is routed via the firewall that responded with the virtual MAC address for the lifetime of the ARP cache. The lifetime of the ARP cache depends on the end host operating system.

If a link or firewall fails, the floating IP address and virtual MAC address move over to the functional firewall. The functional firewall sends gratuitous ARPs to update the MAC table of the connected switches to redirect traffic from the failed firewall to itself. See Use Case: Configure Active/Active HA with ARP Load-Sharing.

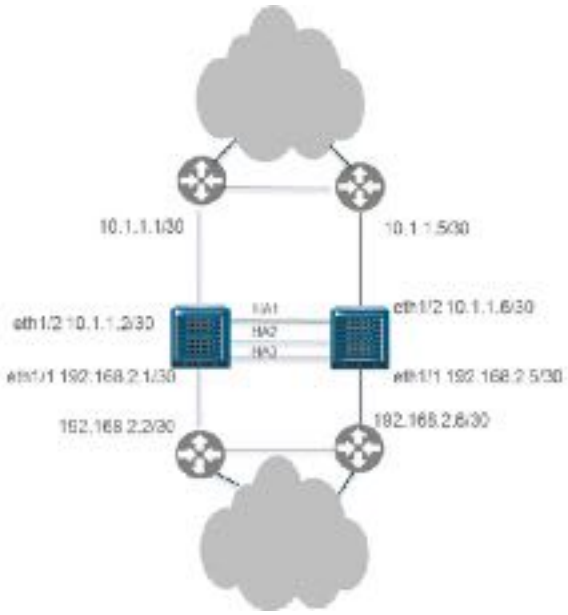
You can configure interfaces on the WAN side of the HA firewalls with floating IP addresses, and configure interfaces on the LAN side of the HA firewalls with a shared IP address for ARP load-sharing. For example, the figure below illustrates floating IP addresses for the upstream WAN edge routers and an ARP load-sharing address for the hosts on the LAN segment.

 © Palo Alto Networks, Inc.


Floating IP 10.1.1.100
eth1/2 10.1.1.253/24
Floating IP 10.1.1.101
eth1/2 10.1.1.252/24
eth1/1 192.168.2.253/24
Shared: 192.168.2.100
eth1/1 192.168.2.252/24
192.168.2.3
GW: 192.168.2.100
192.168.2.4
GW: 192.168.2.100

Route-Based Redundancy

In a Layer 3 interface deployment and active/active HA configuration, the firewalls are connected to routers, not switches. The firewalls use dynamic routing protocols to determine the best path (asymmetric route) and to load share between the HA pair. In such a scenario, no floating IP addresses are necessary. If a link, monitored path, or firewall fails, or if Bidirectional Forwarding Detection (BFD) detects a link failure, the routing protocol (RIP, OSPF, or BGP) handles the rerouting of traffic to the functioning firewall. You configure each firewall interface with a unique IP address. The IP addresses remain local to the firewall where they are configured; they do not move between devices when a firewall fails. See Use Case: Configure Active/Active HA with Route-Based Redundancy.



HA Timers

High availability (HA) timers facilitate a firewall 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.

 © Palo Alto Networks, Inc.

Timers	Description	PA-7000 Series PA-5000 Series PA-4000 Series PA-3000 Series VM-Series	PA-2000 Series PA-500 PA-200	Panorama Virtual Appliance 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 2000/1000 (only for VM-Series in AWS)	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-7000 Series PA-5000 Series PA-4000 Series PA-3000 Series VM-Series	PA-2000 Series PA-500 PA-200	Panorama Virtual Appliance Panorama M-Series
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
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.

Session Owner

In an HA active/active configuration, both firewalls are active simultaneously, which means packets can be distributed between them. Such distribution requires the firewalls to fulfill two functions: session ownership and session setup. Typically, each firewall of the pair performs one of these functions, thereby avoiding race conditions that can occur in asymmetrically routed environments.

You configure the session owner of sessions to be either the firewall that receives the First Packet of a new session from the end host or the firewall that is in active-primary state (the Primary device). If Primary device is configured, but the firewall that receives the first packet is not in active-primary state, the firewall forwards the packet to the peer firewall (the session owner) over the HA3 link.

The session owner performs all Layer 7 processing, such as App-ID, Content-ID, and threat scanning for the session. The session owner also generates all traffic logs for the session.

If the session owner fails, the peer firewall becomes the session owner. The existing sessions fail over to the functioning firewall and no Layer 7 processing is available for those sessions. When a firewall recovers from a failure, by default, all sessions it owned before the failure revert back to that original firewall; Layer 7 processing does not resume.

If you configure session ownership to be Primary device, the session setup defaults to Primary device also.

Palo Alto Networks recommends setting the Session Owner to First Packet and the Session Setup to IP Modulo unless otherwise indicated in a specific use case.

Setting Session Owner and Session Setup to Primary Device causes the active-primary firewall to perform all traffic processing. You might want to configure this for one of these reasons:

- You are troubleshooting and capturing logs and pcaps, so that packet processing is not split between the firewalls.
- You want to force the active/active HA pair to function like an active/passive HA pair. See Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall.

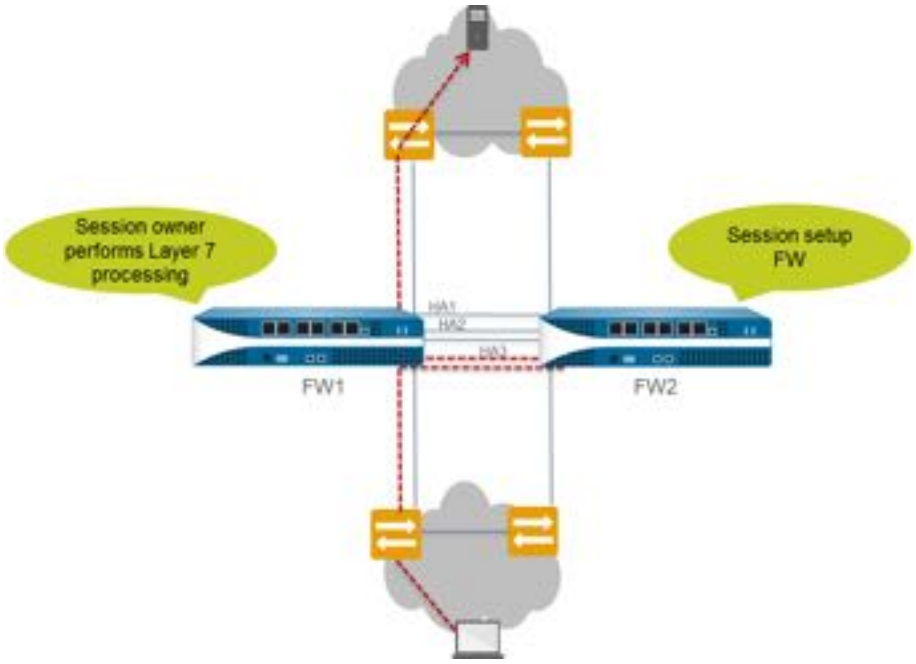
Session Setup

The session setup firewall performs the Layer2 through Layer 4 processing necessary to set up a new session. The session setup firewall also performs NAT using the NAT pool of the session owner. You determine the session setup firewall in an active/active configuration by selecting one of the following session setup load sharing options.

Session Setup Option	Description
IP Modulo	The firewall distributes the session setup load based on parity of the source IP address. This is a deterministic method of sharing the session setup.
IP Hash	The firewall uses a hash of the source and destination IP addresses to distribute session setup responsibilities.
Primary Device	The active-primary firewall always sets up the session; only one firewall performs all session setup responsibilities.
First Packet	The firewall that receives the first packet of a session performs session setup.

- If you want to load-share the session owner and session setup responsibilities, set session owner to First Packet and session setup to IP modulo. These are the recommended settings.
- If you want to do troubleshooting or capture logs or pcaps, or if you want an active/active HA pair to function like an active/passive HA pair, set both the session owner and session setup to Primary device so that the active-primary device performs all traffic processing. See Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall.

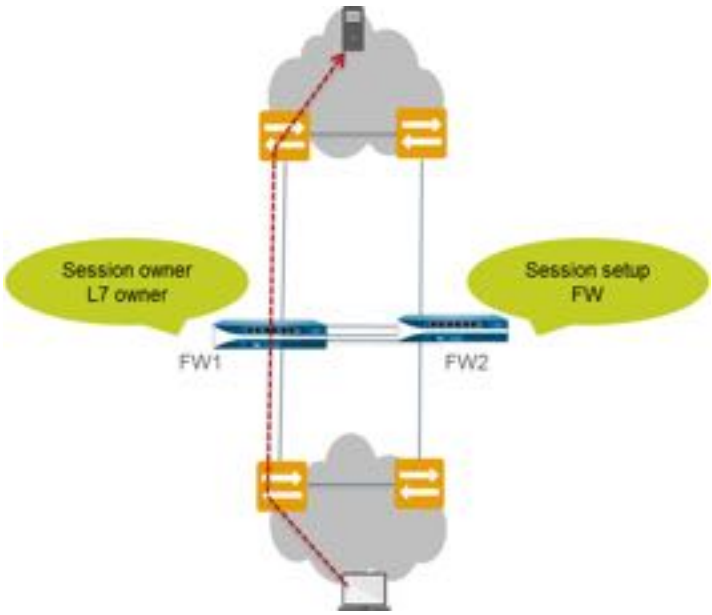
The firewall uses the HA3 link to send packets to its peer for session setup if necessary. The following figure and text describe the path of a packet that firewall FW1 receives for a new session. The red dotted lines indicate FW1 forwarding the packet to FW2 and FW2 forwarding the packet back to FW1 over the HA3 link.



- The end host sends a packet to FW1.
- FW1 examines the contents of the packet to match it to an existing session. If there is no session match, FW1 determines that it has received the first packet for a new session and therefore becomes the session owner (assuming **Session Owner Selection** is set to **First Packet**).
- FW1 uses the configured session setup load-sharing option to identify the session setup firewall. In this example, FW2 is configured to perform session setup.
- FW1 uses the HA3 link to send the first packet to FW2.
- FW2 sets up the session and returns the packet to FW1 for Layer 7 processing, if any.
- FW1 then forwards the packet out the egress interface to the destination.

 © Palo Alto Networks, Inc.

The following figure and text describe the path of a packet that matches an existing session:



1. The end host sends a packet to FW1.
2. FW1 examines the contents of the packet to match it to an existing session. If the session matches an existing session, FW1 processes the packet and sends the packet out the egress interface to the destination.

NAT in Active/Active HA Mode

In an active/active HA configuration:

- You must bind each Dynamic IP (DIP) NAT rule and Dynamic IP and Port (DIPP) NAT rule to either Device ID 0 or Device ID 1.
- You must bind each static NAT rule to either Device ID 0, Device ID 1, both Device IDs, or the firewall in active-primary state.

Thus, when one of the firewalls creates a new session, the Device ID **0** or Device ID **1** binding determines which NAT rules match the firewall. The device binding must include the session owner firewall to produce a match.

The session setup firewall performs the NAT policy match, but the NAT rules are evaluated based on the session owner. That is, the session is translated according to NAT rules that are bound to the session owner firewall. While performing NAT policy matching, a firewall skips all NAT rules that are not bound to the session owner firewall.

For example, suppose the firewall with Device ID 1 is the session owner and session setup firewall. When the firewall with Device ID 1 tries to match a session to a NAT rule, it skips all rules bound to Device ID 0. The firewall performs the NAT translation only if the session owner and the Device ID in the NAT rule match.

You will typically create device-specific NAT rules when the peer firewalls use different IP addresses for translation.

If one of the peer firewalls fails, the active firewall continues to process traffic for synchronized sessions from the failed firewall, including NAT traffic. In a source NAT configuration, when one firewall fails:

- The floating IP address that is used as the Translated IP address of the NAT rule transfers to the surviving firewall. Hence, the existing sessions that fail over will still use this IP address.
- All new sessions will use the device-specific NAT rules that the surviving firewall naturally owns. That is, the surviving firewall translates new sessions using only the NAT rules that match its Device ID; it ignores any NAT rules bound to the failed Device ID.

If you want the firewalls to perform dynamic NAT using the same IP address simultaneously, a best practice is to create a duplicate NAT rule that is bound to the peer firewall also. The result is two NAT rules with the same translation IP addresses, one bound to Device ID 0 and one bound to Device ID 1. Thus, the configuration allows the current firewall to perform new session setup and perform NAT policy matching for NAT rules that are bound to its Device ID. Without the duplicate NAT rule, the firewall will not find its own device-specific rules and will skip all NAT rules that are not bound to its Device ID when it attempts to match a NAT policy.

For examples of active/active HA with NAT, see:

- Use Case: Configure Active/Active HA with Source DIPP NAT Using Floating IP Addresses
- Use Case: Configure Separate Source NAT IP Address Pools for Active/Active HA Firewalls
- Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT
- Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT in Layer 3

ECMP in Active/Active HA Mode

When an active/active HA peer fails, its sessions transfer to the new active-primary firewall, which tries to use the same egress interface that the failed firewall was using. If the firewall finds that interface among the ECMP paths, the transferred sessions will take the same egress interface and path. This behavior occurs 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 firewall select a new ECMP path.

If you did not configure the same interfaces on the active/active peers, upon failover the active-primary firewall selects the next best path from the FIB table. Consequently, the existing sessions might not be distributed according to the ECMP algorithm.

 © 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 should be running the same PAN-OS version and must each be up-to-date on the application, URL, and threat databases.
- The same multi virtual system capability—Both firewalls must have **Multi Virtual System Capability** either enabled or not enabled. When enabled, each firewall requires its own multiple virtual systems licenses.
- 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.



As a best practice, 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, Reset the Firewall to Factory Default Settings on the new firewall. This ensures 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 configuration.

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, see Reference: HA Synchronization.

The following checklist details the settings that you must configure identically on both firewalls:

- You must enable HA on both firewalls.
- You must configure the same Group ID value on both firewalls. The firewall uses the Group ID value to create a virtual MAC address for all the configured interfaces. See Floating IP Address and Virtual MAC Address for information about virtual MAC addresses. When a new active firewall takes over, it sends Gratuitous ARP messages from each of its connected interfaces to inform the connected Layer 2 switches of the virtual MAC address' new location.
- If you are using in-band ports as HA links, you must set the interfaces for the HA1 and HA2 links to type HA.
- Set the HA Mode to Active Passive on both firewalls.
- If required, enable preemption on both firewalls. The device priority value, however, must not be identical.
- If required, configure encryption on the HA1 link (for communication between the HA peers) on both firewalls.
- 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:



HA functionality (HA1 and HA1 backup) is not supported on the management interface if it's configured for DHCP addressing (**IP Type** set to **DHCP Client**), except for AWS.

  - HA1: Dedicated HA1 port
    HA1 Backup: In-band port
    Recommendation: Enable Heartbeat Backup
  - HA1: Dedicated HA1 port
    HA1 Backup: Management port
    Recommendation: Do not enable Heartbeat Backup
  - HA1: In-band port
    HA1 Backup: In-band port
    Recommendation: Enable Heartbeat Backup
  - HA1: Management port
    HA1 Backup: In-band port
    Recommendation: Do not enable Heartbeat Backup

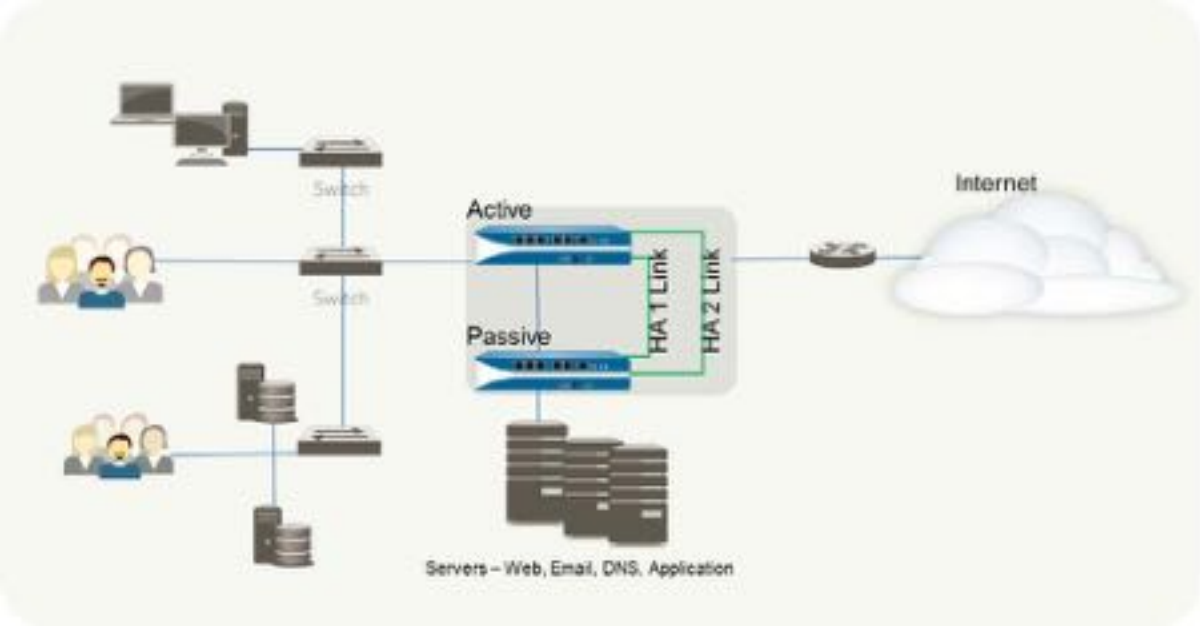
 © Palo Alto Networks, Inc.

The following table lists the HA settings that you must configure independently on each firewall. See Reference: HA Synchronization for more information about other configuration settings are not automatically synchronized between peers.

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 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. 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. 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 the firewalls have the same device priority value, they use the MAC address of their HA1 as the tie-breaker.	If PeerB is passive, set the device priority value to a number larger than the setting 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. 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.



To configure an active/passive HA pair, first complete the following workflow on the first firewall and then repeat the steps on the second firewall.

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. • 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. Use the management port for the HA1 link and ensure that the management ports can connect to each other across your network.
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 edit the Management Interface Settings. 2. Select **Ping** as a service that is permitted on the interface.
Step 3	If the firewall does not have dedicated HA ports, set up the data ports to function as HA ports. For firewalls with dedicated HA ports continue to the next step.	1. Select **Network > Interfaces**. 2. Confirm that the link is up on the ports that you want to use. 3. Select the interface and set **Interface Type** to **HA**. 4. Set the **Link Speed** and **Link Duplex** settings, as appropriate.

 © Palo Alto Networks, Inc.

Connect and Configure the Firewalls (Continued)		
Step 4	Set the HA mode and group ID.	1. Select **Device > High Availability > General** and edit the Setup section. 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 multiple HA pairs that share the same broadcast domain you must set a unique Group ID for each pair. 3. Set the mode to **Active Passive**.
Step 5	Set up the control link connection. This example shows an in-band port that is set to interface type HA. 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. 2. Select the **Port** that you have cabled for use as the HA1 link. 3. Set the **IPv4/IPv6 Address** and **Netmask**. 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
Step 6	(Optional) Enable encryption for the control link connection. 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. a. Select **Device > Certificate Management > Certificates**. b. Select **Export HA key**. Save the HA key to a network location that the peer can access. 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. 2. Select **Device > High Availability > General**, edit the Control Link (HA1) section. 3. Select **Encryption Enabled**.
Step 7	Set up the backup control link connection.	1. In **Device > High Availability > General**, edit the Control Link (HA1 Backup) section. 2. Select the HA1 backup interface and set the **IPv4/IPv6 Address** and **Netmask**.

Connect and Configure the Firewalls (Continued)		
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. 2. Select the **Port** to use for the data link connection. 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. 4. If you use IP or UDP as the transport method, enter the **IPv4/IPv6 Address** and **Netmask**. 5. Verify that **Enable Session Synchronization** is selected. 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.  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. 7. Edit the **Data Link (HA2 Backup)** section, select the interface, and add the **IPv4/IPv6 Address** and **Netmask**.
Step 9	Enable heartbeat backup if your control link uses a dedicated HA port or an in-band port. 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. 2. Select **Heartbeat Backup**. 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.  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.

 © Palo Alto Networks, Inc.

Connect and Configure the Firewalls (Continued)	
Step 10 Set the device priority and enable preemption. 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. 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. 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. 3. Select **Preemptive**. You must enable preemptive on both the active firewall and the passive firewall.
Step 11 (Optional) Modify the HA Timers. 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. 2. Select the **Aggressive** profile for triggering failover faster; select **Advanced** to define custom values for triggering failover in your set up. 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 12 (Optional, only configured on the passive firewall) Modify the link status of the HA ports on the passive firewall. 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. To enable the link status on the passive firewall to stay up and reflect the cabling status on the physical interface: 1. In **Device > High Availability > General**, edit the Active Passive Settings. 2. Set the **Passive Link State** to **Auto**. The auto option decreases the amount of time it takes for the passive firewall to take over when a failover occurs. Although the interface displays green (as cabled and up) it continues to discard all traffic until a failover is triggered. 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.

Connect and Configure the Firewalls (Continued)	
Step 13 Enable HA.	1. Select **Device > High Availability > General** and edit the Setup section. 2. Select **Enable HA**. 3. Select **Enable Config Sync**. This setting enables the synchronization of the configuration settings between the active and the passive firewall. 4. Enter the IP address assigned to the control link of the peer in **Peer HA1 IP Address**. 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. 5. Enter the **Backup HA1 IP Address**.
Step 14 (Optional) Enable LACP and LLDP Pre-Negotiation for Active/Passive HA for faster failover if your network uses LACP or LLDP. Enable LACP and LLDP before configuring HA pre-negotiation for the protocol if you want pre-negotiation to function in active mode.	1. Ensure that in Step 12 you set the link state to **Auto**. 2. Select **Network > Interfaces > Ethernet**. 3. To enable LACP active pre-negotiation: a. Select an AE interface in a Layer 2 or Layer 3 deployment. b. Select the **LACP** tab. c. Select **Enable in HA Passive State**. d. Click **OK**. You cannot also select **Same System MAC Address for Active-Passive HA** because pre-negotiation requires unique interface MAC addresses on the active and passive firewalls. 4. To enable LACP passive pre-negotiation: a. Select an Ethernet interface in a virtual wire deployment. b. Select the **Advanced** tab. c. Select the **LACP** tab. d. Select **Enable in HA Passive State**. e. Click **OK**. 5. To enable LLDP active pre-negotiation: a. Select an Ethernet interface in a Layer 2, Layer 3, or virtual wire deployment. b. Select the **Advanced** tab. c. Select the **LLDP** tab. d. Select **Enable in HA Passive State**. e. Click **OK**. If you want to allow LLDP passive pre-negotiation for a virtual wire deployment, perform Step 5 but do not enable LLDP itself.
Step 15 Save your configuration changes.	Click **Commit**.

 © 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. 2. On the active firewall, click the **Sync to peer** link. 3. Confirm that the firewalls are paired and synced, as shown as follows: • 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**.

Define HA Failover Conditions

Configure the Failover Triggers	
Step 1 To configure link monitoring, define the interfaces you want to monitor. A change in the link state of these interfaces will trigger a failover.	1. Select **Device > High Availability > Link and Path Monitoring** and **Add** a Link Group. 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. By default, the firewall will trigger a failover when any monitored link fails.	1. Select the **Link Monitoring** section. 2. Set the **Failure Condition** to **All**. 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. 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.
Step 4 (Optional) Modify the failure condition for all Path Groups configured on the firewall. By default, the firewall will trigger a failover when any monitored path fails.	Set the **Failure Condition** to **All**. 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 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.
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.
Step 3	Restore the suspended firewall to a functional state. Wait for a couple of 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. 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.

 © Palo Alto Networks, Inc.

Set Up Active/Active HA

- Prerequisites for Active/Active HA
- Configure Active/Active HA
- Determine Your Active/Active Use Case

Prerequisites for Active/Active HA

To set up active/active HA on your firewalls, you need a pair of firewalls that meet the following requirements:

- The same model—The firewalls in the pair must be of the same hardware model.
- The same PAN-OS version—The firewalls must be running the same PAN-OS version and must each be up-to-date on the application, URL, and threat databases.
- The same multi virtual system capability—Both firewalls must have **Multi Virtual System Capability** either enabled or not enabled. When enabled, each firewall requires its own multiple virtual systems licenses.
- 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.
  - The HA interfaces must be configured with static IP addresses only, not IP addresses obtained from DHCP (except AWS can use DHCP addresses). Determine the IP address for the HA1 (control) connection between the HA peers. The HA1 IP address for the 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.
  - Each firewall needs a dedicated interface for the HA3 link. PA-7000 Series firewalls use the HSCI port. On the remaining platforms, you can configure aggregate interfaces as the HA3 link for redundancy.
- 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. You will also have to configure local IP addresses.

Configure Active/Active HA

The following procedure describes the basic workflow for configuring your firewalls in an active/active configuration. However, before you begin, Determine Your Active/Active Use Case for configuration examples more tailored to your specific network environment.

To configure active/active, first complete the following steps on one peer and then complete them on the second peer, ensuring that you set the Device ID to different values (0 or 1) on each peer.

Configure Active/Active HA		

Step 1	Connect the HA ports to set up a physical connection between the firewalls.
	  For each use case, the firewalls could be any hardware platform; choose the HA3 step that corresponds with your platform.
		• 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.
		• 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.
		Use the management port for the HA1 link and ensure that the management ports can connect to each other across your network.
		• For HA3:
		  • On PA-7000 Series firewalls, connect the High Speed Chassis Interconnect (HSCI-A) on the first chassis to the HSCI-A on the second chassis, and the HSCI-B on the first chassis to the HSCI-B on the second chassis.
		  • On any other hardware platform, use dataplane interfaces for HA3.

Step 2	Enable ping on the management port. Enabling ping allows the management port to exchange heartbeat backup information.	1. In **Device > Setup > Management**, edit Management Interface Settings. 2. Select **Ping** as a service that is permitted on the interface.
Step 3	If the firewall does not have dedicated HA ports, set up the data ports to function as HA ports. For firewalls with dedicated HA ports continue to the next step.	1. Select **Network > Interfaces**. 2. Confirm that the link is up on the ports that you want to use. 3. Select the interface and set **Interface Type** to **HA**. 4. Set the **Link Speed** and **Link Duplex** settings, as appropriate.
Step 4	Enable active/active HA and set the group ID.	1. In **Device > High Availability > General**, edit Setup. 2. Select **Enable HA**. 3. Enter a **Group ID**, which must be the same for both firewalls. The firewall uses the Group ID to calculate the virtual MAC address (range is 1-63). 4. (Optional) Enter a **Description**. 5. For **Mode**, select **Active Active**.

 © Palo Alto Networks, Inc.

Configure Active/Active HA (Continued)		
Step 5	Set the Device ID, enable synchronization, and identify the control link on the peer firewall	1. In **Device > High Availability > General**, edit Setup. 2. Select **Device ID** as follows: • When configuring the first peer, set the **Device ID** to **0**. • When configuring the second peer, set the **Device ID** to **1**. 3. Select **Enable Config Sync**. This setting is required to synchronize the two firewall configurations (enabled by default). 4. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall. 5. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall. 6. Click **OK**.
Step 6	Determine whether or not the firewall with the lower Device ID preempts the active-primary firewall upon recovery from a failure.	1. In **Device > High Availability > General**, edit Election Settings. 2. Select **Preemptive** to cause the firewall with the lower Device ID to automatically resume active-primary operation after either firewall recovers from a failure. Both firewalls must have **Preemptive** selected for preemption to occur. Leave **Preemptive** unselected if you want the active-primary role to remain with the current firewall until you manually make the recovered firewall the active-primary firewall.
Step 7	Enable heartbeat backup if your control link uses a dedicated HA port or an in-band port. You need not enable heartbeat backup if you are using the management port for the control link.	1. In **Device > High Availability > General**, edit Election Settings. 2. Select **Heartbeat Backup**. 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.  Enabling heartbeat backup 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 the other is down and attempts to start services that are running, thereby causing a split brain. Enabling heartbeat backup prevents split brain because redundant heartbeats and hello messages are transmitted over the management port.
Step 8	(Optional) Modify the **HA Timers**. 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 Election Settings. 2. Select **Aggressive** to trigger faster failover. Select **Advanced** to define custom values for triggering failover in your setup.  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.

Configure Active/Active HA (Continued)	
Step 9 Set up the control link connection. This example uses an in-band port that is set to interface type HA. 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 Control Link (HA1). 2. Select the **Port** that you have cabled for use as the HA1 link. 3. Set the **IPv4/IPv6 Address** and **Netmask**. 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.
Step 10 (Optional) Enable encryption for the control link connection. 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. a. Select **Device > Certificate Management > Certificates**. b. Select **Export HA key**. Save the HA key to a network location that the peer can access. 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. 2. In **Device > High Availability > General**, edit the Control Link (HA1). 3. Select **Encryption Enabled**.
Step 11 Set up the backup control link connection.	1. In **Device > High Availability > General**, edit Control Link (HA1 Backup). 2. Select the HA1 backup interface and set the **IPv4/IPv6 Address** and **Netmask**.
Step 12 Set up the data link connection (HA2) and the backup HA2 connection between the firewalls.	1. In **Device > High Availability > General**, edit Data Link (HA2). 2. Select the **Port** to use for the data link connection. 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. 4. If you use IP or UDP as the transport method, enter the **IPv4/IPv6 Address** and **Netmask**. 5. Verify that **Enable Session Synchronization** is selected. 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.  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. 7. Edit the **Data Link (HA2 Backup)** section, select the interface, and add the **IPv4/IPv6 Address** and **Netmask**. 8. Click **OK**.

 © Palo Alto Networks, Inc.

Configure Active/Active HA (Continued)	
Step 13 Configure the HA3 link for packet forwarding.	1. In **Device > High Availability > Active/Active Config**, edit Packet Forwarding. 2. For **HA3 Interface**, select the interface you want to use to forward packets between active/active HA peers. It must be a dedicated interface capable of Layer 2 transport and set to **Interface Type HA**. 3. Select **VR Sync** to force synchronization of all virtual routers configured on the HA peers. Select when the virtual router is not configured for dynamic routing protocols. Both peers must be connected to the same next-hop router through a switched network and must use static routing only. 4. Select **QoS Sync** to synchronize the QoS profile selection on all physical interfaces. Select when both peers 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.
Step 14 (Optional) Modify the Tentative Hold time.	1. In **Device > High Availability > Active/Active Config**, edit Packet Forwarding. 2. For **Tentative Hold Time (sec)**, enter the number of seconds that a firewall stays in Tentative state after it fails (range is 10-600, default is 60).
Step 15 Configure Session Owner and Session Setup.	1. In **Device > High Availability > Active/Active Config**, edit Packet Forwarding. 2. For **Session Owner Selection**, select one of the following: • **First Packet**—The firewall that receives the first packet of a new session is the session owner (recommended setting). This setting minimizes traffic across HA3 and load shares traffic across peers. • **Primary Device**—The firewall that is in active-primary state is the session owner. 3. For **Session Setup**, select one of the following: • **IP Modulo**— Distributes session setup load based on parity of the source IP address (recommended setting). • **Primary Device**—The active-primary firewall sets up all sessions. • **First Packet**—The firewall that receives the first packet of a new session performs session setup. • **IP Hash**—The firewall uses a hash of either the source IP address or a combination of the source and destination IP addresses to distribute session setup responsibilities. 4. Click **OK**.

Configure Active/Active HA (Continued)	
Step 16 Configure an HA virtual address. You need a virtual address to use a Floating IP Address and Virtual MAC Address or ARP Load-Sharing.	1. In **Device > High Availability > Active/Active Config**, **Add** a Virtual Address. 2. Enter or select an **Interface**. 3. Select the **IPv4** or **IPv6** tab and click **Add**. 4. Enter an **IPv4 Address** or **IPv6 Address**. 5. For **Type**: • Select **Floating** to configure the virtual IP address to be a floating IP address. • Select **ARP Load Sharing** to configure the virtual IP address to be a shared IP address and skip to Configure ARP Load-Sharing.
Step 17 Configure the floating IP address.	1. Do not select **Floating IP bound to the Active-Primary device** unless you want the active/active HA pair to behave like an active/passive HA pair. 2. For **Device 0 Priority** and **Device 1 Priority**, enter a priority for the firewall configured with Device ID 0 and Device ID 1, respectively. The relative priorities determine which peer owns the floating IP address you just configured (range is 0-255). The firewall with the lowest priority value (highest priority) owns the floating IP address. 3. Select **Failover address if link state is down** to cause the firewall to use the failover address when the link state on the interface is down. 4. Click **OK**.
Step 18 Configure ARP Load-Sharing. The device selection algorithm determines which HA firewall responds to the ARP requests to provide load sharing.	1. For **Device Selection Algorithm**, select one of the following: • **IP Modulo**—The firewall that will respond to ARP requests is based on the parity of the ARP requester's IP address. • **IP Hash**—The firewall that will respond to ARP requests is based on a hash of the ARP requester's IP address. 2. Click **OK**.
Step 19 Enable jumbo frames on firewalls other than PA-7000 Series firewalls. Switch ports that connect the HA3 link must support jumbo frames to handle the overhead associated with the MAC-in-MAC encapsulation on the HA3 link. The jumbo frame packet size on the firewall must match the setting on the switch.	1. Select **Device > Setup > Session**. 2. In the Session Settings section, select **Enable Jumbo Frames**. 3. Click **OK**. 4. Repeat on any intermediary networking devices.
Step 20 Define HA Failover Conditions.	
Step 21 Save the configuration.	Click **Commit**.
Step 22 Reboot the firewall after changing the jumbo frame configuration.	1. Select **Device > Setup > Operations**. 2. Click **Reboot Device**.

 © Palo Alto Networks, Inc.

Determine Your Active/Active Use Case

Determine which type of use case you have and then select the corresponding procedure to configure active/active HA.

If you are using Route-Based Redundancy, Floating IP Address and Virtual MAC Address, or ARP Load-Sharing, select the corresponding procedure:

- Use Case: Configure Active/Active HA with Route-Based Redundancy
- Use Case: Configure Active/Active HA with Floating IP Addresses
- Use Case: Configure Active/Active HA with ARP Load-Sharing

If you want a Layer 3 active/active HA deployment that behaves like an active/passive deployment, select the following procedure:

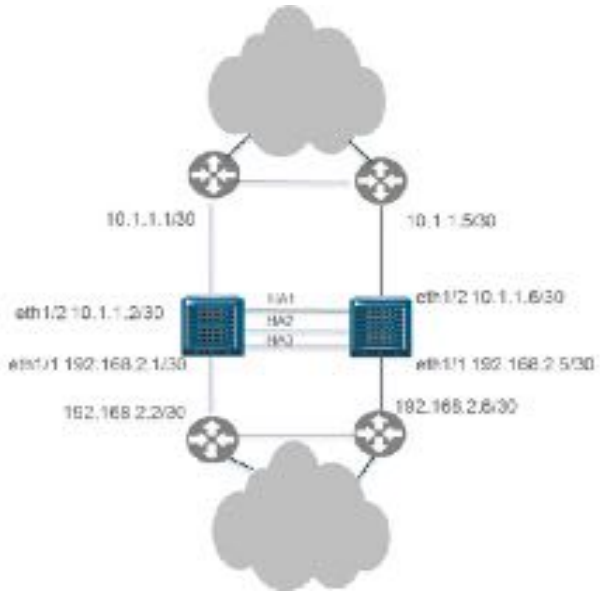
- Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall

If you are configuring NAT in Active/Active HA Mode, see the following procedures:

- Use Case: Configure Active/Active HA with Source DIPP NAT Using Floating IP Addresses
- Use Case: Configure Separate Source NAT IP Address Pools for Active/Active HA Firewalls
- Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT
- Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT in Layer 3

Use Case: Configure Active/Active HA with Route-Based Redundancy

The following Layer 3 topology illustrates two PA-7050 firewalls in an active/active HA environment that use Route-Based Redundancy. The firewalls belong to an OSPF area. When a link or firewall fails, OSPF handles the redundancy by redirecting traffic to the functioning firewall.

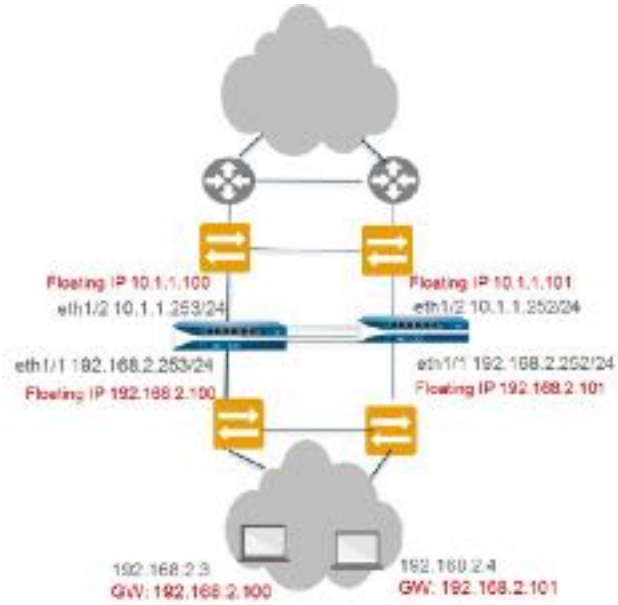


Configure Active/Active HA with Route-Based Redundancy		
Step 1	Configure Active/Active HA.	Perform Step 1 through Step 15.
Step 2	Configure OSPF.	See OSPF.
Step 3	Define HA failover conditions.	Define HA Failover Conditions.
Step 4	Save the configuration.	Click **Commit**.
Step 5	Configure the peer firewall in the same way, except in Step 5, if you selected Device ID **0** for the first firewall, select Device ID **1** for the peer firewall.	

 © Palo Alto Networks, Inc.

Use Case: Configure Active/Active HA with Floating IP Addresses

In this Layer 3 interface example, the HA firewalls connect to switches and use floating IP addresses to handle link or firewall failures. The end hosts are each configured with a gateway, which is the floating IP address of one of the HA firewalls. See Floating IP Address and Virtual MAC Address.

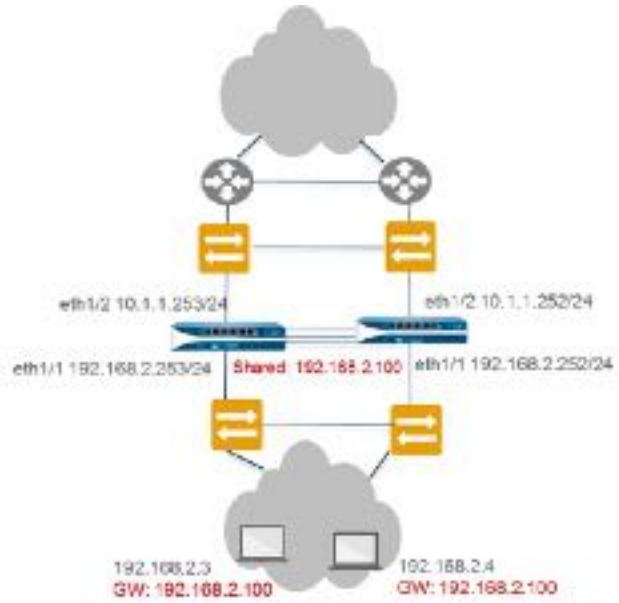


Configure Active/Active HA with Floating IP Addresses		
Step 1	Configure Active/Active HA.	Perform Step 1 through Step 15.
Step 2	Configure an HA virtual address. You need a virtual address to use a Floating IP Address and Virtual MAC Address.	1. In **Device > High Availability > Active/Active Config**, **Add** a Virtual Address. 2. Enter or select an **Interface**. 3. Select the **IPv4** or **IPv6** tab and click **Add**. 4. Enter an **IPv4 Address** or **IPv6 Address**. 5. For **Type**, select **Floating** to configure the virtual IP address to be a floating IP address.
Step 3	Configure the floating IP address.	1. Do not select **Floating IP bound to the Active-Primary device**. 2. For **Device 0 Priority** and **Device 1 Priority**, enter a priority for the firewall configured with Device ID 0 and Device ID 1, respectively. The relative priorities determine which peer owns the floating IP address you just configured (range is 0-255). The firewall with the lowest priority value (highest priority) owns the floating IP address. 3. Select **Failover address if link state is down** to cause the firewall to use the failover address when the link state on the interface is down. 4. Click **OK**.

Configure Active/Active HA with Floating IP Addresses (Continued)		
Step 4	Enable jumbo frames on firewalls other than PA-7000 Series firewalls.	Perform Step 19 of Configure Active/Active HA.
Step 5	Define HA failover conditions.	Define HA Failover Conditions.
Step 6	Save the configuration.	Click **Commit**.
Step 7	Configure the peer firewall in the same way, except selecting a different Device ID. For example, if you selected Device ID **0** for the first firewall, select Device ID **1** for the peer firewall.	

Use Case: Configure Active/Active HA with ARP Load-Sharing

In this example, hosts in a Layer 3 deployment need gateway services from the HA firewalls. The firewalls are configured with a single shared IP address, which allows ARP Load-Sharing. The end hosts are configured with the same gateway, which is the shared IP address of the HA firewalls.



Configure Active/Active HA with ARP Load-Sharing		
Step 1	Perform Step 1 through Step 15 of Configure Active/Active HA.	

 © Palo Alto Networks, Inc.

Configure Active/Active HA with ARP Load-Sharing (Continued)		
Step 2	Configure an HA virtual address. The virtual address is the shared IP address that allows ARP Load-Sharing.	1. Select **Device > High Availability > Active/Active Config > Virtual Address** and click **Add**. 2. Enter or select an **Interface**. 3. Select the **IPv4** or **IPv6** tab and click **Add**. 4. Enter an **IPv4 Address** or **IPv6 Address**. 5. For **Type**, select **ARP Load Sharing**, which allows both peers to use the virtual IP address for ARP Load-Sharing.
Step 3	Configure ARP Load-Sharing. The device selection algorithm determines which HA firewall responds to the ARP requests to provide load sharing.	1. For **Device Selection Algorithm**, select one of the following: • **IP Modulo**—The firewall that will respond to ARP requests is based on the parity of the ARP requester's IP address. • **IP Hash**—The firewall that will respond to ARP requests is based on a hash of the ARP requester's IP address. 2. Click **OK**.
Step 4	Enable jumbo frames on firewalls other than PA-7000 Series firewalls.	Enable jumbo frames on firewalls other than PA-7000 Series firewalls.
Step 5	Define HA failover conditions.	Define HA Failover Conditions.
Step 6	Save the configuration.	Click **Commit**.
Step 7	Configure the peer firewall in the same way, except selecting a different Device ID. For example, if you selected Device ID **0** for the first firewall, select Device ID **1** for the peer firewall.	

Use Case: Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall

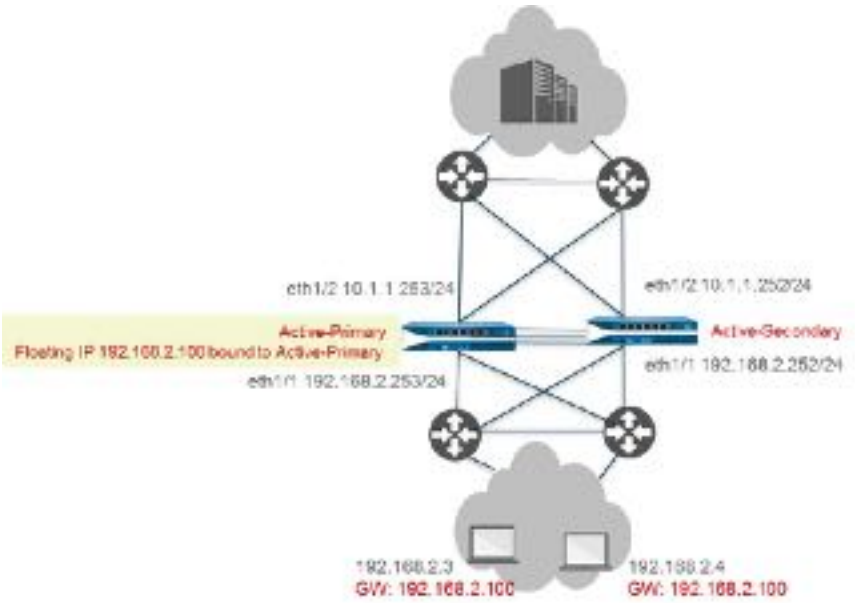
In mission-critical data centers, you may want both Layer 3 HA firewalls to participate in path monitoring so that they can detect path failures upstream from both firewalls. Additionally, you prefer to control if and when the floating IP address returns to the recovered firewall after it comes back up, rather than the floating IP address returning to the device ID to which it is bound. (That default behavior is described in Floating IP Address and Virtual MAC Address.)

In this use case, you control when the floating IP address and therefore the active-primary role move back to a recovered HA peer. The active/active HA firewalls share a single floating IP address that you bind to whichever firewall is in the active-primary state. With only one floating IP address, network traffic flows predominantly to a single firewall, so this active/active deployment functions like an active/passive deployment.

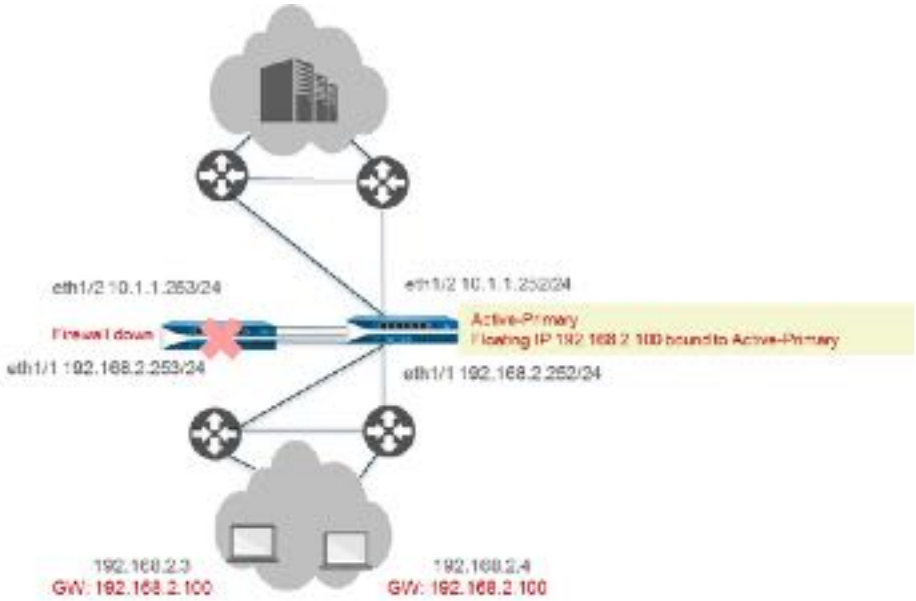
In this use case, Cisco Nexus 7010 switches with virtual PortChannels (vPCs) operating in Layer 3 connect to the firewalls. You must configure the Layer 3 switches (router peers) north and south of the firewalls with a route preference to the floating IP address. That is, you must design your network so the route tables of the router peers have the best path to the floating IP address. This example uses static routes with the proper

metrics so that the route to the floating IP address uses a lower metric (the route to the floating IP address is preferred) and receives the traffic. An alternative to using static routes would be to design the network to redistribute the floating IP address into the OSPF routing protocol (if you are using OSPF).

The following topology illustrates the floating IP address bound to the active-primary firewall, which is initially Peer A, the firewall on the left.



Upon a failover, when the active-primary firewall (Peer A) goes down and the active-secondary firewall (Peer B) takes over as the active-primary peer, the floating IP address moves to Peer B (shown in the following figure). Peer B remains the active-primary firewall and traffic continues to go to Peer B, even when Peer A recovers and becomes the active-secondary firewall. You decide if and when to make Peer A the active-primary firewall again.



 © Palo Alto Networks, Inc.

Binding the floating IP address to the active-primary firewall provides you with more control over how the firewalls determine floating IP address ownership as they move between various HA Firewall States. The following advantages result:

- You can have an active/active HA configuration for path monitoring out of both firewalls, but have the firewalls function like an active/passive HA configuration because traffic directed to the floating IP address always goes to the active-primary firewall.

When you disable preemption on both firewalls, you have the following additional benefits:

- The floating IP address does not move back and forth between HA firewalls if the active-secondary firewall flaps up and down.
- You can review the functionality of the recovered firewall and the adjacent components before manually directing traffic to it again, which you can do at a convenient down time.
- You have control over which firewall owns the floating IP address so that you keep all flows of new and existing sessions on the active-primary firewall, thereby minimizing traffic on the HA3 link.



- We strongly recommended you configure HA link monitoring on the interface(s) that support the floating IP address(es) to allow each HA peer to quickly detect a link failure and fail over to its peer. Both HA peers must have link monitoring for it to function.
- We strongly recommend you configure HA path monitoring to notify each HA peer when a path has failed so a firewall can fail over to its peer. Because the floating IP address is always bound to the active-primary firewall, the firewall cannot automatically fail over to the peer when a path goes down and path monitoring is not enabled.



You cannot configure NAT for a floating IP address that is bound to an active-primary firewall.

Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall		
Step 1	Perform Step 1 through Step 5 of Configure Active/Active HA.	
Step 2	(Optional) Disable preemption.  Disabling preemption allows you full control over when the recovered firewall becomes the active-primary firewall.	1. In **Device > High Availability > General**, edit the Election Settings. 2. Clear **Preemptive** if it is enabled. 3. Click **OK**.
Step 3	Perform Step 7 through Step 14 of Configure Active/Active HA.	

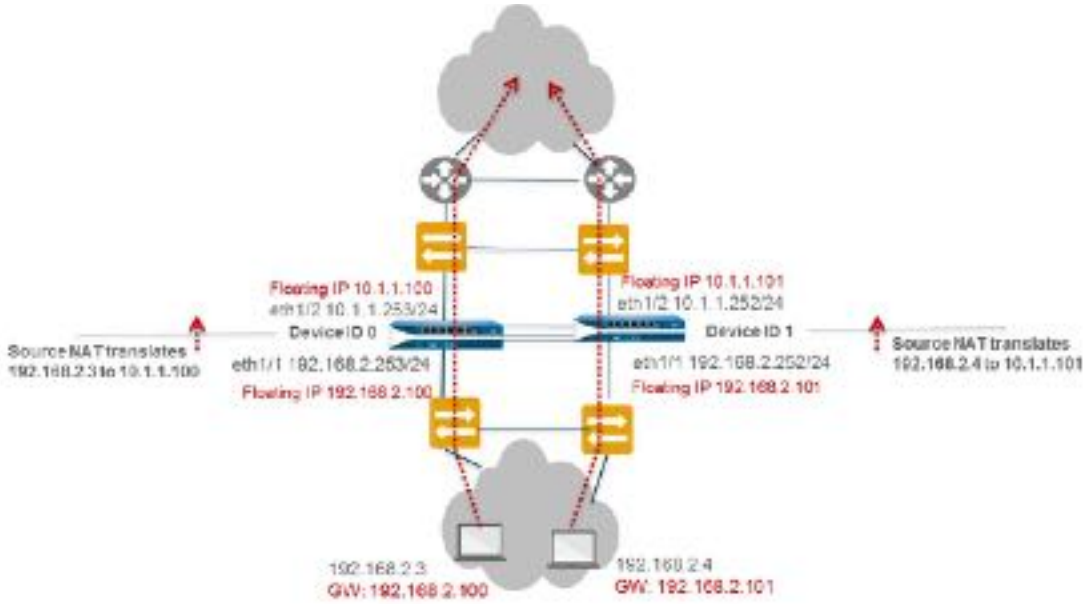
Configure Active/Active HA with Floating IP Address Bound to Active-Primary Firewall (Continued)		
Step 4	Configure Session Owner and Session Setup.	1. In **Device > High Availability > Active/Active Config**, edit Packet Forwarding. 2. For **Session Owner Selection**, we recommend you select **Primary Device**. The firewall that is in active-primary state is the session owner. Alternatively, for **Session Owner Selection** you can select **First Packet** and then for **Session Setup**, select **Primary Device** or **First Packet**. 3. For **Session Setup**, select **Primary Device**—The active-primary firewall sets up all sessions. This is the recommended setting if you want your active/active configuration to behave like an active/passive configuration because it keeps all activity on the active-primary firewall.  You must also engineer your network to eliminate the possibility of asymmetric traffic going to the HA pair. If you don't do so and traffic goes to the active-secondary firewall, setting **Session Owner Selection** and **Session Setup** to **Primary Device** causes the traffic to traverse HA3 to get to the active-primary firewall for session ownership and session setup. 4. Click **OK**.
Step 5	Configure an HA virtual address.	1. Select **Device > High Availability > Active/Active Config > Virtual Address** and click **Add**. 2. Enter or select an **Interface**. 3. Select the **IPv4** or **IPv6** tab and **Add** an **IPv4 Address** or **IPv6 Address**. 4. For **Type**, select **Floating**, which configures the virtual IP address to be a floating IP address. 5. Click **OK**.
Step 6	Bind the floating IP address to the active-primary firewall.	1. Select **Floating IP bound to the Active-Primary device**. 2. Select **Failover address if link state is down** to cause the firewall to use the failover address when the link state on the interface is down. 3. Click **OK**.
Step 7	Enable jumbo frames on firewalls other than PA-7000 Series firewalls.	Enable jumbo frames on firewalls other than PA-7000 Series firewalls.
Step 8	Save the configuration.	Click **Commit**.
Step 9	Configure the peer firewall in the same way, except selecting a different Device ID. For example, if you selected Device ID **0** for the first firewall, select Device ID **1** for the peer firewall.	

 © Palo Alto Networks, Inc.

Use Case: Configure Active/Active HA with Source DIPP NAT Using Floating IP Addresses

This Layer 3 interface example uses source NAT in Active/Active HA Mode. The Layer 2 switches create broadcast domains to ensure users can reach everything north and south of the firewalls.

PA-3050-1 has Device ID 0 and its HA peer, PA-3050-2, has Device ID 1. In this use case, NAT translates the source IP address and port number to the floating IP address configured on the egress interface. Each host is configured with a default gateway address, which is the floating IP address on Ethernet1/1 of each firewall. The configuration requires two source NAT rules, one bound to each Device ID, although you configure both NAT rules on a single firewall and they are synchronized to the peer firewall.



Configure Active/Active HA with Source DIPP NAT Using Floating IP Address		
Step 1	On PA-3050-2 (Device ID 1), perform Step 1 through Step 3 of Configure Active/Active HA.	

Configure Active/Active HA with Source DIPP NAT Using Floating IP Address (Continued)	
Step 2 Enable active/active HA.	1. In **Device > High Availability > General**, edit Setup. 2. Select **Enable HA**. 3. Enter a **Group ID**, which must be the same for both firewalls. The firewall uses the Group ID to calculate the virtual MAC address (range is 1-63). 4. For **Mode**, select **Active Active**. 5. Set the **Device ID** to **1**. 6. Select **Enable Config Sync**. This setting is required to synchronize the two firewall configurations (enabled by default). 7. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall. 8. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall. 9. Click **OK**.
Step 3 Configure Active/Active HA.	Complete Step 6 through Step 14.
Step 4 Configure Session Owner and Session Setup.	1. In **Device > High Availability > Active/Active Config**, edit Packet Forwarding. 2. For **Session Owner Selection**, select **First Packet**—The firewall that receives the first packet of a new session is the session owner. 3. For **Session Setup**, select **IP Modulo**— Distributes session setup load based on parity of the source IP address. 4. Click **OK**.
Step 5 Configure an HA virtual address.	1. Select **Device > High Availability > Active/Active Config > Virtual Address** and click **Add**. 2. Select **Interface** eth1/1. 3. Select **IPv4** and **Add** an **IPv4 Address** of 10.1.1.101. 4. For **Type**, select **Floating**, which configures the virtual IP address to be a floating IP address.
Step 6 Configure the floating IP address.	1. Do not select **Floating IP bound to the Active-Primary device**. 2. Select **Failover address if link state is down** to cause the firewall to use the failover address when the link state on the interface is down. 3. Click **OK**.
Enable jumbo frames on firewalls other than PA-7000 Series firewalls.	
Define HA Failover Conditions.	
Step 7 Save the configuration.	Click **Commit**.

 © Palo Alto Networks, Inc.

Configure Active/Active HA with Source DIPP NAT Using Floating IP Address (Continued)		
Step 8	Configure the peer firewall, PA-3050-1, with the same settings, except for the following changes: • Select **Device ID 0**. • Configure an HA virtual address of 10.1.1.100. • For **Device 1 Priority**, enter 255. For **Device 0 Priority**, enter 0. In this example, Device ID 0 has a lower priority value so a higher priority; therefore, the firewall with Device ID 0 (PA-3050-1) owns the floating IP address 10.1.1.100.	
Step 9	Still on PA-3050-1, create the source NAT rule for Device ID 0.	1. Select **Policies > NAT** and click **Add**. 2. Enter a **Name** for the rule that in this example identifies it as a source NAT rule for Device ID 0. 3. For **NAT Type**, select **ipv4** (default). 4. On the **Original Packet**, for **Source Zone**, select **Any**. 5. For **Destination Zone**, select the zone you created for the external network. 6. Allow **Destination Interface**, **Service**, **Source Address**, and **Destination Address** to remain set to **Any**. 7. For the **Translated Packet**, select **Dynamic IP And Port** for **Translation Type**. 8. For **Address Type**, select **Interface Address**, in which case the translated address will be the IP address of the interface. Select an **Interface** (eth1/1 in this example) and an **IP Address** of the floating IP address 10.1.1.100. 9. On the **Active/Active HA Binding** tab, for **Active/Active HA Binding**, select **0** to bind the NAT rule to Device ID 0. 10. Click **OK**.

Configure Active/Active HA with Source DIPP NAT Using Floating IP Address (Continued)	
Step 10 Create the source NAT rule for Device ID 1.	1. Select **Policies > NAT** and click **Add**. 2. Enter a **Name** for the policy rule that in this example helps identify it as a source NAT rule for Device ID 1. 3. For **NAT Type**, select **ipv4** (default). 4. On the **Original Packet**, for **Source Zone**, select **Any**. For **Destination Zone**, select the zone you created for the external network. 5. Allow **Destination Interface**, **Service**, **Source Address**, and **Destination Address** to remain set to **Any**. 6. For the **Translated Packet**, select **Dynamic IP And Port** for **Translation Type**. 7. For **Address Type**, select **Interface Address**, in which case the translated address will be the IP address of the interface. Select an **Interface** (eth1/1 in this example) and an **IP Address** of the floating IP address 10.1.1.101. 8. On **Active/Active HA Binding** tab, for the **Active/Active HA Binding**, select **1** to bind the NAT rule to Device ID 1. 9. Click **OK**.
Step 11 Save the configuration.	Click **Commit**.

Use Case: Configure Separate Source NAT IP Address Pools for Active/Active HA Firewalls

If you want to use IP address pools for source NAT in Active/Active HA Mode, each firewall must have its own pool, which you then bind to a Device ID in a NAT rule.

Address objects and NAT rules are synchronized (in both active/passive and active/active mode), so they need to be configured on only one of the firewalls in the HA pair.

This example configures an address object named Dyn-IP-Pool-dev0 containing the IP address pool 10.1.1.140-10.1.1.150. It also configures an address object named Dyn-IP-Pool-dev1 containing the IP address pool 10.1.1.160-10.1.1.170. The first address object is bound to Device ID 0; the second address object is bound to Device ID 1.

Create Address Objects for IP Address Pools for Source NAT in an Active/Active HA Configuration	
Step 1 On one HA firewall, create address objects.	1. Select **Objects > Addresses** and **Add** an address object **Name**, in this example, Dyn-IP-Pool-dev0. 2. For **Type**, select **IP Range** and enter the range 10.1.1.140-10.1.1.150. 3. Click **OK**. 4. Repeat this step to configure another address object named Dyn-IP-Pool-dev1 with the **IP Range** of 10.1.1.160-10.1.1.170.

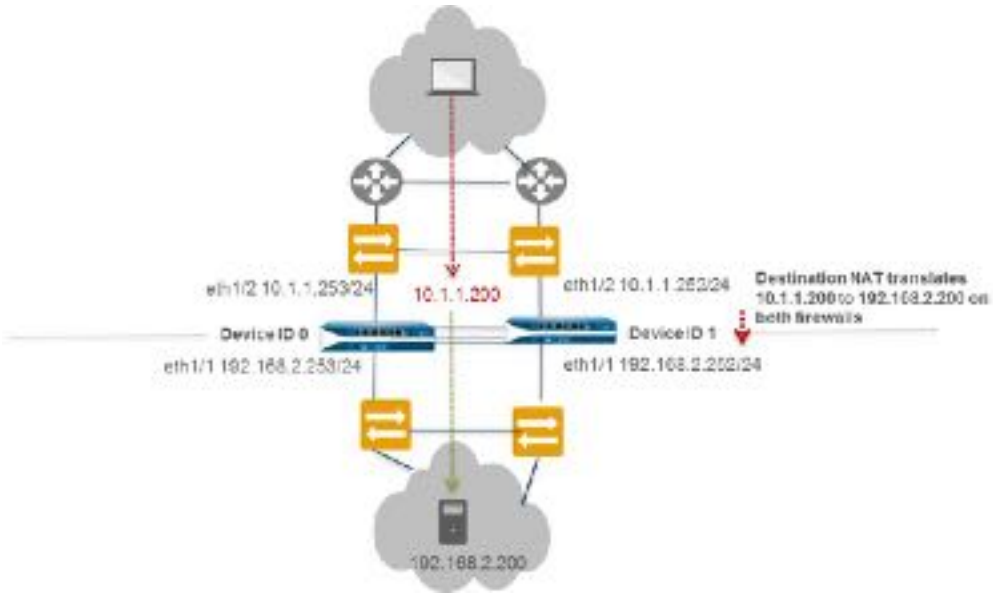
 © Palo Alto Networks, Inc.

Create Address Objects for IP Address Pools for Source NAT in an Active/Active HA Configuration (Continued)		
Step 2	Create the source NAT rule for Device ID 0.	1. Select **Policies > NAT** and **Add** a NAT policy rule with a **Name**, for example, Src-NAT-dev0. 2. For **Original Packet**, for **Source Zone**, select **Any**. 3. For **Destination Zone**, select the destination zone for which you want to translate the source address, such as Untrust. 4. For **Translated Packet**, for **Translation Type**, select **Dynamic IP and Port**. 5. For **Translated Address**, **Add** the address object you created for the pool of addresses belonging to Device ID 0: Dyn-IP-Pool-dev0. 6. For **Active/Active HA Binding**, select **0** to bind the NAT rule to Device ID 0. 7. Click **OK**.
Step 3	Create the source NAT rule for Device ID 1.	1. Select **Policies > NAT** and **Add** a NAT policy rule with a **Name**, for example, Src-NAT-dev1. 2. For **Original Packet**, for **Source Zone**, select **Any**. 3. For **Destination Zone**, select the destination zone for which you want to translate the source address, such as Untrust. 4. For **Translated Packet**, for **Translation Type**, select **Dynamic IP and Port**. 5. For **Translated Address**, **Add** the address object you created for the pool of addresses belonging to Device ID 1: Dyn-IP-Pool-dev1. 6. For **Active/Active HA Binding**, select **1** to bind the NAT rule to Device ID 1. 7. Click **OK**.
Step 4	Save the configuration.	Select **Commit**.

Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT

This Layer 3 interface example uses NAT in Active/Active HA Mode and ARP Load-Sharing with destination NAT. Both HA firewalls respond to an ARP request for the destination NAT address with the ingress interface MAC address. Destination NAT translates the public, shared IP address (in this example, 10.1.1.200) to the private IP address of the server (in this example, 192.168.2.200).

When the HA firewalls receive traffic for the destination 10.1.1.200, both firewalls could possibly respond to the ARP request, which could cause network instability. To avoid the potential issue, configure the firewall that is in active-primary state to respond to the ARP request by binding the destination NAT rule to the active-primary firewall.



Configure Active/Active HA for ARP Load-Sharing with Destination NAT	
Step 1 On PA-3050-2 (Device ID 1), perform Step 1 through Step 3 of Configure Active/Active HA.	
Step 2 Enable active/active HA.	1. In **Device > High Availability > General**, edit Setup. 2. Select **Enable HA**. 3. Enter a **Group ID**, which must be the same for both firewalls. The firewall uses the Group ID to calculate the virtual MAC address (range is 1-63). 4. (Optional) Enter a **Description**. 5. For **Mode**, select **Active Active**. 6. Select **Device ID** to be **1**. 7. Select **Enable Config Sync**. This setting is required to synchronize the two firewall configurations (enabled by default). 8. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall. 9. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall. 10. Click **OK**.
Step 3 Perform Step 6 through Step 15 in Configure Active/Active HA.	

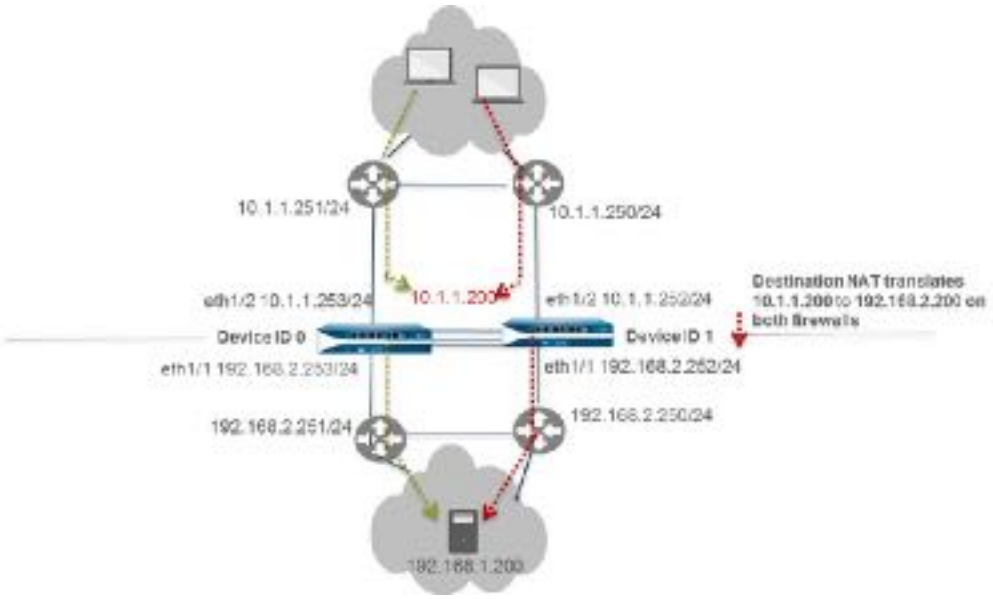
 © Palo Alto Networks, Inc.

Configure Active/Active HA for ARP Load-Sharing with Destination NAT (Continued)		
Step 4	Configure an HA virtual address.	1. Select **Device > High Availability > Active/Active Config > Virtual Address** and click **Add**. 2. Select **Interface** eth1/1. 3. Select **IPv4** and **Add** an **IPv4 Address** of 10.1.1.200. 4. For **Type**, select **ARP Load Sharing**, which configures the virtual IP address to be for both peers to use for ARP Load-Sharing.
Step 5	Configure ARP Load-Sharing. The device selection algorithm determines which HA firewall responds to the ARP requests to provide load sharing.	1. For **Device Selection Algorithm**, select **IP Modulo**. The firewall that will respond to ARP requests is based on the parity of the ARP requester's IP address. 2. Click **OK**.
Step 6	Enable jumbo frames on firewalls other than PA-7000 Series firewalls.	
Step 7	Define HA Failover Conditions.	
Step 8	Save the configuration.	Click **Commit**.
Step 9	Configure the peer firewall, PA-3050-1 (Device ID 0), with the same settings, except in Step 2 select **Device ID 0**.	
Step 10	Still on PA-3050-1 (Device ID 0), create the destination NAT rule so that the active-primary firewall responds to ARP requests.	1. Select **Policies > NAT** and click **Add**. 2. Enter a **Name** for the rule that, in this example, identifies it as a destination NAT rule for Layer 2 ARP. 3. For **NAT Type**, select **ipv4** (default). 4. On the **Original Packet**, for **Source Zone**, select **Any**. 5. For **Destination Zone**, select the Untrust zone you created for the external network. 6. Allow **Destination Interface**, **Service**, and **Source Address** to remain set to **Any**. 7. For **Destination Address**, specify 10.1.1.200. 8. For the **Translated Packet**, Source Address Translation remains **None**. 9. For **Destination Address Translation**, enter the private IP address of the destination server, in this example, 192.168.1.200. 10. On the **Active/Active HA Binding** tab, for **Active/Active HA Binding**, select **primary** to bind the NAT rule to the firewall in active-primary state. 11. Click **OK**.
Step 11	Save the configuration.	Click **Commit**.

Use Case: Configure Active/Active HA for ARP Load-Sharing with Destination NAT in Layer 3

This Layer 3 interface example uses NAT in Active/Active HA Mode and ARP Load-Sharing. PA-3050-1 has Device ID 0 and its HA peer, PA-3050-2, has Device ID 1.

In this use case, both of the HA firewalls must respond to an ARP request for the destination NAT address. Traffic can arrive at either firewall from either WAN router in the untrust zone. Destination NAT translates the public-facing, shared IP address to the private IP address of the server. The configuration requires one destination NAT rule bound to both Device IDs so that both firewalls can respond to ARP requests.



Configure Active/Active HA for ARP Load-Sharing with Destination NAT in Layer 3	
Step 1 On PA-3050-2 (Device ID 1), perform Step 1 through Step 3 of Configure Active/Active HA.	

 © Palo Alto Networks, Inc.

Configure Active/Active HA for ARP Load-Sharing with Destination NAT in Layer 3 (Continued)		
Step 2	Enable active/active HA.	1. Select **Device > High Availability > General > Setup** and edit. 2. Select **Enable HA**. 3. Enter a **Group ID**, which must be the same for both firewalls. The firewall uses the Group ID to calculate the virtual MAC address (range is 1-63). 4. (Optional) Enter a **Description**. 5. For **Mode**, select **Active Active**. 6. Select **Device ID** to be **1**. 7. Select **Enable Config Sync**. This setting is required to synchronize the two firewall configurations (enabled by default). 8. Enter the **Peer HA1 IP Address**, which is the IP address of the HA1 control link on the peer firewall. 9. (Optional) Enter a **Backup Peer HA1 IP Address**, which is the IP address of the backup control link on the peer firewall. 10. Click **OK**.
Step 3	Configure Active/Active HA.	Perform Step 6 through Step 15.
Step 4	Configure an HA virtual address.	1. Select **Device > High Availability > Active/Active Config > Virtual Address** and click **Add**. 2. Select **Interface** eth1/2. 3. Select **IPv4** and **Add** an **IPv4 Address** of 10.1.1.200. 4. For **Type**, select **ARP Load Sharing**, which configures the virtual IP address to be for both peers to use for ARP Load-Sharing.
Step 5	Configure ARP Load-Sharing. The device selection algorithm determines which HA firewall responds to the ARP requests to provide load sharing.	1. For **Device Selection Algorithm**, select one of the following • **IP Modulo**—The firewall that will respond to ARP requests is based on the parity of the ARP requester's IP address. • **IP Hash**—The firewall that will respond to ARP requests is based on a hash of the ARP requester's source IP address and destination IP address. 2. Click **OK**.
Step 6	Enable jumbo frames on firewalls other than PA-7000 Series firewalls.	
Step 7	Define HA Failover Conditions.	
Step 8	Save the configuration.	Click **Commit**.
Step 9	Configure the peer firewall, PA-3050-1 (Device ID 0), with the same settings, except set the **Device ID** to **0** instead of **1**.	

Configure Active/Active HA for ARP Load-Sharing with Destination NAT in Layer 3 (Continued)	
Step 10 Still on PA-3050-1 (Device ID 0), create the destination NAT rule for both Device ID 0 and Device ID 1.	1. Select **Policies > NAT** and click **Add**. 2. Enter a **Name** for the rule that in this example identifies it as a destination NAT rule for Layer 3 ARP. 3. For **NAT Type**, select **ipv4** (default). 4. On the **Original Packet**, for **Source Zone**, select **Any**. 5. For **Destination Zone**, select the Untrust zone you created for the external network. 6. Allow **Destination Interface**, **Service**, and **Source Address** to remain set to **Any**. 7. For **Destination Address**, specify 10.1.1.200. 8. For the **Translated Packet**, Source Address Translation remains None. 9. For **Destination Address Translation**, enter the private IP address of the destination server, in this example 192.168.1.200. 10. On the **Active/Active HA Binding** tab, for **Active/Active HA Binding**, select **both** to bind the NAT rule to both Device ID 0 and Device ID 1. 11. Click **OK**.
Step 11 Save the configuration.	Click **Commit**.

 © Palo Alto Networks, Inc.

HA Firewall States

An HA firewall can be in one of the following states:

HA Firewall State	Occurs In	Description
Initial	A/P or A/A	Transient state of a firewall when it joins the HA pair. The firewall remains in this state after boot-up until it discovers a peer and negotiations begins. After a timeout, the firewall becomes active if HA negotiation has not started.
Active	A/P	State of the active firewall in an active/passive configuration.
Passive	A/P	State of the passive firewall in an active/passive configuration. The passive firewall is ready to become the active firewall with no disruption to the network. Although the passive firewall is not processing other traffic: • If passive link state auto is configured, the passive firewall is running routing protocols, monitoring link and path state, and the passive firewall will pre-negotiate LACP and LLDP if LACP and LLDP pre-negotiation are configured, respectively. • The passive firewall is synchronizing flow state, runtime objects, and configuration. • The passive firewall is monitoring the status of the active firewall using the hello protocol.
Active-Primary	A/A	In an active/active configuration, state of the firewall that connects to User-ID agents, runs DHCP server and DHCP relay, and matches NAT and PBF rules with the Device ID of the active-primary firewall. A firewall in this state can own sessions and set up sessions.
Active-Secondary	A/A	In an active/active configuration, state of the firewall that connects to User-ID agents, runs DHCP server, and matches NAT and PBF rules with the Device ID of the active-secondary firewall. A firewall in active-secondary state does not support DHCP relay. A firewall in this state can own sessions and set up sessions.

HA Firewall State	Occurs In	Description
Tentative	A/A	State of a firewall (in an active/active configuration) caused by one of the following: • Failure of a firewall. • Failure of a monitored object (a link or path). • The firewall leaves suspended or non-functional state. A firewall in tentative state synchronizes sessions and configurations from the peer. • In a virtual wire deployment, when a firewall enters tentative state due to a path failure and receives a packet to forward, it sends the packet to the peer firewall over the HA3 link for processing. The peer firewall processes the packet and sends it back over the HA3 link to the firewall to be sent out the egress interface. This behavior preserves the forwarding path in a virtual wire deployment. • In a Layer 3 deployment, when a firewall in tentative state receives a packet, it sends that packet over the HA3 link for the peer firewall to own or set up the session. Depending on the network topology, this firewall either sends the packet out to the destination or sends it back to the peer in tentative state for forwarding. After the failed path or link clears or as a failed firewall transitions from tentative state to active-secondary state, the **Tentative Hold Time** is triggered and routing convergence occurs. The firewall attempts to build routing adjacencies and populate its route table before processing any packets. Without this timer, the recovering firewall would enter active-secondary state immediately and would blackhole packets because it would not have the necessary routes. When a firewall leaves suspended state, it goes into tentative state for the **Tentative Hold Time** after links are up and able to process incoming packets. **Tentative Hold Time range (sec)** can be disabled (which is 0 seconds) or in the range 10-600; default is 60.
Non-functional	A/P or A/A	Error state due to a dataplane failure or a configuration mismatch, such as only one firewall configured for packet forwarding, VR sync or QoS sync. In active/passive mode, all of the causes listed for Tentative state cause non-functional state.
Suspended	A/P or A/A	Administratively disabled state. In this state, an HA firewall cannot participate in the HA election process.

 © 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.



Only committed configurations synchronize between HA peers. Any configuration in the commit queue at the time of an HA sync will not be synchronized.

The following topics identify which configuration settings you must configure on each firewall independently (these settings are not 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 firewall, including: • **Device > Setup > Management > General Settings**—Hostname, Domain, Login Banner, SSL/TLS Service Profile, Time Zone, Locale, Date, Time, Latitude, Longitude. The configuration for the associated SSL/TLS Service profile (**Device > Certificate Management > SSL/TLS Service Profile** and the associated certificates (**Device > Certificate Management > Certificates**) is synchronized. It is just the setting of which SSL/TLS Service Profile to use on the Management interface that does not sync. • **Device > Setup > Management > Management Interface Settings**—IP Type, 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. You must also enable **Multi Virtual System Capability** on each firewall (**Device > Setup > Management > General Settings**).

Configuration Item	What Doesn't Sync in Active/Passive?
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**). • **Panorama Servers** • **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**
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 firewall, or download them on one peer and sync the update to the other peer. You must install the update on each peer. • **Device > Software**
GlobalProtect Agent Package	With GlobalProtect client updates, you can either download and install them separately on each firewall, or download them to one peer and sync the update to the other peer. You must activate separately on each peer. • **Device > GlobalProtect Client**
Content Updates	With content updates, you can either download and install them separately on each firewall, or download them on one peer and sync the update to the other peer. You must install the update on each peer. • **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 firewall (**Device > Master Key and Diagnostics**). 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.
HA settings	• **Device > High Availability**

 © 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 in Active/Active?
Management Interface Settings	You must configure all management settings individually on each firewall, including: • **Device > Setup > Management > General Settings**—Hostname, Domain, Login Banner, SSL/TLS Service Profile, Time Zone, Locale, Date, Time, Latitude, Longitude.  The configuration for the associated SSL/TLS Service profile (**Device > Certificate Management > SSL/TLS Service Profile** and the associated certificates (**Device > Certificate Management > Certificates**) is synchronized. It is just the setting of which SSL/TLS Service Profile to use on the Management interface that does not sync. • **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. 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**). • **Panorama Servers** • **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 in Active/Active?
Software Updates	With software updates, you can either download and install them separately on each firewall, or download them on one peer and sync the update to the other peer. You must install the update on each peer. • **Device > Software**
GlobalProtect Agent Package	With GlobalProtect client updates, you can either download and install them separately on each firewall, or download them to one peer and sync the update to the other peer. You must activate separately on each peer. • **Device > GlobalProtect Client**
Content Updates	With content updates, you can either download and install them separately on each firewall, or download them on one peer and sync the update to the other peer. You must install the update on each peer. • **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.
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.

 © Palo Alto Networks, Inc.

Configuration Item	What Doesn't Sync in Active/Active?
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 > Network Profiles > LLDP**).
BFD	No BFD configuration or BFD session data is synchronized in an active/active configuration (**Network > Network Profiles > BFD Profile**).
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 firewall (**Device > Master Key and Diagnostics**). 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.
HA settings	• **Device > High Availability** • (The exception is **Device > High Availability > Active/Active Configuration > Virtual Addresses**, which do sync.)

Synchronization of System Runtime Information

The following table summarizes what system runtime information is synchronized between HA peers.

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	No	No	N/A	
FQDN Refresh	No	No	N/A	
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. • Active/active peers do not sync host session, multicast session, or BFD session information.
ARP Table	Yes	No	HA2	Upon upgrade to PAN-OS 7.1, the ARP table capacity automatically increases. To avoid a mismatch, upgrade both peers within a short period of time. As a best practice, clear the ARP cache (`clear arp`) on both peers prior to upgrading to PAN-OS 7.1.
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.

Monitoring

In order 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

Use the Dashboard

The **Dashboard** tab widgets show general firewall 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 firewall 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. 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 firewall—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.

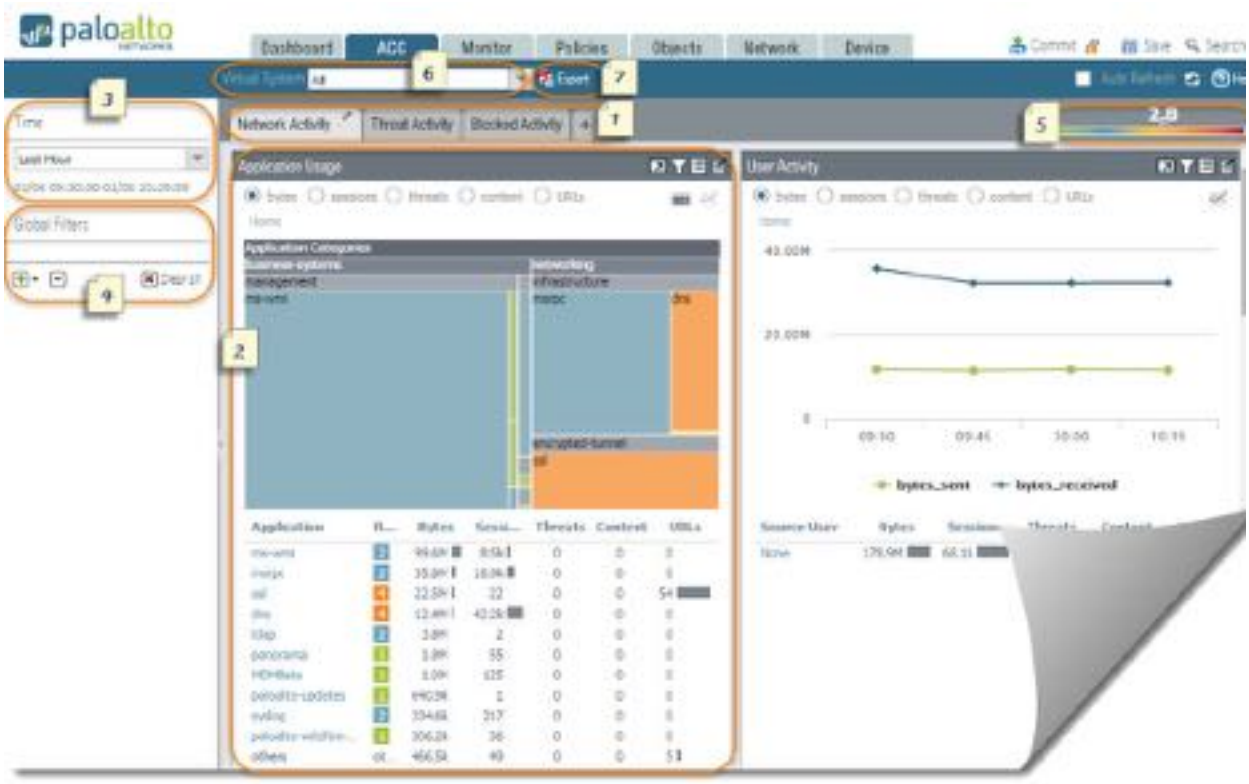
Use the Application Command Center

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: • bytes (in and out) • sessions • content (files and data) • URL categories • threats (and count) 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. 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. 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. 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. 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: • Top applications in use • Top users who generate traffic (with a drill down into the bytes, content, threats or URLs accessed by the user) • Most used security rules against which traffic matches occur 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.

Tab	Description
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.

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.



 © Palo Alto Networks, Inc.

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. 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.
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: • 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. • Hover over the attribute in the table and use the options available in the drop-down. 
4	**Actions**	**Maximize view**— Allows you enlarge the widget and view the table in a larger screen space and with more viewable information. **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. **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. **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.

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.	

Widget	Description
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. Sort attributes: bytes, sessions, threats, content, URLs 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. Sort attributes: bytes, sessions, threats, content, URLs 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. Sort attributes: bytes, sessions, threats, content, URLs 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. Sort attributes: bytes, sessions, threats, content, URLs 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. Sort attributes: bytes, sessions, threats, content, URLs 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. Sort attributes: bytes, sessions, threats, content, URLs Charts available: map, bar
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. Sort attributes: profiles, objects, operating systems 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. Sort attributes: bytes, sessions, threats, content, URLs Charts available: line

 © Palo Alto Networks, Inc.

Widget	Description
Ingress Interfaces	Displays the firewall interfaces that are most used for allowing traffic into the network. Sort attributes: bytes, bytes sent, bytes received Charts available: line
Egress Interfaces	Displays the firewall interfaces that are most used by traffic exiting the network. Sort attributes: bytes, bytes sent, bytes received Charts available: line
Source Zones	Displays the zones that are most used for allowing traffic into the network. Sort attributes: bytes, sessions, threats, content, URLs Charts available: line
Destination Zones	Displays the zones that are most used by traffic going outside the network. Sort attributes: bytes, sessions, threats, content, URLs 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. Available on the PA-3000 Series, PA-5000 Series, PA-7000 Series, and Panorama. 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. Sort attributes: count Charts available: line
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. Sort attributes: count 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. Sort attributes: threats 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. Sort attributes: malicious, benign Charts available: bar, line

Widget	Description
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. Sort attributes: malicious, benign 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. Sort attributes: bytes, sessions, threats, content, URLs 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. 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. Sort attributes: bytes, sessions, threats, content, URLs Charts available: treemap, line

Blocked Activity—Focuses on traffic that was prevented from coming into the network

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. Sort attributes: threats, content, URLs 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 rule. Sort attributes: threats, content, URLs 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. Sort attributes: threats 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. Sort attributes: files, data Charts available: bar, area, column

 © Palo Alto Networks, Inc.

Widget	Description
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. Sort attributes: threats, content, URLs Charts available: bar, area, column

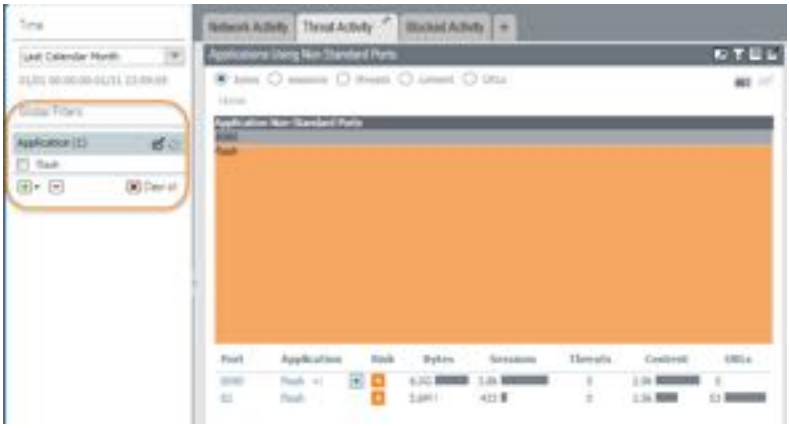
ACC Filters

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.
- **Add a widget filter to a global filter**—Hover over the attribute and click the arrow icon to the right of the attribute. 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.

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. 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 .  Editing a tab allows 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. 2. Select the **Add Widget** drop-down and verify the widgets that have the check boxes selected.

 © Palo Alto Networks, Inc.

Work with the Tabs and Widgets (Continued)	
• Add a widget or a widget group.	1. Add a new tab or edit a predefined tab. 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. 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. 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.  You cannot delete a predefined tab. 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. 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.
• Use the table drop-down to find more information on an attribute.	1. Hover over an attribute in a table to see the drop-down. 2. Click into the drop-down to view the available options. • **Global Find**—Use Global Find to Search the Firewall or Panorama Management Server for references to the attribute (username/IP address, object name, policy rule name, threat ID, or application name) anywhere in the candidate configuration. • **Value**—Displays the details of the threat ID, or application name, or address object. • **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. • **Search HIP Report**—Uses the username or IP address to find matches in a HIP Match report.
• Set a widget filter. 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. 2. Click the icon to add the filters you want to apply. 3. Click **Apply**. These filters are persistent across reboots. The active widget filters are indicated next to the widget name.

Work with the Tabs and Widgets (Continued)	
• Negate a widget filter	1. Click the icon to display the Setup Local Filters dialog. 2. Add a filter, and then click the negate icon.
• Set a global filter from a table.	Hover over an attribute in the table below the chart and click the arrow icon to the right of the attribute. 
• Set a global filter using the Global Filters pane. Watch global filters in action.	1. Locate the **Global Filters** pane on the left side of the ACC.  2. Click the icon to view the list of filters you can apply.
• 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. 2. To promote the filter to be a global filter, select the arrow to the right of the filter. 
• Remove a filter.	Click the icon to remove a filter. • For global filters: It is located in the Global Filters pane. • For widget filters: Click the icon to display the Setup Local Filters dialog, then select the filter, and click the icon.
• Clear all filters.	• For global filters: Click the **Clear All** button under Global Filters. • For widget filters: Select a widget and click the icon. Then click the **Clear All** button in the Setup Local Filters dialog.

 © Palo Alto Networks, Inc.

Work with the Tabs and Widgets (Continued)	
• See what filters are in use.	• For global filters: The number of global filters applied are displayed on the left pane under Global Filters. • 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. 

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.



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.

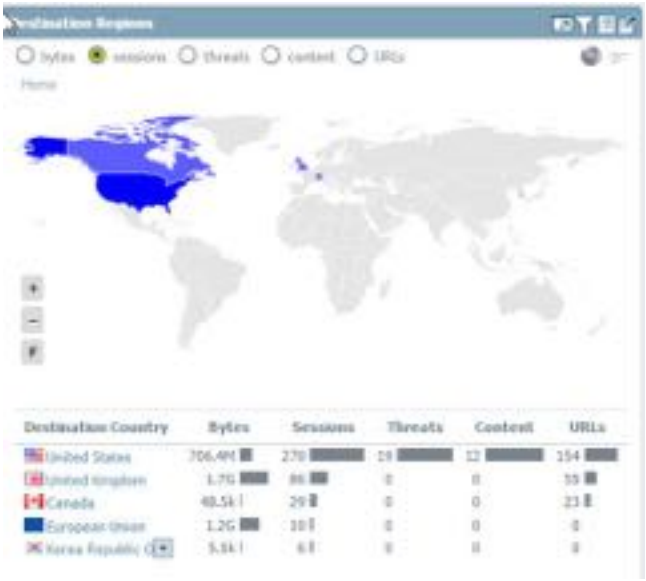
Consider whether you want to sanction rapidshare for company use. Should you allow uploads to this site and do you need a QoS policy to limit bandwidth?

To view which IP addresses Marsha has communicated with, check the **Destination IP Activity** widget, and view the data by bytes and by URLs.

 © Palo Alto Networks, Inc.



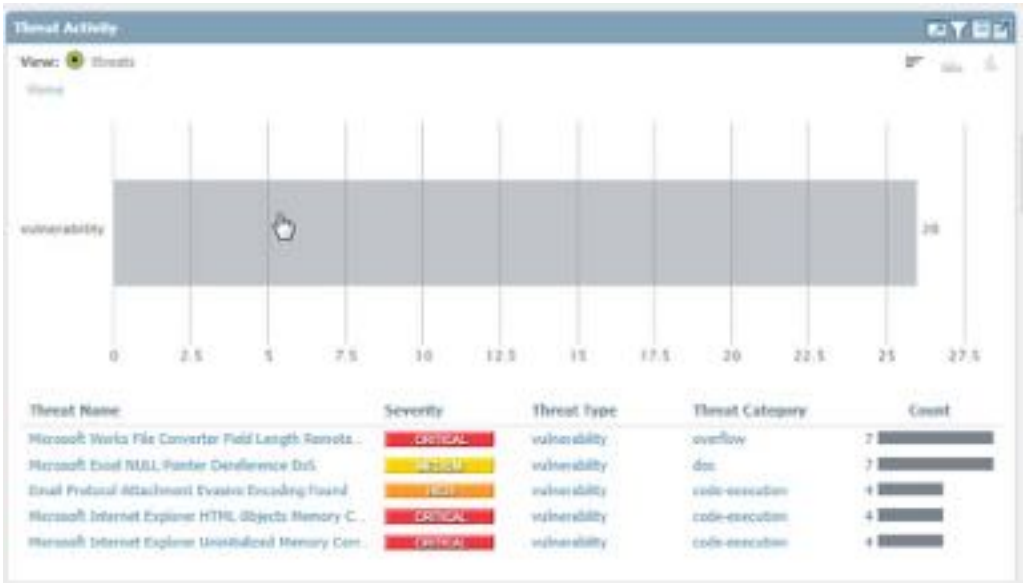
To find out which countries Marsha communicated with, sort on **sessions** in the **Destination Regions** widget.



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.



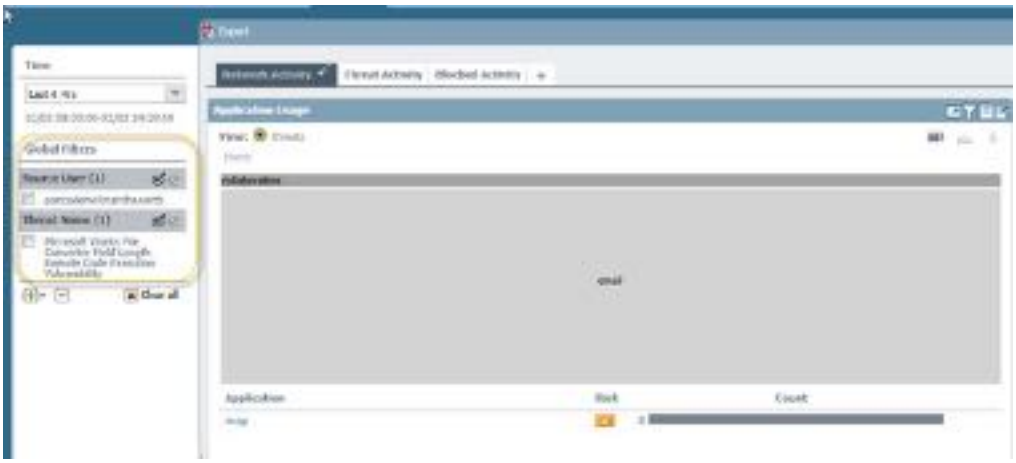


To further drill-down into each vulnerability, click into the graph and narrow the scope of your investigation. Each click automatically applies a local filter on the widget.



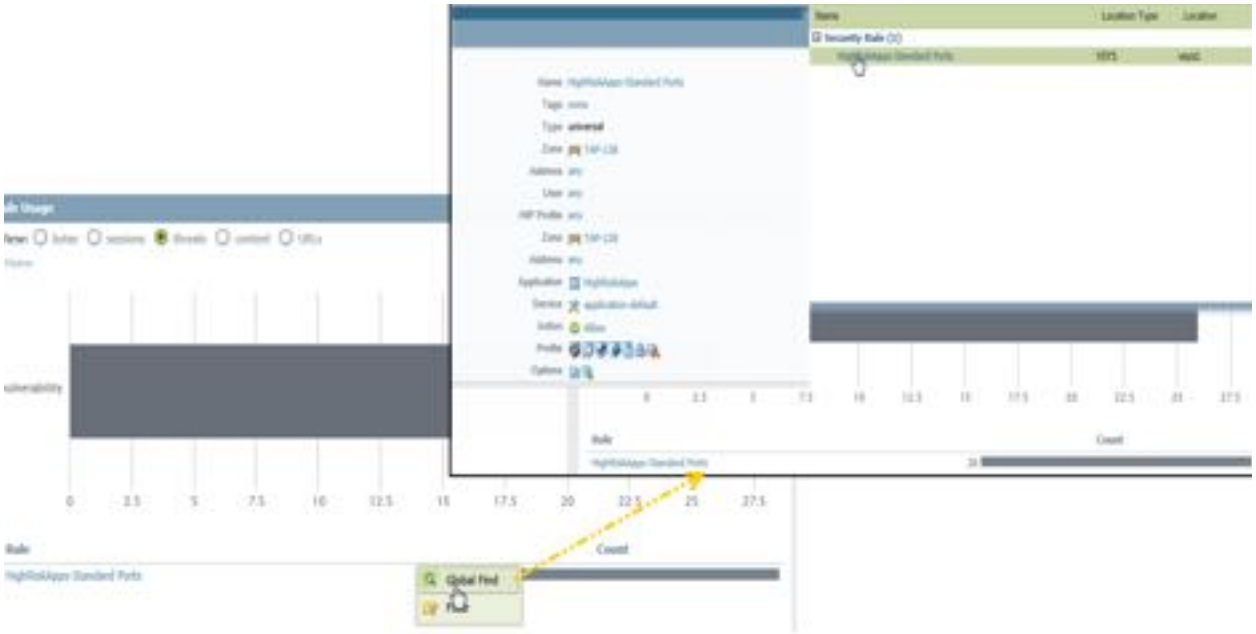
To investigate each threat by name, you can create a global filter for say, **Microsoft Works File Converter Field Length Remote Code Execution Vulnerability**. Then, view the **User Activity widget** in the **Network Activity** tab. The tab is automatically filtered to display threat activity for Marsha (notice the global filters in the screenshot).

 © Palo Alto Networks, Inc.



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.

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**).

 © Palo Alto Networks, Inc.



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.

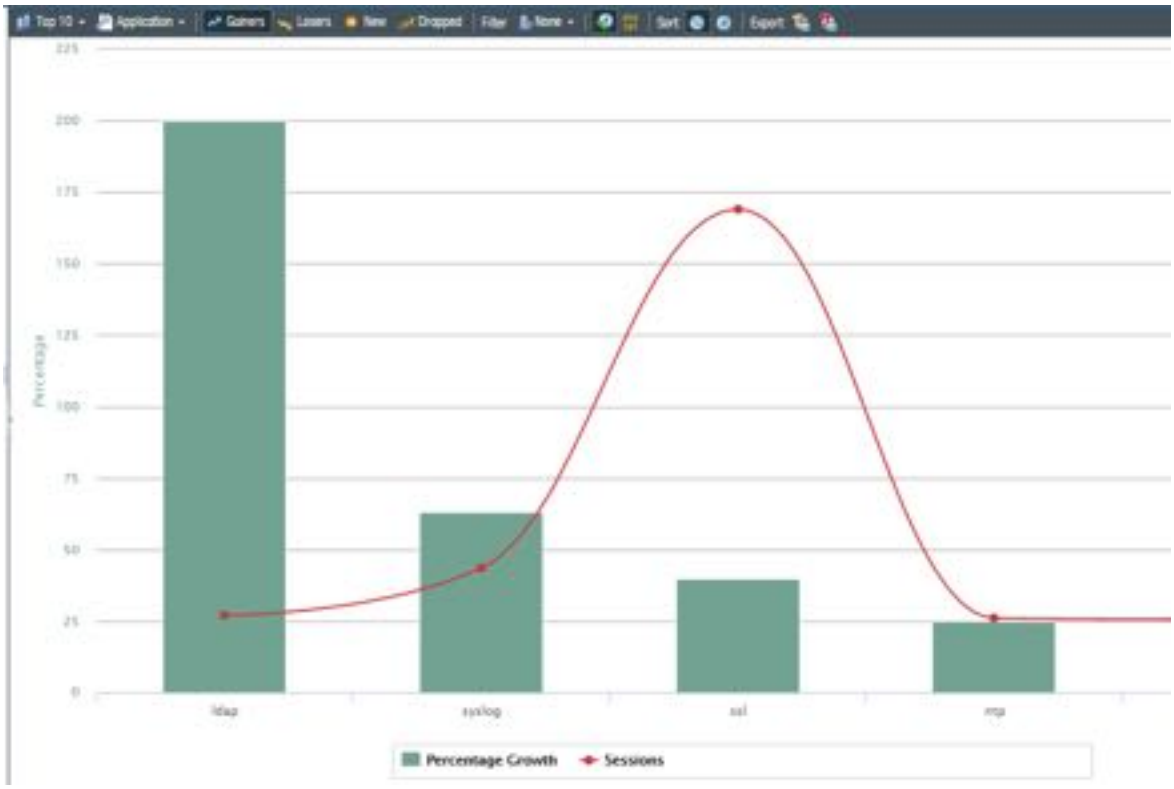
Summary Report

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.

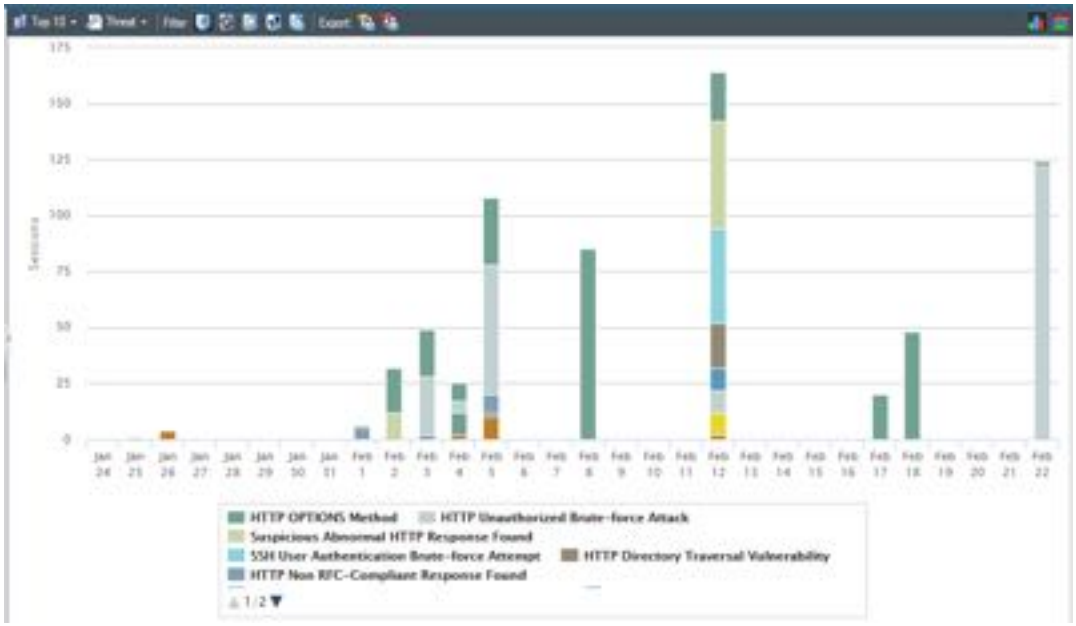
Button	Description
Top 10	Determines the number of records with the highest measurement included in the chart.
Application	Determines the type of item reported: Application, Application Category, Source, or Destination.
Gainers	Displays measurements of items that have increased over the measured period.
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
Filter	Applies a filter to display only the selected item. **None** displays all entries.
	Determines whether to display session or byte information.
Sort	Determines whether to sort entries by percentage or raw growth.
Export	Exports the graph as a .png image or as a PDF.
Compare	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
Top 10	Determines the number of records with the highest measurement included in the chart.
Threats	Determines the type of item measured: Threat, Threat Category, Source, or Destination.

Button	Description
Filter	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.
Export	Exports the graph as a .png image or as a PDF.
	Specifies the period over which the measurements are taken.

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.



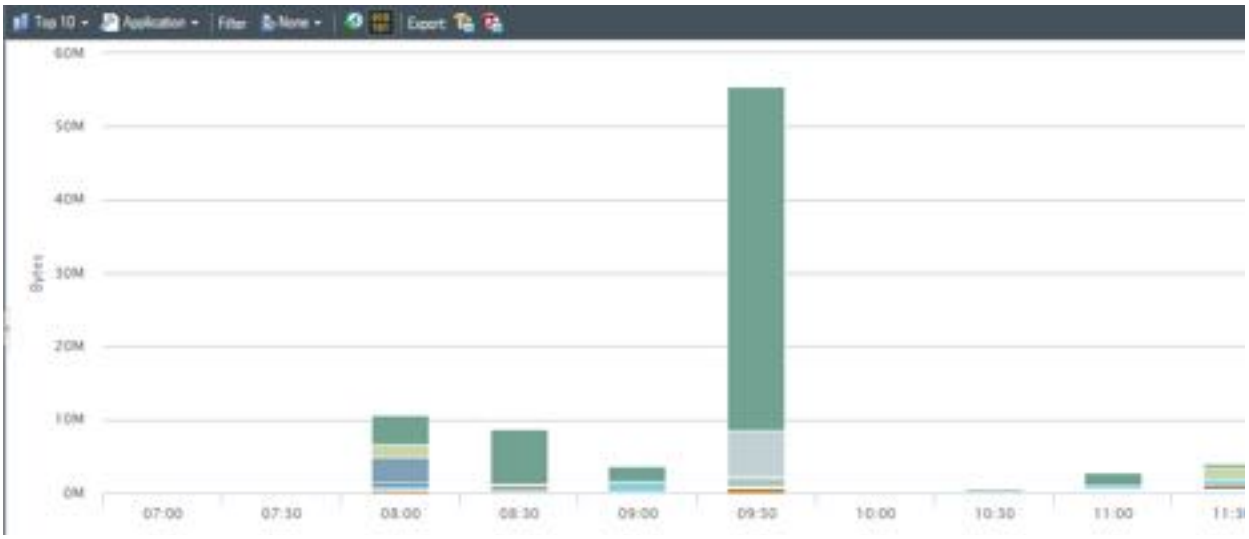
The Threat Map report contains the following buttons and options.

Button	Description
Top 10	Determines the number of records with the highest measurement included in the chart.
Incoming threats	Displays incoming threats.
Outdoing threats	Displays outgoing threats.
Filer	Applies a filter to display only the selected type of items.
Zoom In and Zoom Out	Zoom in and zoom out of the map.
Export	Exports the graph as a .png image or as a PDF.
	Indicates the period over which the measurements are taken.

 © Palo Alto Networks, Inc.

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
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.
Filter	Applies a filter to display only the selected item. **None** displays all entries.
	Determines whether to display session or byte information.
Export	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.

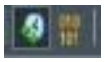
Traffic Map Report

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.
	Determines whether to display session or byte information.
Zoom In and Zoom Out	Zoom in and zoom out of the map.
Export	Exports the graph as a .png image or as a PDF.
	Indicates the period over which the change measurements are taken.

 © Palo Alto Networks, Inc.

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

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.

© Palo Alto Networks, Inc.

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.

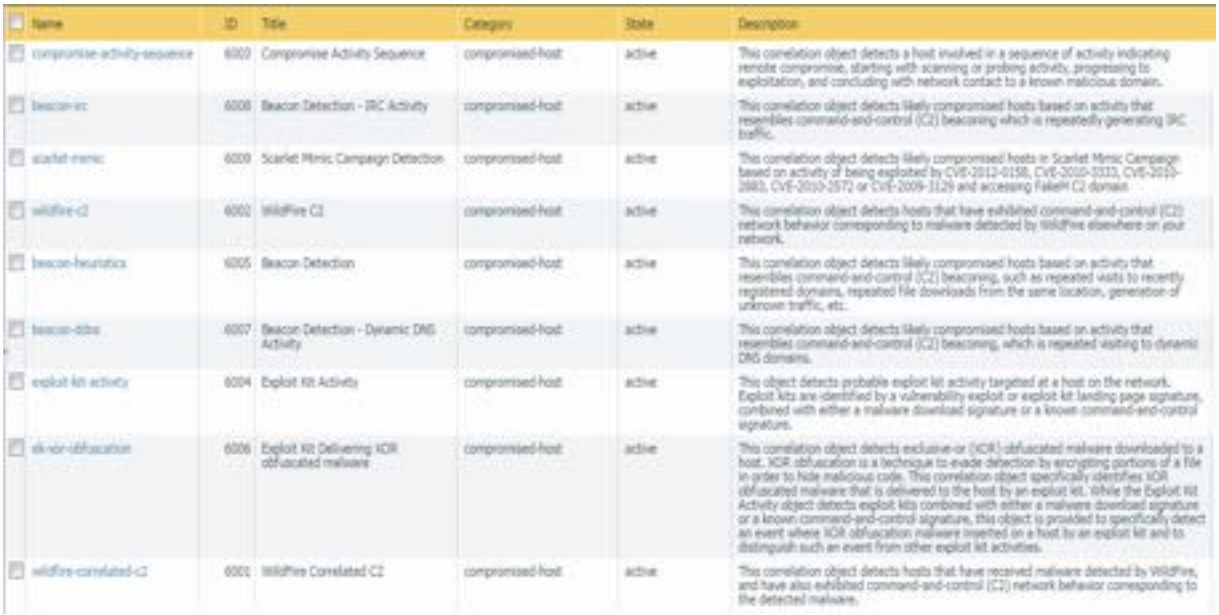
Correlated Events

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.



Name	ID	Title	Category	State	Description
compromise-activity-sequence	6002	Compromise Activity Sequence	compromised-host	active	This correlation object detects a host involved in a sequence of activity indicating remote compromise, starting with scanning or probing activity, progressing to exploitation, and concluding with network contact to a known malicious domain.
beacon-irc	6008	Beacon Detection - IRC Activity	compromised-host	active	This correlation object detects likely compromised hosts based on activity that resembles command-and-control (C2) beaconing which is repeatedly generating IRC traffic.
scarlet-mimic	6009	Scarlet Mimic Campaign Detection	compromised-host	active	This correlation object detects likely compromised hosts in Scarlet Mimic Campaign based on activity of being exploited by CVE-2012-0158, CVE-2010-3333, CVE-2010-2883, CVE-2010-2572 or CVE-2009-3129 and accessing FakeIM C2 domain
wildfire-c2	6002	WildFire C2	compromised-host	active	This correlation object detects hosts that have exhibited command-and-control (C2) network behavior corresponding to malware detected by WildFire elsewhere on your network.
beacon-heuristics	6005	Beacon Detection	compromised-host	active	This correlation object detects likely compromised hosts based on activity that resembles command-and-control (C2) beaconing, such as repeated visits to recently registered domains, repeated file downloads from the same location, generation of unknown traffic, etc.
beacon-ddns	6007	Beacon Detection - Dynamic DNS Activity	compromised-host	active	This correlation object detects likely compromised hosts based on activity that resembles command-and-control (C2) beaconing, which is repeated visiting to dynamic DNS domains.
exploit-kit-activity	6004	Exploit Kit Activity	compromised-host	active	This object detects probable exploit kit activity targeted at a host on the network. Exploit kits are identified by a vulnerability exploit or exploit kit landing page signature, combined with either a malware download signature or a known command-and-control signature.
xk-xor-obfuscation	6006	Exploit Kit Delivering XOR obfuscated malware	compromised-host	active	This correlation object detects exclusive-or (XOR) obfuscated malware downloaded to a host. XOR obfuscation is a technique to evade detection by encrypting portions of a file in order to hide malicious code. This correlation object specifically identifies XOR obfuscated malware that is delivered to the host by an exploit kit. While the Exploit Kit Activity object detects exploit kits combined with either a malware download signature or a known command-and-control signature, this object is provided to specifically detect an event where XOR obfuscation malware inserted on a host by an exploit kit and to distinguish such an event from other exploit kit activities.
wildfire-correlated-c2	6001	WildFire Correlated C2	compromised-host	active	This correlation object detects hosts that have received malware detected by WildFire, and have also exhibited command-and-control (C2) network behavior corresponding to the detected malware.

 © Palo Alto Networks, Inc.

View the Correlation Objects Available on the Firewall
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.

Interpret Correlated Events

You can view and analyze the logs generated for each correlated event in the **Monitor > Automated Correlation Engine > Correlated Events** tab.

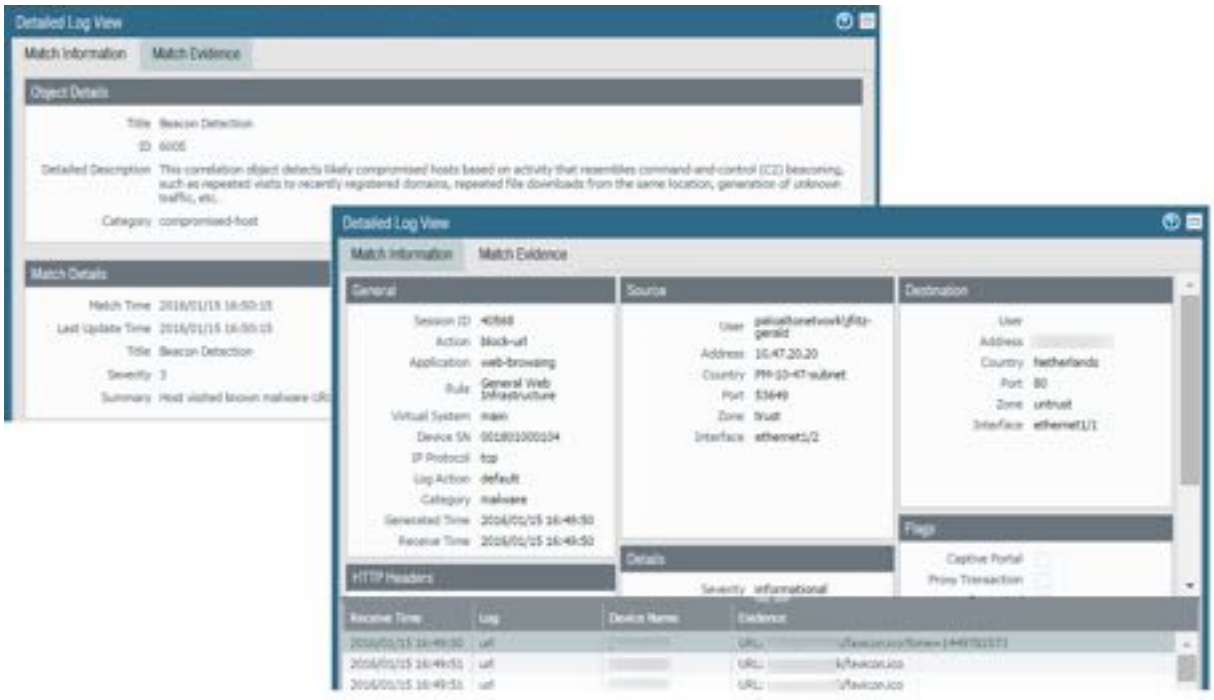


Correlated Events includes the following details:

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 which the traffic originated.
Source User	The user and user group information from the directory server, if User-ID is enabled.

Field	Description
Severity 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: • **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. • **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. • **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. • **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, but the event is not necessarily significant on its own.
Summary	A description that summarizes the evidence gathered on the correlated event.

Click the icon to see the detailed log view, which includes all the evidence on a match:



 © Palo Alto Networks, Inc.

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.

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.

- Types of Packet Captures
- Disable Hardware Offload
- Take a Custom Packet Capture
- Take a Threat Packet Capture
- Take an Application Packet Capture
- Take a Packet Capture on the Management Interface

Types of Packet Captures

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 firewall 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-3060, 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: `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: `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.

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.



The following example shows how to use a packet capture to troubleshoot a Telnet connectivity issue from a user in the Trust zone to a server in the DMZ zone.



 © Palo Alto Networks, Inc.

Take a Custom Packet Capture (Continued)

Step 2 Set packet capture filters, so the firewall only captures traffic you are interested in.

Using filters makes 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**.

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`.

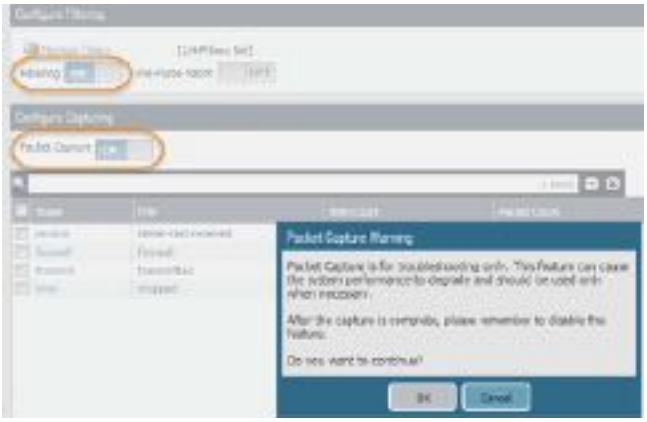


2. Continue to **Add** each **Stage** you want to capture (**receive, firewall**, **transmit**, and **drop**) and set a unique **File** name for each stage.



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.



Step 6 Generate traffic that matches the filters that you defined.

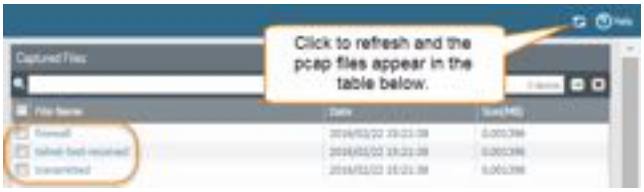
For this example, generate traffic from the source system to the Telnet-enabled server by running the following command from the source system (192.168.2.10):

```
telnet 10.43.14.55
```

 © Palo Alto Networks, Inc.

Take a Custom Packet Capture (Continued)

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.

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.

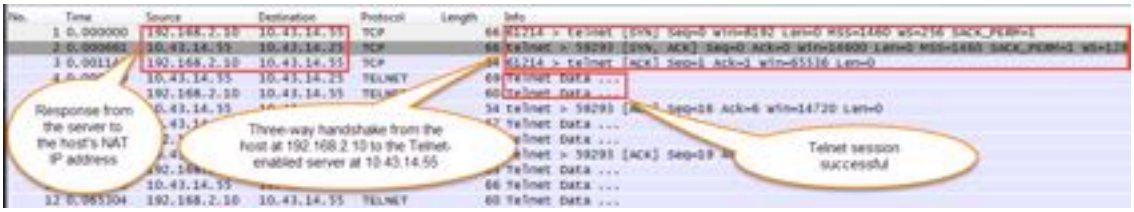
Step 11 Generate traffic that will trigger the packet capture.

Run the Telnet session again from the source system to the Telnet-enabled server

```
telnet 10.43.14.55
```

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.

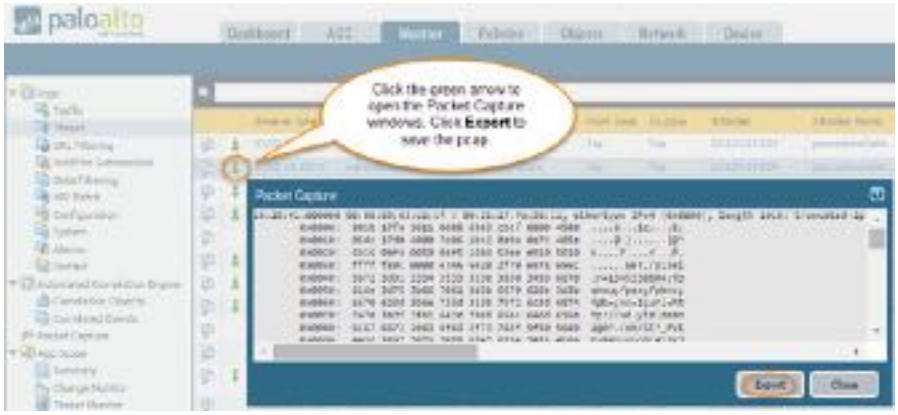
Take a Threat Packet Capture		
Step 1	Enable the packet capture option in the security profile. Some security profiles allow you to define a single-packet capture, or extended-capture. If you choose extended-capture, define the capture length. This will allow the firewall to capture more packets to provide 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**.
Step 2	Add the security profile (with packet capture enabled) to a Security Policy rule.	1. Select **Policies > Security** and select a rule. 2. Select the **Actions** tab. 3. In the Profile Settings section, select a profile that has packet capture enabled. For example, click the **Antivirus** drop-down and select a profile that has packet capture enabled.

 © Palo Alto Networks, Inc.

Take a Threat Packet Capture (Continued)

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 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"
   ```

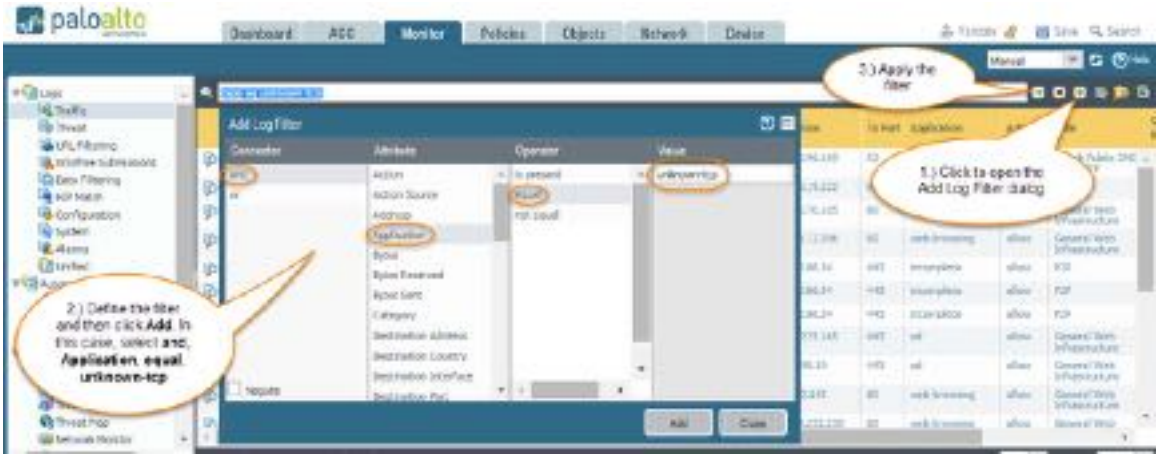
2. If the unknown capture setting option is off, enable it:

   ```
   admin@PA-200> set application dump-unknown yes
   ```

Identify Unknown Applications in Traffic Logs and View Packet Captures (Continued)

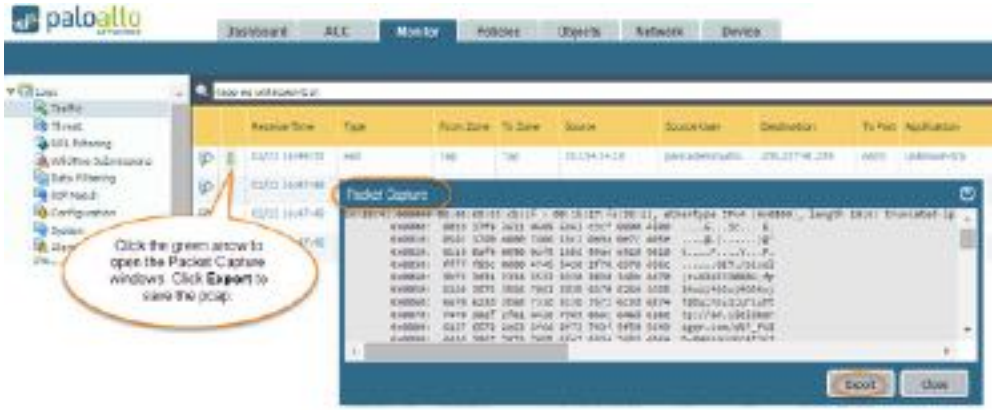
Step 2 Locate unknown application by filtering the traffic logs.

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.



 © Palo Alto Networks, Inc.

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.
	`admin@PA-200> set application dump on application <application-name> rule <rule-name>`
	For example, to capture packets for the facebook-base application that matches the security rule named rule1, run the following CLI command:
	`admin@PA-200> set application dump on application facebook-base rule rule1`
	You can also apply other filters, such as source IP address and destination IP address.

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.
Step 4	Access Facebook.com from a web browser to generate Facebook traffic and then turn off application packet capture by running the following CLI command: `admin@PA-200> set application dump off`

```
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
```

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 computer. 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.

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.

 © Palo Alto Networks, Inc.

Take a Management Interface Packet Capture (Continued)

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
```

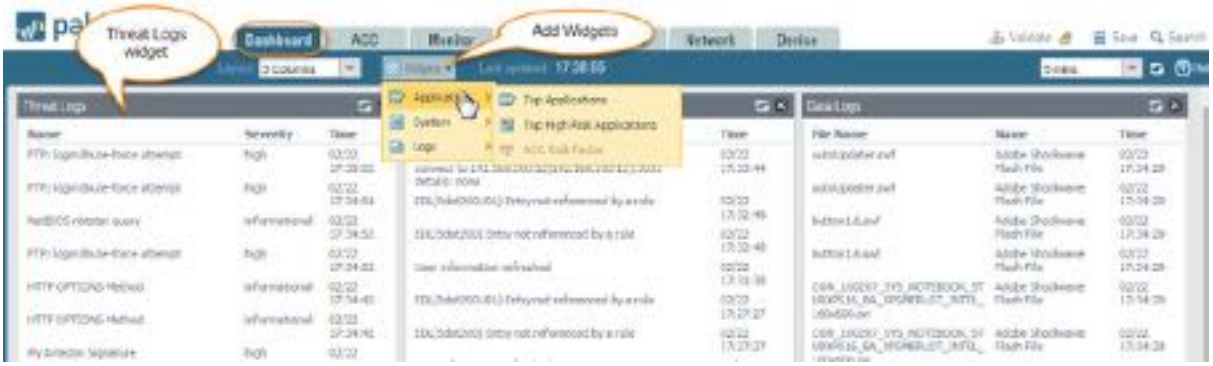
Enter the login name and password for the account on the SCP server to enable the firewall to copy the packet capture to the c:\temp-SCP folder on the SCP-enabled.

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.



View AutoFocus Threat Data for Logs to check whether logged events on the firewall pose a security risk. The AutoFocus intelligence summary shows the prevalence of properties, activities, or behaviors associated with logs in your network and on a global scale, as well as the WildFire verdict and AutoFocus tags linked to them. With an active AutoFocus subscription, you can use this information to create customized AutoFocus Alerts that track specific threats on your network.

 © Palo Alto Networks, Inc.

Monitor and Manage Logs

A log is an automatically generated, time-stamped file that provides an audit trail for system events on the firewall or network traffic events that the firewall monitors. Log entries contain *artifacts*, which are properties, activities, or behaviors associated with the logged event, such as the application type or the IP address of an attacker. Each log type records information for a separate event type. For example, the firewall generates a Threat log to record traffic that matches a spyware, vulnerability, or virus signature or a DoS attack that matches the thresholds configured for a port scan or host sweep activity on the firewall.

- Log Types and Severity Levels
- Work with Logs
- Configure Log Storage Quotas and Expiration Periods
- Schedule Log Exports to an SCP or FTP Server

Log Types and Severity Levels

You can see the following log types in the **Monitor > Logs** pages.

- Traffic Logs
- Threat Logs
- URL Filtering Logs
- WildFire Submissions Logs
- Data Filtering Logs
- Correlation Logs
- Config Logs
- System Logs
- HIP Match Logs
- Alarms Logs
- Unified Logs

Traffic Logs

Traffic logs display an entry for the start and end of each session. Each entry includes the following information: date and time; source and destination zones, addresses and ports; application name; security rule applied to the traffic flow; rule action (allow, deny, or drop); ingress and egress interface; number of bytes; and session end reason.

The Type column indicates whether the entry is for the start or end of the session. The Action column indicates whether the firewall allowed, denied, or dropped the session. A *drop* indicates the security rule that blocked the traffic was specified for *any* application, while a *deny* indicates that the rule identified a specific application. If the firewall drops traffic before identifying the application, such as when a rule drops all traffic for a specific service, the Application column displays not-applicable.

© Palo Alto Networks, Inc.

Click beside an entry to view additional details about the session, such as whether an ICMP entry aggregates multiple sessions between the same source and destination (in which case the Count column value is greater than one).

Threat Logs

Threat logs display entries when traffic matches one of the Security Profiles attached to a security rule on the firewall. Each entry includes the following information: date and time; type of threat (such as virus or spyware); threat description or URL (Name column); source and destination zones, addresses, and ports; application name; alarm action (such as allow or block); and severity level.

To see more details on individual Threat log entries:

- Click beside a threat entry to view details such as whether the entry aggregates multiple threats of the same type between the same source and destination (in which case the Count column value is greater than one).
- If you configured the firewall to Take Packet Captures, click beside an entry to access the captured packets.

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.

URL Filtering Logs

URL Filtering logs display entries for traffic that matches URL Filtering Profiles attached to security rules. For example, the firewall generates a log if a rule blocks access to specific web sites and web site categories or if you configured a rule to generate an alert when a user accesses a web site.

 © Palo Alto Networks, Inc.

WildFire Submissions Logs

The firewall forwards samples (files and emails links) to the WildFire cloud for analysis based on WildFire Analysis profiles settings (**Objects > Security Profiles > WildFire Analysis**). The firewall generates WildFire Submissions log entries for each sample it forwards after WildFire completes static and dynamic analysis of the sample. WildFire Submissions log entries include the WildFire verdict for the submitted sample.

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).
Malicious	Indicates that the entry received a WildFire analysis verdict of malicious. Samples categorized as malicious are can pose a security threat. Malware can include viruses, worms, Trojans, Remote Access Tools (RATs), rootkits, and botnets. For samples that are identified as malware, the WildFire cloud generates and distributes a signature to prevent against future exposure.

Data Filtering Logs

Data Filtering logs display entries for the security rules that help prevent sensitive information such as credit card numbers from leaving the area that the firewall protects. See Set Up Data Filtering for information on defining Data Filtering profiles.

This log type also shows information for File Blocking Profiles. For example, if a rule blocks .exe files, the log shows the blocked files.

Correlation Logs

The firewall logs a correlated event when the patterns and thresholds defined in a Correlation Object match the traffic patterns on your network. To Interpret Correlated Events and view a graphical display of the events, see Use the Compromised Hosts Widget in the ACC.

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.
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.

Severity	Description
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.

Config Logs

Config logs display entries for changes to the firewall configuration. Each entry includes the date and time, the administrator username, the IP address from where the administrator made the change, the type of client (Web, CLI, or Panorama), the type of command executed, the command status (succeeded or failed), the configuration path, and the values before and after the change.

System Logs

System logs displays entries for each system event on the firewall. Each entry includes the date and time, event severity, and event description. 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 high availability (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.

HIP Match Logs

The GlobalProtect Host Information Profile (HIP) matching enables you to collect information about the security status of the end devices accessing your network (such as whether they have disk encryption enabled). The firewall can allow or deny access to a specific host based on adherence to the HIP-based security rules you define. HIP Match logs display traffic flows that match a HIP Object or HIP Profile that you configured for the rules.

 © Palo Alto Networks, Inc.

Alarms Logs

An alarm is a firewall-generated message indicating that the number of events of a particular type (for example, encryption and decryption failures) has exceeded the threshold configured for that event type. To enable alarms and configure alarm thresholds, select **Device > Log Settings** and edit the Alarm Settings.

When generating an alarm, the firewall creates an Alarm log and opens the System Alarms dialog to display the alarm. After you **Close** the dialog, you can reopen it anytime by clicking **Alarms** () at the bottom of the web interface. To prevent the firewall from automatically opening the dialog for a particular alarm, select the alarm in the Unacknowledged Alarms list and **Acknowledge** the alarm.

Unified Logs

Unified logs are entries from the Traffic, Threat, URL Filtering, WildFire Submissions, and Data Filtering logs displayed in a single view. Unified log view enables you to investigate and filter the latest entries from different log types in one place, instead of searching through each log type separately. Click Effective Queries () in the filter area to select which log types will display entries in Unified log view.

The Unified log view displays only entries from logs that you have permission to see. For example, an administrator who does not have permission to view WildFire Submissions logs will not see WildFire Submissions log entries when viewing Unified logs. Administrative Roles define these permissions.



When you Set Up Remote Search in AutoFocus to perform a targeted search on the firewall, the search results are displayed in Unified log view.

Work with Logs

- View Logs
- Filter Logs
- Export Logs
- View AutoFocus Threat Data for Logs

View Logs

You can view the different log types on the firewall in a tabular format. The firewall locally stores all log files and automatically generates Configuration and System logs by default. To learn more about the security rules that trigger the creation of entries for the other types of logs, see Log Types and Severity Levels.

To configure the firewall to forward logs as syslog messages, email notifications, or Simple Network Management Protocol (SNMP) traps, Use External Services for Monitoring.

View Logs	
Step 1 Select a log type to view.	1. Select **Monitor > Logs**. 2. Select a log type from the list. The firewall displays only the logs you have permission to see. For example, if your administrative account does not have permission to view WildFire Submissions logs, the firewall does not display that log type when you access the logs pages. Administrative Roles define the permissions.
Step 2 (Optional) Customize the log column display.	1. Click the arrow to the right of any column header, and select **Columns**. 2. Select columns to display from the list. The log updates automatically to match your selections.
Step 3 View additional details about log entries.	• Click the spyglass ( ) for a specific log entry. The Detailed Log View has more information about the source and destination of the session, as well as a list of sessions related to the log entry. • (Threat log only) Click next to an entry to access local packet captures of the threat. To enable local packet captures, see Take Packet Captures.
Next Steps...	• Filter Logs. • Export Logs. • View AutoFocus Threat Data for Logs. • Configure Log Storage Quotas and Expiration Periods.

Filter Logs

Each log has a filter area that allows you to set a criteria for which log entries to display. The ability to filter logs is useful for focusing on events on your firewall that possess particular properties or attributes. Filter logs by artifacts that are associated with individual log entries.



Filter Logs	
Step 1 (Unified logs only) Select the log types to include in the Unified log display.	1. Click Effective Queries ( ). 2. Select one or more log types from the list (**traffic**, **threat**, **url**, **data**, and **wildfire**). 3. Click **OK**. The Unified log updates to show only entries from the log types you have selected.

 © Palo Alto Networks, Inc.

Filter Logs	
Step 2 Add a filter to the filter field.  If the value of the artifact matches the operator (such as **has** or **in**), enclose the value 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" )`.	• Click one or more artifacts (such as the application type associated with traffic and the IP address of an attacker) in a log entry. For example, click the Source **10.0.0.25** and Application **web-browsing** of a log entry to display only entries that contain both artifacts in the log (AND search). • To specify artifacts to add to the filter field, click Add Filter ( ). • To add a previously saved filter, click Load Filter ( ).
Step 3 Apply the filter to the log.	Click Apply Filter ( ). The log will refresh to display only log entries that match the current filter.
Step 4 (Optional) Save frequently used filters.	1. Click Save Filter ( ). 2. Enter a **Name** for the filter. 3. Click **OK**. You can view your saved filters by clicking Load Filter ( ).
Next Steps...	• View Logs. • Export Logs. • View AutoFocus Threat Data for Logs.

Export Logs

You can export the contents of a log type to a comma-separated value (CSV) formatted report. By default, the report contains up to 2,000 rows of log entries.

Export Logs	
Step 1 Set the number of rows to display in the report.	1. Select **Device > Setup > Management**, then edit the Logging and Reporting Settings. 2. Click the **Log Export and Reporting** tab. 3. Edit the number of **Max Rows in CSV Export** (up to 100,000 rows). 4. Click **OK**.
Step 2 Download the log.	1. Click Export to CSV ( ). A progress bar showing the status of the download appears. 2. When the download is complete, click **Download file** to save a copy of the log to your local folder. For descriptions of the column headers in a downloaded log, refer to Syslog Field Descriptions.
Next Step...	Schedule Log Exports to an SCP or FTP Server.

View AutoFocus Threat Data for Logs

With a valid AutoFocus subscription, you can view AutoFocus threat intelligence data for the following artifacts in Traffic, Threat, URL Filtering, WildFire Submissions, Data Filtering, and Unified logs:

- IP address
- URL
- User agent
- Threat name (only for threats of the subtype virus and wildfire-virus)
- Filename
- SHA-256 hash

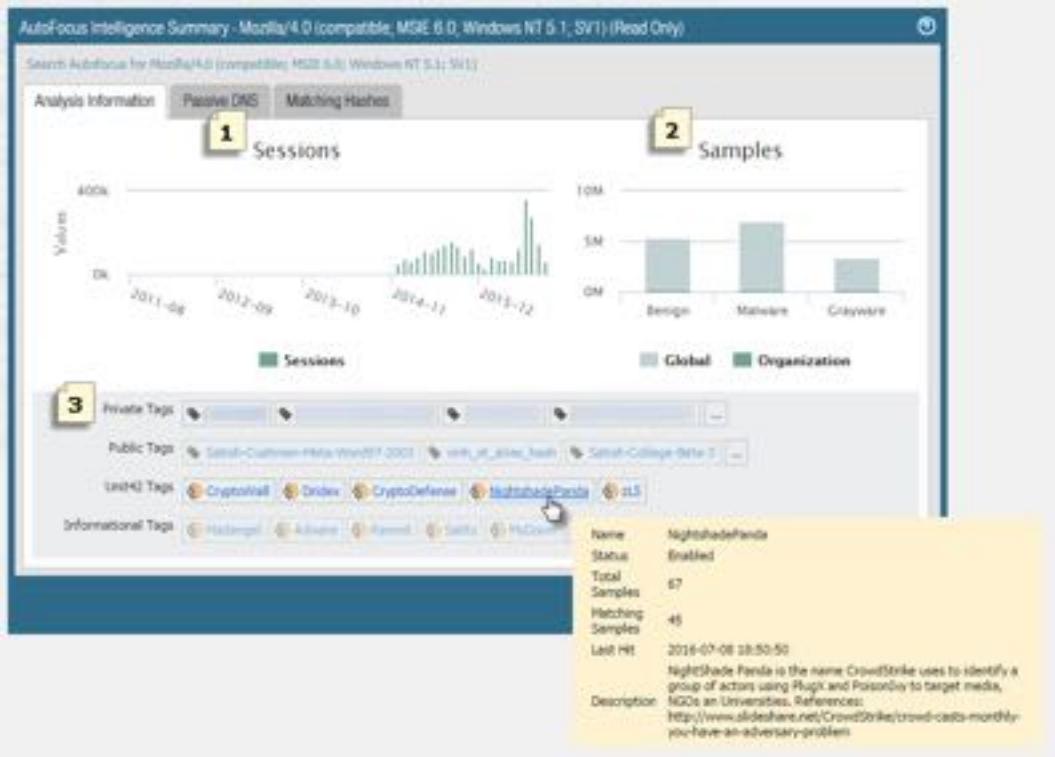
You can also open an AutoFocus search for log artifacts.

View AutoFocus Threat Data for Logs		
Step 1	Connect the firewall to AutoFocus to Enable AutoFocus Threat Intelligence.  Enable AutoFocus in Panorama to view AutoFocus threat data for all Panorama log entries, including those from firewalls that are not connected to AutoFocus and/or are running PAN-OS 7.0 and earlier release versions (**Panorama > Setup > Management > AutoFocus**).	
Step 2	Select a log type to view.	1. Select **Monitor > Logs**. 2. Select one of the following log types: **Traffic**, **Threat**, **URL Filtering**, **WildFire Submissions**, **Data Filtering**, or **Unified**.
Step 3	Open the AutoFocus Intelligence Summary for an artifact.	1. Click the drop-down ( ) for an IP address, URL, user agent, threat name (subtype: virus or wildfire-virus), filename, or SHA-256 hash in any log entry. 2. Click **AutoFocus**.

 © Palo Alto Networks, Inc.

View AutoFocus Threat Data for Logs

Step 4 View the **Analysis Information** available in AutoFocus for the artifact.



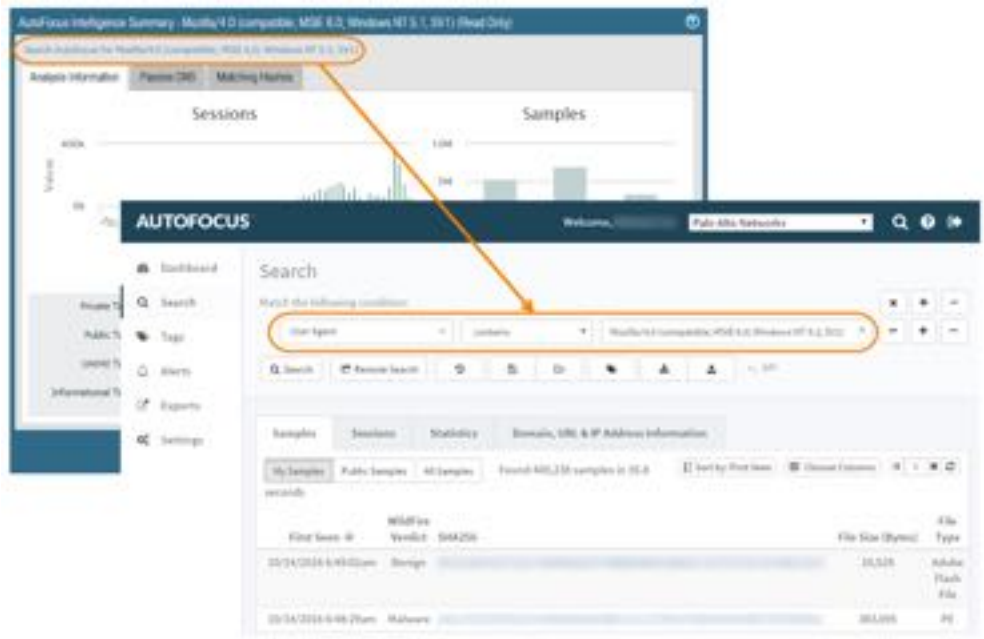
- View the number of sessions (1) logged in your firewall(s) in which the firewall detected samples associated with the artifact.
- Compare the WildFire verdicts (benign, malware, grayware) for global and organization samples (2) associated with the artifact. *Global* refers to samples from all WildFire submissions, while *organization* refers to only samples submitted to WildFire by your organization.
- Review the matching tags (3) for the artifact. AutoFocus Tags indicate whether an artifact is linked to malware or targeted attacks. Hover over a tag to view more details about the tag. Click on the ellipsis to launch AutoFocus search for the artifact. The Tags column in the AutoFocus search results displays more matching tags for the artifact.

Step 5	For an IP address, domain, or URL artifact, view passive DNS history that includes the artifact.	Click the **Passive DNS** tab. The passive DNS history is based on global DNS intelligence in AutoFocus; it is not limited to the DNS activity in your network. Passive DNS history consists of the domain request, the DNS request type, the IP address or domain returned in response to the domain request, the number of times the request was made, and the date and time the request was first seen and last seen.
Step 6	View the latest samples in your network where WildFire found the artifact.	Click the **Matching Hashes** tab, which displays the 5 most recently detected matching samples. Sample information include SHA256 hash, the file type, the date that the sample was first analyzed by WildFire, the WildFire verdict for the sample, and the date that the WildFire verdict was updated (if applicable).

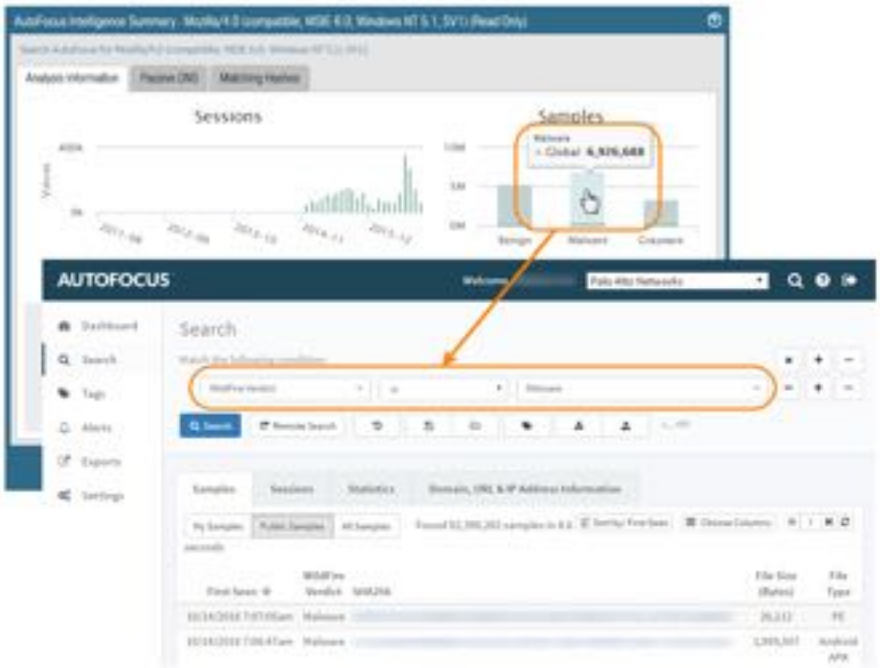
View AutoFocus Threat Data for Logs

Step 7 Launch an AutoFocus Search for firewall artifacts.

- Click the link for the log artifact. AutoFocus opens in a new browser tab, with the log artifact added as a search condition.



- Click a linked artifact in the tables or charts or click any of the matching tags to launch an AutoFocus search for it.



Step 8 Learn more about how to use AutoFocus Search to investigate threats on your network.

 © 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	Select **Log Storage** and enter a **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.  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**.

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.

Schedule Log Exports to an SCP or FTP Server	
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 **Enable FTP Passive Mode** 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**. Before establishing a connection, the firewall must accept the host key for the SCP server.  If you use a Panorama template to configure the log export schedule, you must perform this step after committing the template configuration to the firewalls. After the template commit, log in to each firewall, open the log export schedule, and click **Test SCP server connection**.
Step 12	Click **OK** and **Commit**.

 © 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.

- Report Types
- View Reports
- Configure the Report Expiration Period
- Disable Predefined Reports
- Custom Reports
- Generate Custom Reports
- Generate Botnet Reports
- Generate the SaaS Application Usage Report
- Manage PDF Summary Reports
- Generate User/Group Activity Reports
- Manage Report Groups
- Schedule Reports for Email Delivery

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 grouped into sections (types) on the right-hand side of the page: **Custom Reports**, **Application Reports**, **Traffic Reports**, **Threat Reports**, **URL Filtering Reports**, and **PDF Summary Reports**.
Step 2	Select a report to view. The reports page then displays the report for the previous day. To view reports for other days, select a date in the calendar at the bottom right of the page and select a report. If you select a report in another section, the date selection resets to the current date.
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.

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).  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**.

 © Palo Alto Networks, Inc.

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.

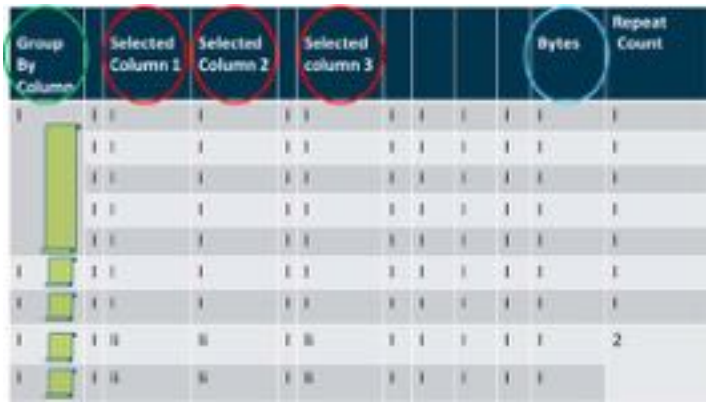
Make sure that no report group or PDF summary report includes the predefined reports you will disable. Otherwise, the firewall will render the PDF summary report or report group without any data.

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**.

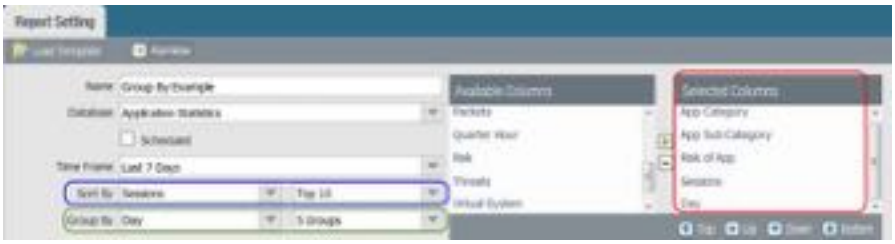
Custom Reports

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. • 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. • 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. 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. 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:  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. 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. 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

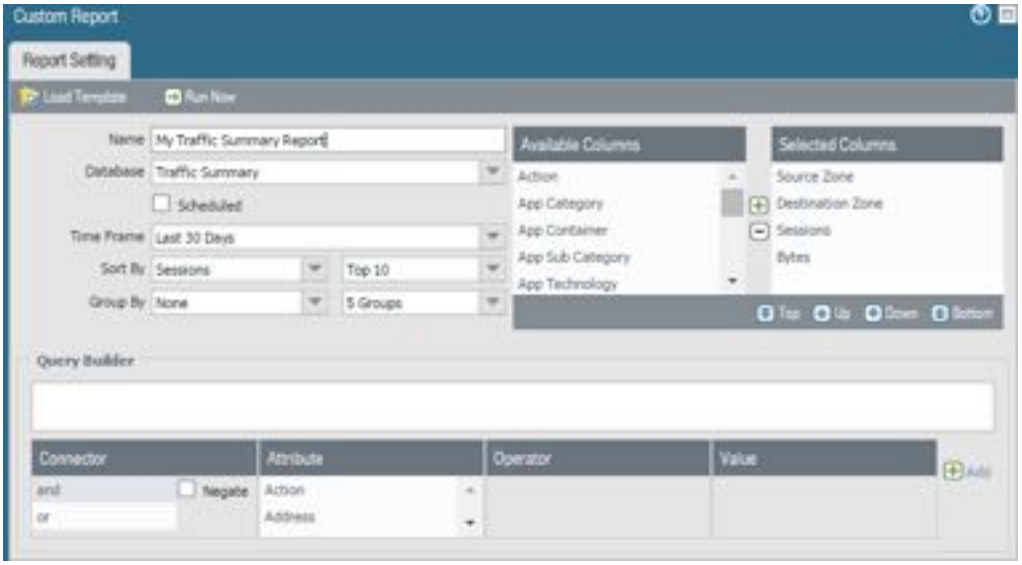
Generate Custom Reports	
Step 1	Select **Monitor > Manage Custom Reports**.
Step 2	Click **Add** and then enter a **Name** for the report. To base a report on an predefined template, click **Load Template** and choose the template. You can then edit the template and save it as a custom report.
Step 3	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:

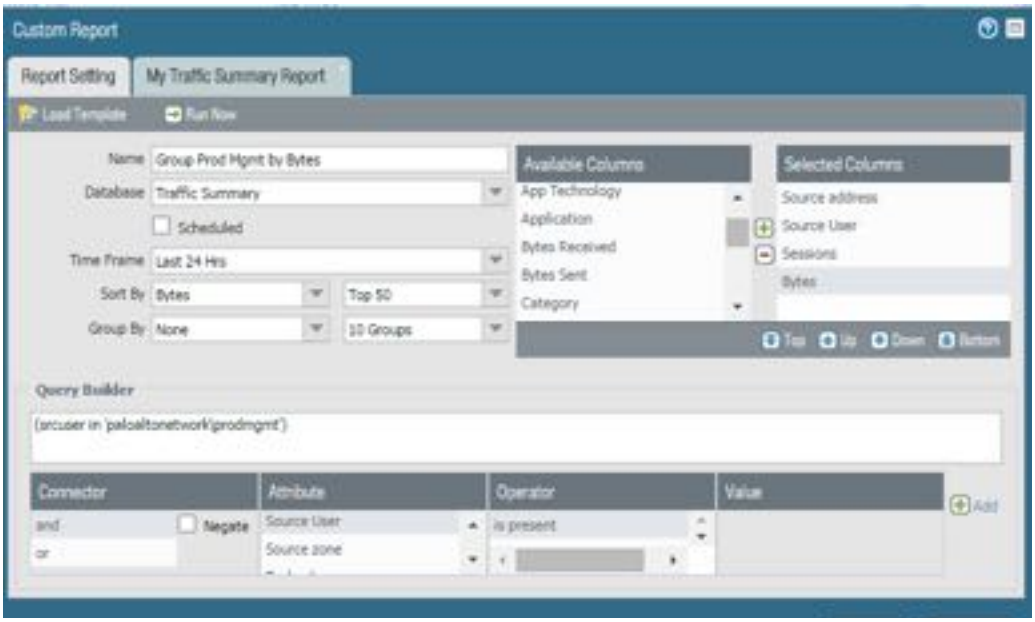
My Traffic Summary Report

ca1demo.paloaltonetworks.com : 2016/01/25 10:34:39 - 2016/02/24 10:34:38

Source Zone	Destination Zone	App Category	Application	Sessions	Bytes
Tap	Tap	general-internet	web-browsing	74.54 M	2.47 T
Tap	Tap	networking	dns	52.03 M	28.93 G
Tap	Tap	networking	ssl	18.01 M	678.13 G
Tap	Tap	general-internet	bittorrent	9.80 M	1.62 T
Tap	Tap	general-internet	google-base	4.48 M	168.99 G
Tap	Tap	unknown	insufficient-data	4.45 M	31.30 G
Tap	Tap	collaboration	facebook-base	4.09 M	99.14 G
Tap	Tap	networking	ntp	4.07 M	3.29 G
Tap	Tap	collaboration	blackboard	2.84 M	186 G
Tap	Tap	collaboration	smtp	1.92 M	172.57 G
Tap	Tap	networking	icmp	1.36 M	320.49 M
Tap	Tap	general-internet	gnutella	1.17 M	17.84 G
Tap	Tap	collaboration	myspace-base	1.10 M	35.22 G
Tap	Tap	general-internet	ping	1.06 M	86.21 M
Tap	Tap	general-internet	flash	1.01 M	168.14 G

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.

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.

 © Palo Alto Networks, Inc.

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. 2. **Enable** and define the **Count** for each type of HTTP Traffic that the report will include. 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. 3. Define the thresholds that determine whether the report will include hosts associated with traffic involving Unknown TCP or Unknown UDP applications. 4. Select the **IRC** check box to include traffic involving IRC servers. 5. Click **OK** to save the report configuration.
Step 2 Schedule the report or run it on demand.	1. Click **Report Setting** on the right side of the page. 2. Select a time interval for the report in the **Test Run Time Frame** drop-down. 3. Select the **No. of Rows** to include in the report. 4. (Optional) **Add** queries to the Query Builder to filter the report output by attributes such as source/destination IP addresses, users, or zones. For example, if you know in advance that traffic initiated from the IP address 10.3.3.15 contains no potential botnet activity, 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. 5. Select **Scheduled** to run the report daily or click **Run Now** to run the report immediately. 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.

Generate the SaaS Application Usage Report

The SaaS Application Usage PDF report is a two-part report that is based on the notion of sanctioned and unsanctioned applications. A sanctioned application is an application that you formally approve for use on your network; a SaaS application is an application that has the characteristic SaaS=yes in the applications details page in **Objects > Applications**, all other applications are considered as non-SaaS. To indicate that you have sanctioned a SaaS or non-SaaS application, you must tag it with the new predefined tag named Sanctioned. The firewall and Panorama consider any application without this predefined tag as unsanctioned for use on the network.

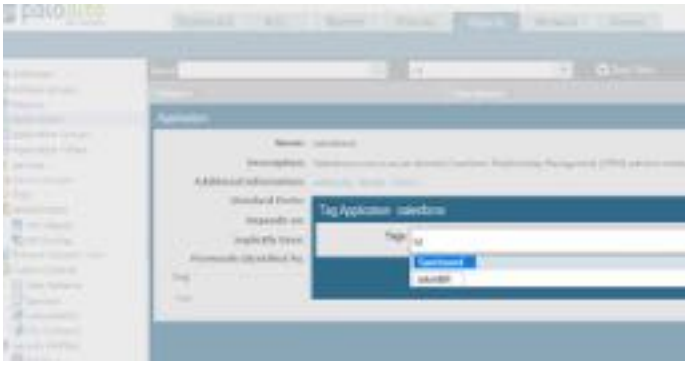
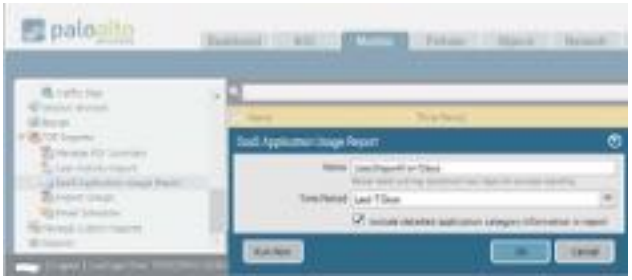
- The first part of the report (8 pages) focuses on the SaaS applications used on your network during the reporting period. It presents a comparison of sanctioned versus unsanctioned SaaS applications by total number of applications used on your network, bandwidth consumed by these applications, and the number of users using these applications. This first part of the report also highlights the top SaaS application subcategories listed in order by maximum number of applications used, the number of users, and the amount of data (bytes) transferred in each application subcategory.
- The second part of the report focuses on the detailed browsing information for SaaS and non-SaaS applications for each application subcategory listed in the first-part of the report. For each application in a subcategory, it also includes information about the top users who transferred data, the top blocked or alerted file types, and the top threats for each application. In addition, this section of the report tallies samples for each application that the firewall submitted for WildFire analysis, and the number of samples determined to be benign and malicious.

Use the insights from this report to consolidate the list of business-critical and approved SaaS applications and to enforce policies for controlling unsanctioned applications that pose an unnecessary risk for malware propagation and data leaks.



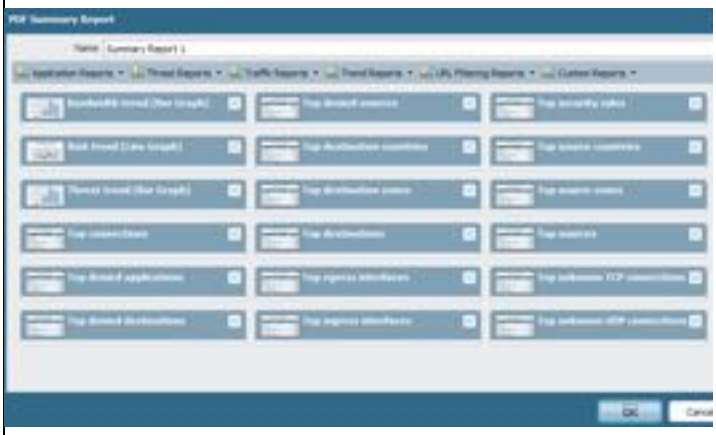
The predefined SaaS application usage report introduced in PAN-OS 7.0 is still available as a daily report that lists the top 100 SaaS applications (with the SaaS application characteristic, SaaS=yes) running on your network on a given day.

 © Palo Alto Networks, Inc.

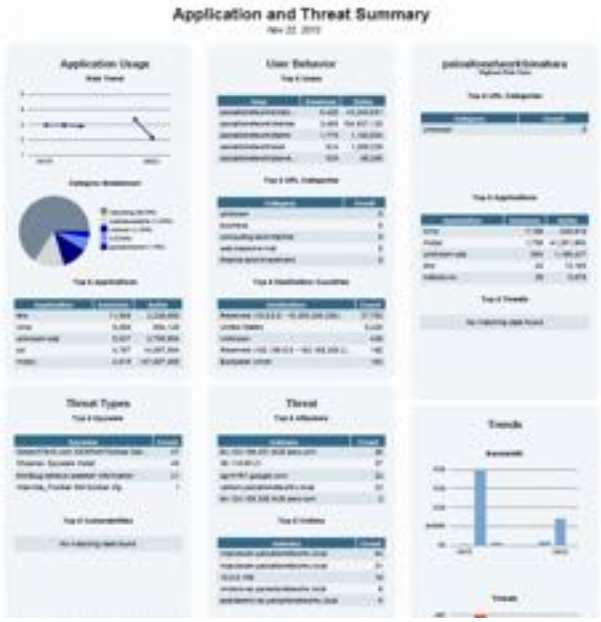
Generate the SaaS Application Usage Report	
Step 1 Tag applications that you approve for use on your network as Sanctioned. The accuracy of the report depends on whether you have tagged an application as Sanctioned. You can tag both SaaS and non-SaaS applications as Sanctioned; the detailed browsing section of the SaaS Application Usage report displays whether the application is SaaS and whether it is sanctioned.	1. Select **Object > Applications**. 2. Click the application **Name** to edit an application and select **Edit** in the Tag section. 3. Select **Sanctioned** from the **Tags** drop-down. You must use the predefined **Sanctioned** tag (with the azure colored background). If you use any other tag to indicate that you sanctioned an application, the firewall will fail to recognize the tag and the report will be inaccurate.  4. Click **OK** and **Close** to exit all open dialogs.
Step 2 Configure the SaaS Application Usage report.	1. Select **Monitor > PDF Reports > SaaS Application Usage**. 2. Click **Add**, enter a **Name**, and select a **Time Period** for the report (default is **Last 7 Days**). By default, the report includes detailed information on the top SaaS and non-SaaS application subcategories, which can make the report large by page count and file size. Clear the **Include detailed application category information in report** check box if you want to reduce the file size and restrict the page count to eight pages.  3. To generate the report on-demand, click **Run Now**. Make sure that the pop-up blocker is disabled on your browser because the report opens in a new tab. 4. Click **OK** to save your changes.
Step 3 Schedule Reports for Email Delivery.	On the PA-200, PA-500, and PA-2000 Series firewalls, the SaaS Application Usage report is not sent as a PDF attachment in the email. Instead, the email includes a link that you must click to open the report in a web browser.

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**. 2. Click **Add** and then enter a **Name** for the report. 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.  • To remove an element from the report, click the **x** icon or clear the selection from the drop-down for the appropriate report group. • To rearrange the reports, drag and drop the element icons to another area of the report. 4. Click **OK t**o save the report. 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 except for the **Browsing Summary by URL Category** and **Browse time calculations**, which only the User Activity report includes.

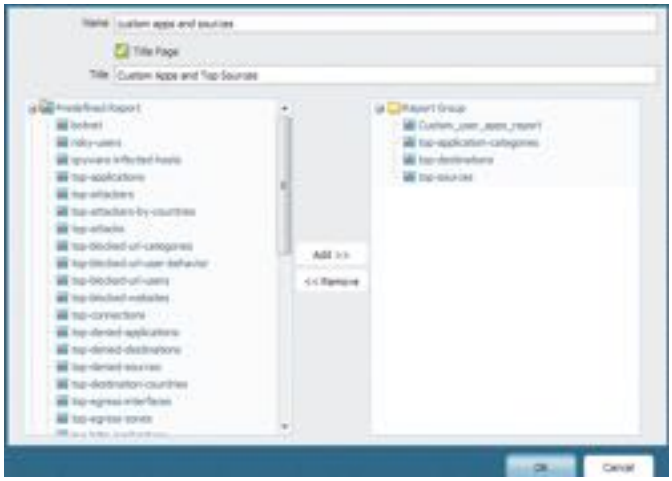
You must configure User-ID on the firewall to access the list of users and user groups.

Generate User/Group Activity Reports	
Step 1 Configure the browse times and number of logs for User/Group Activity reports. 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. 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. 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. 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. 5. Click **OK** to save your changes.
Step 2 Generate the User/Group Activity report.	1. Select **Monitor > PDF Reports > User Activity Report**. 2. Click **Add** and then enter a **Name** for the report. 3. Create the report: • User Activity Report—Select **User** and enter the **Username** or **IP address** (IPv4 or IPv6) of the user. • Group Activity Report—Select **Group** and select the **Group Name** of the user group. 4. Select the **Time Period** for the report. 5. Optionally, select the **Include Detailed Browsing** check box (default is cleared) 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 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. You must set up a **Report Group** to email report(s).	1. Create an Email server profile. 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. a. Select **Monitor > Report Group**. b. Click **Add** and then enter a **Name** for the report group. c. (Optional) Select **Title Page** and add a **Title** for the PDF output. d. Select reports from the left column and click **Add** to move each report to the report group on the right.  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. 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. e. Click **OK** to save the settings. 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 Email Addresses** field allows you to send this report exclusively to the specified recipients. When you add recipients to the field, the firewall does not send the report 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.
Step 7	Click **OK** and **Commit**.

 © 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 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 Manage Reporting, 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. 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. • 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. • 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. 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**. 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. 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. 3. Perform the following steps for each log type and each severity level or WildFire verdict: 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.) 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. To trigger log generation and forwarding, the rules require certain Security Profiles according to log type: • Traffic logs—No security profile is necessary; the traffic only needs to match a specific security rule. • Threat logs—The traffic must match any security profile assigned to a security rule. • 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: 1. Select **Policies > Security** and click the rule. 2. Select the **Actions** tab and select the **Log Forwarding** profile you just created. 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. 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 for System, Config, HIP Match, and Correlation logs.	1. Select **Device > Log Settings**. 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 editing the section. a. Select the **Panorama** check box if you want to aggregate System, Config, and HIP Match logs on Panorama. Optionally, you can then configure Panorama to forward the logs to the external services. Panorama generates Correlation logs based on the firewall logs it receives, rather than aggregating Correlation logs from firewalls. 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**. 2. Select the **Slot** and **Interface Name**. 3. For the **Interface Type**, select **Log Card**. 4. Enter the **IP Address**, **Default Gateway**, and (for IPv4 only) **Netmask**. 5. Select **Advanced** and specify the **Link Speed**, **Link Duplex**, and **Link State**. 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). 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. 2. Verify the log destinations you configured are receiving firewall logs: • 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. • Email server—Verify that the specified recipients are receiving logs as email notifications. • Syslog server—Refer to the documentation for your syslog server to verify it is receiving logs as syslog messages. • 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.  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**. 2. Click **Add** and then enter a **Name** for the profile. 3. If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where this profile is available. 4. For each Simple Mail Transport Protocol (SMTP) server (email server), click **Add** and define the following information: • **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. • **Email Display Name**—The name to show in the From field of the email. • **From**—The email address from which the firewall sends emails. • **To**—The email address to which the firewall sends emails. • **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. • **Email Gateway**—The IP address or hostname of the SMTP gateway to use for sending emails. 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. 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. a. Select **Objects > Log Forwarding**, click **Add**, and enter a **Name** to identify the profile. b. For each log type and each severity level or WildFire verdict, select the Email server profile and click **OK**. 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**. 2. For System and Correlation logs, click each Severity level, select the **Email** server profile, and click **OK**. 3. For Config and HIP Match logs, edit the section, select the **Email** server profile, and click **OK**. 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 firewalls 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

Configure Syslog Monitoring

To Use Syslog for Monitoring a Palo Alto Networks firewall, create a Syslog server profile and assign it to the 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.  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**. 2. Click **Add** and enter a **Name** for the profile. 3. If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where this profile is available. 4. For each syslog server, click **Add** and enter the information that the firewall requires to connect to it: • **Name**—Unique name for the server profile. • **Syslog Server**—IP address or fully qualified domain name (FQDN) of the syslog server. • **Transport**—Select TCP, UDP, or SSL as the method of communication with the syslog server. • **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. • **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. • **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. 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. 6. Click **OK** to save the server profile.

 © Palo Alto Networks, Inc.

Configure Syslog Monitoring (Continued)		
Step 2	Configure syslog forwarding for Traffic, Threat, and WildFire Submission logs.	1. Create a log forwarding profile. a. Select **Objects > Log Forwarding**, click **Add**, and enter a **Name** to identify the profile. b. For each log type and each severity level or WildFire verdict, select the **Syslog** server profile and click **OK**. 2. Assign the log forwarding profile to security rules.
Step 3	Configure syslog forwarding for System, Config, HIP Match, and Correlation logs.	1. Select **Device > Log Settings**. 2. For System and Correlation logs, click each Severity level, select the **Syslog** server profile, and click **OK**. 3. For Config, HIP Match, and Correlation logs, edit the section, select the **Syslog** server profile, and click **OK**.
Step 4	(Optional) Configure the header format of syslog messages. The log data includes the unique identifier of the firewall 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. This is a global setting and applies to all syslog server profiles configured on the firewall.	1. Select **Device > Setup > Management** and edit the Logging and Reporting Settings. 2. Select the **Log Export and Reporting** tab and select the **Syslog HOSTNAME Format**: • **FQDN** (default)—Concatenates the hostname and domain name defined on the sending firewall. • **hostname**—Uses the hostname defined on the sending firewall. • **ipv4-address**—Uses the IPv4 address of the firewall interface used to send logs. By default, this is the MGT interface. • **ipv6-address**—Uses the IPv6 address of the firewall interface used to send logs. By default, this is the MGT interface. • **none**—Leaves the hostname field unconfigured on the firewall. There is no identifier for the firewall that sent the logs. 3. Click **OK** to save your changes.

Configure Syslog Monitoring (Continued)		
Step 5	Create a certificate to secure syslog communication over SSL. Required only if the syslog server uses client authentication. The syslog server uses the certificate to verify that the firewall is authorized to communicate with the syslog server. Ensure the following conditions are met: • The private key must be available on the sending firewall; the keys can't reside on a Hardware Security Module (HSM). • The subject and the issuer for the certificate must not be identical. • The syslog server and the sending firewall must have certificates that the same trusted certificate authority (CA) signed. Alternatively, you can generate a self-signed certificate on the firewall, export the certificate from the firewall, and import it in to the syslog server.	1. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Generate**. 2. Enter a **Name** for the certificate. 3. In the **Common Name** field, enter the IP address of the firewall sending logs to the syslog server. 4. In **Signed by**, select the trusted CA or the self-signed CA that the syslog server and the sending firewall both trust. The certificate can't be a **Certificate Authority** nor an **External Authority** (certificate signing request [CSR]). 5. Click **Generate**. The firewall generates the certificate and key pair. 6. Click the certificate Name to edit it, select the **Certificate for Secure Syslog** check box, and click **OK**.
Step 6	Commit your changes and review the logs on the syslog server.	1. Click **Commit**. 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 firewalls 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 firewalls do not currently implement.



WildFire Submissions logs are a subtype of Threat log and use the same syslog format.

- Traffic Log Fields
- Threat Log Fields
- HIP Match Log Fields
- Config Log Fields
- System Log Fields
- Correlated Events Log Fields
- Custom Log/Event Format

 © Palo Alto Networks, Inc.

▲ Escape Sequences

Traffic Log Fields

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 firewall 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 • Start—session started • End—session ended • Drop—session dropped before the application is identified and there is no rule that allows the session. • 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

Field Name	Description
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: • 0x80000000—session has a packet capture (PCAP) • 0x02000000—IPv6 session • 0x01000000—SSL session was decrypted (SSL Proxy) • 0x00800000—session was denied via URL filtering • 0x00400000—session has a NAT translation performed (NAT) • 0x00200000—user information for the session was captured via the captive portal (Captive Portal) • 0x00080000—X-Forwarded-For value from a proxy is in the source user field • 0x00040000—log corresponds to a transaction within a http proxy session (Proxy Transaction) • 0x00008000—session is a container page access (Container Page) • 0x00002000—session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above. • 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: • Allow—session was allowed by policy • Deny—session was denied by policy • Drop—session was dropped silently • Drop ICMP—session was silently dropped with an ICMP unreachable message to the host or application • Reset both—session was terminated and a TCP reset is sent to both the sides of the connection • Reset client—session was terminated and a TCP reset is sent to the client • 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 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 Available on all models except the PA-4000 Series

 © Palo Alto Networks, Inc.

Field Name	Description
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-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 Available on all models except the PA-4000 Series
Packets Received (pkts_received)	Number of server-to-client packets for the session Available on all models except the PA-4000 Series

Field Name	Description

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):
	• threat—The firewall detected a threat associated with a reset, drop, or block (IP address) action.
	• policy-deny—The session matched a security rule with a deny or drop action.
	• decrypt-cert-validation—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when the session uses client authentication or when the session uses a server certificate with any of the following conditions: expired, untrusted issuer, unknown status, or status verification time-out. This session end reason also displays when the server certificate produces a fatal error alert of type bad_certificate, unsupported_certificate, certificate_revoked, access_denied, or no_certificate_RESERVED (SSLv3 only).
	• decrypt-unsupport-param—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when the session uses an unsupported protocol version, cipher, or SSH algorithm. This session end reason is displays when the session produces a fatal error alert of type unsupported_extension, unexpected_message, or handshake_failure.
	• decrypt-error—The session terminated because you configured the firewall to block SSL forward proxy decryption or SSL inbound inspection when firewall resources or the hardware security module (HSM) were unavailable. This session end reason is also displayed when you configured the firewall to block SSL traffic that has SSH errors or that produced any fatal error alert other than those listed for the decrypt-cert-validation and decrypt-unsupport-param end reasons.
	• tcp-rst-from-client—The client sent a TCP reset to the server.
	• tcp-rst-from-server—The server sent a TCP reset to the client.
	• 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.
	• tcp-fin—One host or both hosts in the connection sent a TCP FIN message to close the session.
	• tcp-reuse—A session is reused and the firewall closes the previous session.
	• decoder—The decoder detects a new connection within the protocol (such as HTTP-Proxy) and ends the previous connection.
	• aged-out—The session aged out.
	• unknown—This value applies in the following situations:
	  • Session terminations that the preceding reasons do not cover (for example, a `clear session all` command).
	  • 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.
	  • In Panorama, logs received from firewalls for which the PAN-OS version does not support session end reasons will have a value of `unknown`.
	• n/a—This value applies when the traffic log type is not `end`.

 © Palo Alto Networks, Inc.

Field Name	Description
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	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. 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: CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.
Action Source (action_source)	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.

Threat Log Fields

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 firewall that generated the log
Type (type)	Specifies type of log; values are traffic, threat, config, system and hip-match

Field Name	Description
Subtype (subtype)	Subtype of threat log. Values include the following: • data—Data pattern matching a Data Filtering profile. • file—File type matching a File Blocking profile. • flood—Flood detected via a Zone Protection profile. • packet—Packet-based attack protection triggered by a Zone Protection profile. • scan—Scan detected via a Zone Protection profile. • spyware—Spyware detected via an Anti-Spyware profile. • url—URL filtering log. • virus—Virus detected via an Antivirus profile. • vulnerability—Vulnerability exploit detected via a Vulnerability Protection profile. • 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. • 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
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

 © Palo Alto Networks, Inc.

Field Name	Description
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: • 0x80000000—session has a packet capture (PCAP) • 0x02000000—IPv6 session • 0x01000000—SSL session was decrypted (SSL Proxy) • 0x00800000—session was denied via URL filtering • 0x00400000—session has a NAT translation performed (NAT) • 0x00200000—user information for the session was captured via the captive portal (Captive Portal) • 0x00080000—X-Forwarded-For value from a proxy is in the source user field • 0x00040000—log corresponds to a transaction within a http proxy session (Proxy Transaction) • 0x00008000—session is a container page access (Container Page) • 0x00002000—session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above • 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 alert, allow, deny, drop, drop-all-packets, reset-client, reset-server, reset-both, block-url. • Alert—threat or URL detected but not blocked • Allow—flood detection alert • Deny—flood detection mechanism activated and deny traffic based on configuration • Drop—threat detected and associated session was dropped • Drop-all-packets—threat detected and session remains, but drops all packets • Reset-client—threat detected and a TCP RST is sent to the client • Reset-server—threat detected and a TCP RST is sent to the server • Reset-both—threat detected and a TCP RST is sent to both the client and the server • 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 The actual URI when the subtype is URL File name or file type when the subtype is file File name when the subtype is virus File name when the subtype is WildFire

Field Name	Description
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: • 8000 – 8099—scan detection • 8500 – 8599—flood detection • 9999—URL filtering log • 10000 – 19999—sypware phone home detection • 20000 – 29999—spyware download detection • 30000 – 44999—vulnerability exploit detection • 52000 – 52999—filetype detection • 60000 – 69999—data filtering detection • 100000 – 2999999—virus detection • 3000000 – 3999999—WildFire signature feed • 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
Direction (direction)	Indicates the direction of the attack, client-to-server or server-to-client: • 0—direction of the threat is client to server • 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. 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 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. 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.

 © Palo Alto Networks, Inc.

Field Name	Description
URL Index (url_idx)	Used in URL Filtering and WildFire subtypes. 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. 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. 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. Specifies the type of file that the firewall forwarded for WildFire analysis.
X-Forwarded-For (xff)	Only for the URL Filtering subtype; all other types do not use this field. 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. 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. 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. 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. 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. Identifies the analysis request on the WildFire cloud or the WildFire appliance.

Field Name	Description
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	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. 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: CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.

HIP Match Log Fields

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 firewall 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

 © 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
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	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. 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: CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.

Config Log Fields

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

Field Name	Description
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. It contains the full xpath before the configuration change.
After Change Detail (after_change_detail)	This field is in custom logs only; it is not in the default format. It contains the full xpath after the configuration change.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	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. 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: CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.

System Log Fields

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 firewall 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

 © Palo Alto Networks, Inc.

Field Name	Description
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)	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. 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: CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.
Device Name (device_name)	The hostname of the firewall on which the session was logged.

Correlated Events Log Fields

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
Log ID(logid)	Time the log was received at the management plane
ID (id)	Serial number of the device that generated the log
Match OID (match_oid)	Type of log; values are traffic, threat, config, system and hip-match
Object ID (objectid)	Name of the object associated with the system event

Field Name	Description
Version (version)	The version of the Correlation objects content update, as pushed by Palo Alto Networks.
Virtual System (vsys)	Virtual System associated with the configuration log
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	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. 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: CLI command in configure mode: **`show readonly dg-meta-data`** API query: `/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show>`
Window (window)	
Source User (srcuser)	Username of the user who initiated the event.
Source (src)	IP address of the user who initiated the event.
Last Update Time (last_update_time)	The last time the events in the correlated event were updated with more information.
Severity (severity)	Severity associated with the event; values are informational, low, medium, high, critical
Match Time (match_time)	The time that the event match was recorded.
Object Name (objectname)	Name of the correlation object that was matched on
Summary (summary)	A summary statement that indicates 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

 © Palo Alto Networks, Inc.

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.

SNMP Monitoring and Traps

The following topics describe how Palo Alto Networks firewalls, Panorama, and WF-500 appliances implement Simple Network Management Protocol (SNMP), and the procedures to configure SNMP monitoring and trap delivery.

- SNMP Support
- Use an SNMP Manager to Explore MIBs and Objects
- Enable SNMP Services for Firewall-Secured Network Elements
- Monitor Statistics Using SNMP
- Forward Traps to an SNMP Manager
- Supported MIBs

SNMP Support

You can use a Simple Network Management Protocol (SNMP) manager to monitor event-driven alerts and operational statistics for the firewall, Panorama, or WF-500 appliance and for the traffic they process. The statistics and traps can help you identify resource limitations, system changes or failures, and malware attacks. You configure alerts by forwarding log data as traps, and enable the delivery of 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.



When an event triggers SNMP trap generation (for example, an interface goes down), the firewall, Panorama virtual appliance, M-Series appliance, and WF-500 appliance respond by updating the corresponding SNMP object (for example, the interfaces MIB) instead of waiting for the periodic update of all objects that occurs every ten seconds. This ensures that your SNMP manager displays the latest information when polling an object to confirm an event.

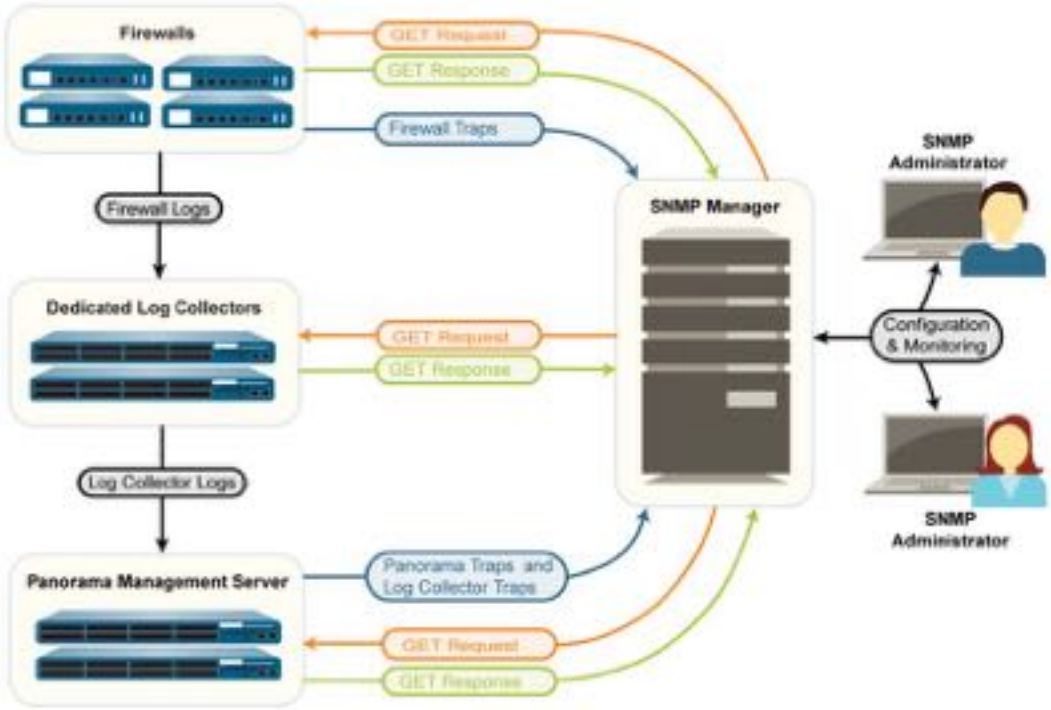
The firewall, Panorama, and WF-500 appliance support SNMP Version 2c and Version 3. Decide which to use based on the version that other devices in your network support and on your network security requirements. SNMPv3 is more secure and enables more granular access control for system statistics than SNMPv2c. The following table summarizes the security features of each version. You select the version and configure the security features when you Monitor Statistics Using SNMP and Forward Traps to an SNMP Manager.

SNMP Version	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

 © Palo Alto Networks, Inc.

Figure: SNMP Implementation 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 firewall, Panorama, or WF-500 appliance. 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.

Figure: SNMP Implementation



Use an SNMP Manager to Explore MIBs and Objects

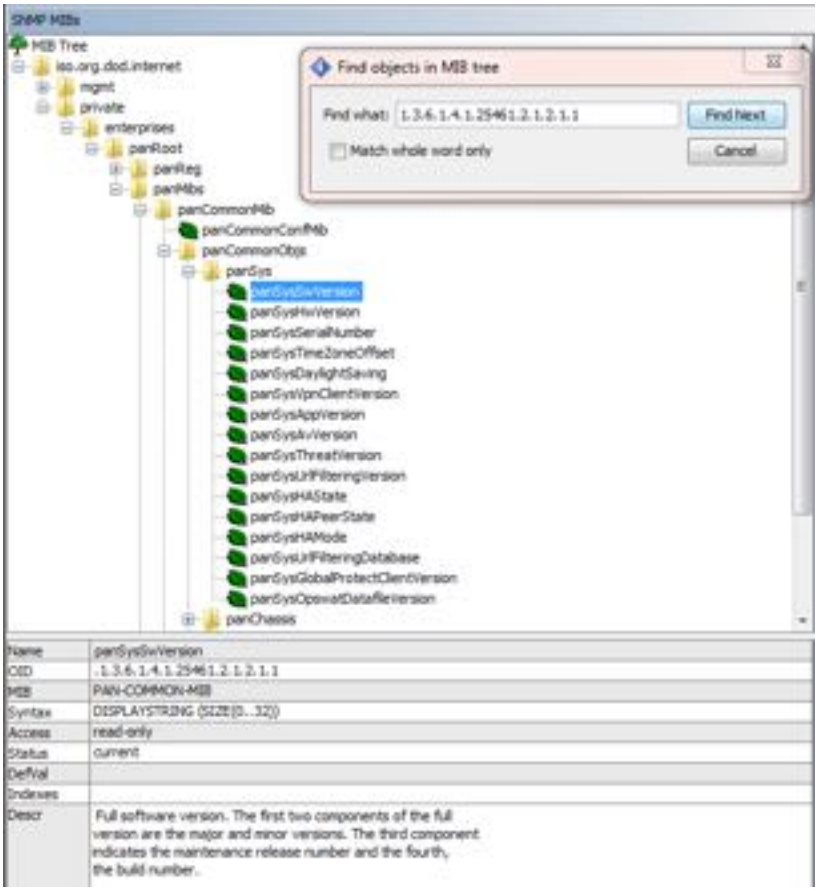
To use SNMP for monitoring Palo Alto Networks firewalls, Panorama, or WF-500 appliances, you must first load the Supported MIBs into your SNMP manager and determine which object identifiers (OIDs) correspond to the system statistics and traps you want to monitor. The following topics provide an overview 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 System Statistic or Trap

Identify a MIB Containing a Known OID

If you already know the OID for a particular SNMP object (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.
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.

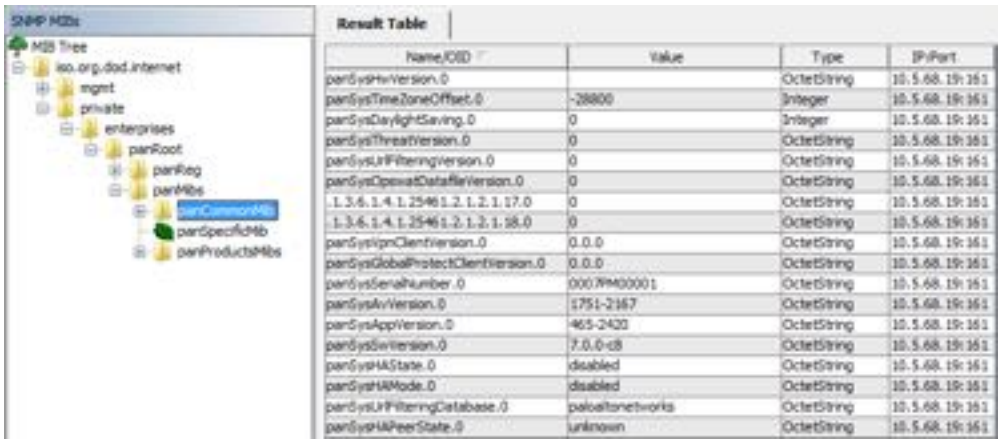


Step 3	Optionally, Walk a MIB to display all its objects.

Walk a MIB

If you want to see which SNMP objects (system 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 firewalls, Panorama, and WF-500 appliance 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 statistics:

 © Palo Alto Networks, Inc.



Identify the OID for a System Statistic or Trap

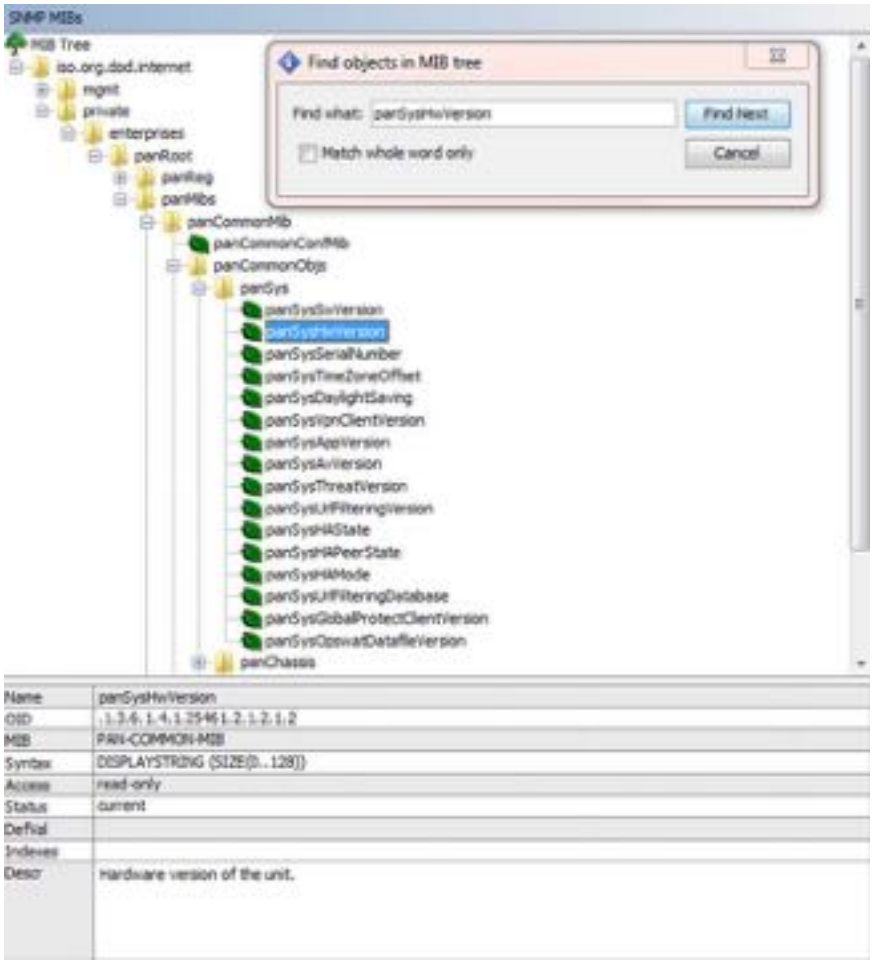
To use an SNMP manager for monitoring Palo Alto Networks firewalls, Panorama, or WF-500 appliances, you must know the OIDs of the system statistics and traps you want to monitor.

Identify the OID for a 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 hardware version information. The panCommonEventEventsV2 MIB contains all the traps that Palo Alto Networks firewalls, Panorama, and WF-500 appliances 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}`

Identify the OID for a 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.



 © Palo Alto Networks, Inc.

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 firewalls, Panorama, or WF-500 appliances. For details, see Monitor 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**. 2. Enter a **Name** to identify the application group. 3. Click **Add**, type `snmp`, and select **snmp** and **snmp-trap** from the drop-down. 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**. 2. In the **General** tab, enter a **Name** for the rule. 3. In the **Source** and **Destination** tabs, click **Add** and enter a **Source Zone** and a **Destination Zone** for the traffic. 4. In the **Applications** tab, click **Add**, type the name of the applications group you just created, and select it from the drop-down. 5. In the **Actions** tab, verify that the **Action** is set to **Allow**, and then click **OK** and **Commit**.

Monitor Statistics Using SNMP

The statistics that a Simple Network Management Protocol (SNMP) manager collects from Palo Alto Networks firewalls can help you gauge the health of your network (systems 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 firewalls (using SET messages), only to collect statistics from them (using GET messages).
For details on how SNMP is implemented for Palo Alto Networks firewalls, see SNMP Support.

Monitor Statistics Using SNMP	
Step 1 Configure the SNMP Manager to get statistics from firewalls.	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. 1. To enable the SNMP manager to interpret firewall statistics, load the Supported MIBs for Palo Alto Networks firewalls and, if necessary, compile them. 2. For each firewall 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 firewall. Note that all Palo Alto Networks firewalls use port 161. The SNMP manager can use the same or different connection and authentication settings for multiple firewalls. The settings must match those you define when you configure SNMP on the firewall (see Step 3). For example, if you use SNMPv2c, the community string you define when configuring the firewall must match the community string you define in the SNMP manager for that firewall. 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. 4. Configure the SNMP manager to monitor the desired OIDs.
Step 2 Enable SNMP traffic on a firewall interface. This is the interface that will receive statistics requests from the SNMP manager. PAN-OS doesn't synchronize management (MGT) interface settings for firewalls in a high availability (HA) configuration. You must configure the interface for each HA peer.	Perform this step in the firewall web interface. • To enable SNMP traffic on the MGT interface, select **Device > Setup > Management**, edit the Management Interface Settings, select **SNMP**, and then click **OK** and **Commit**. • To enable SNMP traffic on any other interface, create an interface management profile for SNMP services and assign the profile to the interface that will receive the SNMP requests. The interface type must be Layer 3 Ethernet.

 © Palo Alto Networks, Inc.

Monitor Statistics Using SNMP (Continued)		
Step 3	Configure the firewall to respond to statistics requests from an SNMP manager. PAN-OS doesn't synchronize SNMP response settings for firewalls 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**. 2. Select the SNMP **Version** and configure the authentication values as follows. For version details, see SNMP Support. • **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. As a best practice, don't use the default community string `public`; it's well known and therefore not secure. • **V3**—Create at least one SNMP view group and one user. User accounts and views provide authentication, privacy, and access control when firewalls forward traps and SNMP managers get firewall statistics. – **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**. – **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. 3. Click **OK** and **Commit**.
Step 4	Monitor the firewall statistics in an SNMP manager.	Refer to the documentation of your SNMP manager for details. 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 firewalls) or to threats (traffic that matches a firewall security rule) that require immediate attention.

To see the list of traps that Palo Alto Networks firewalls 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 firewalls implement SNMP, see SNMP Support.

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 firewalls 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. The profile defines how the firewall accesses the SNMP managers (trap servers). You can define up to four SNMP managers for each profile. Optionally, configure separate SNMP Trap server profiles for different log types, severity levels, and WildFire verdicts.	1. Log in to the firewall web interface. 2. Select **Device > Server Profiles > SNMP Trap**. 3. Click **Add** and enter a **Name** for the profile. 4. If the firewall has more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where this profile is available. 5. Select the SNMP **Version** and configure the authentication values as follows. For version details, see SNMP Support. • **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. As a best practice, don't use the default community string `public`; it's well known and therefore not secure. • **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 firewall (you can leave the field blank to use the firewall 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. 6. Click **OK** to save the server profile.
Step 3 Configure log forwarding.	1. Configure the destinations of Traffic, Threat, and WildFire traps: a. Create a log forwarding profile. For each log type and each severity level or WildFire verdict, select the **SNMP Trap** server profile. b. Assign the log forwarding profile to security rules. The rules will trigger trap generation and forwarding. 2. Configure the destinations for System, Config, HIP Match, and Correlation logs. For each log (trap) type and severity level, select the **SNMP Trap** server profile. 3. Click **Commit**.
Step 4 Monitor the traps in an SNMP manager.	Refer to the documentation of your SNMP manager. 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.

 © Palo Alto Networks, Inc.

Supported MIBs

The following table lists the Simple Network Management Protocol (SNMP) management information bases (MIBs) that Palo Alto Networks firewalls, Panorama, and WF-500 appliances support. You must load these MIBs into your SNMP manager to monitor the objects (system 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.  Palo Alto Networks firewalls, Panorama, and WF-500 appliances 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 IF-MIB HOST-RESOURCES-MIB ENTITY-MIB ENTITY-SENSOR-MIB ENTITY-STATE-MIB IEEE 802.3 LAG MIB LLDP-V2-MIB.my BFD-STD-MIB
Enterprise—You can download the enterprise MIBs from the Palo Alto Networks Technical Documentation portal.	PAN-COMMON-MIB.my PAN-GLOBAL-REG-MIB.my PAN-GLOBAL-TC-MIB.my PAN-LC-MIB.my PAN-PRODUCT-MIB.my PAN-ENTITY-EXT-MIB.my 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 systems 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 firewalls, Panorama, and WF-500 appliances support only the following object groups:

Object Group	Description
system	Provides system 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.

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 firewalls, Panorama, and WF-500 appliances 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 example, checking the current CPU load (hrProcessorLoad object) can help you troubleshoot performance issues on the firewall.

Palo Alto Networks firewalls, Panorama, and WF-500 appliances 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 system 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 system (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 firewalls, Panorama, and WF-500 appliances support only portions of the entPhysicalTable group:

Object	Description
entPhysicalIndex	A single namespace that includes disk slots and disk drives.
entPhysicalDescr	The component description.

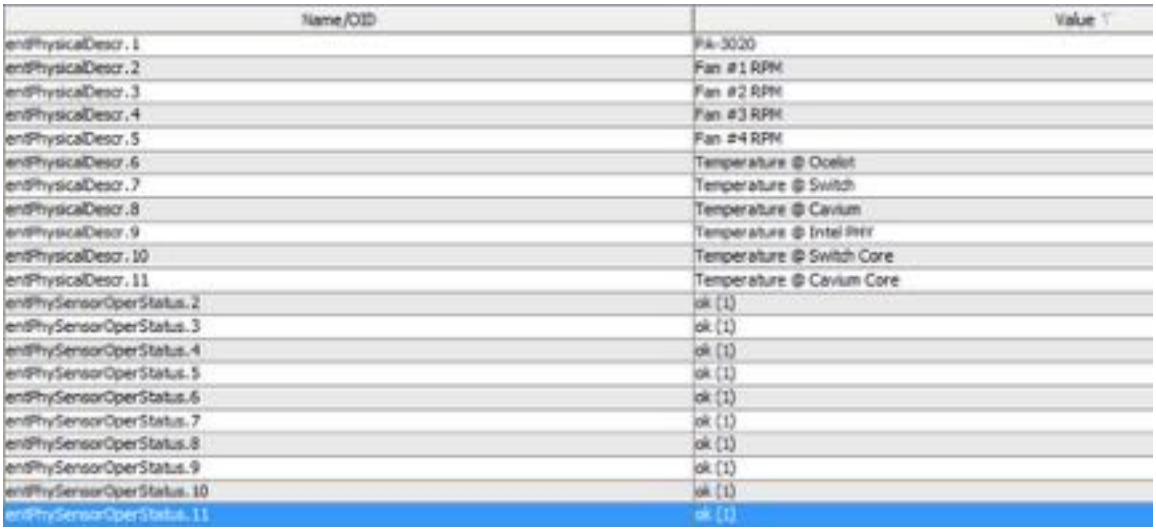
 © Palo Alto Networks, Inc.

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 system (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:



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 platforms. Use the ENTITY-MIB for the targeted platform to match the value to the description.

Palo Alto Networks firewalls, Panorama, and WF-500 appliances support only portions of the entPhySensorTable group. The supported portions vary by platform and include 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 firewall 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 (ECMP) 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.

 © Palo Alto Networks, Inc.

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. 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. 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. 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.

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 firewalls 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.

BFD-STD-MIB

Use the Bidirectional Forwarding Detection (BFD) MIB to monitor and receive failure alerts for the bidirectional path between two forwarding engines, such as interfaces, data links, or the actual engines. For example, you can check the bfdSessState object to see the state of a BFD session between forwarding engines. In the Palo Alto Networks implementation, one of the forwarding engines is a firewall interface and the other is an adjacent configured BFD peer.

RFC 7331 defines this MIB.

PAN-COMMON-MIB.my

Use the PAN-COMMON-MIB to monitor the following information for Palo Alto Networks firewalls, Panorama, and WF-500 appliances:

Object Group	Description
panSys	Contains such objects as system software/hardware versions, dynamic content versions, serial number, HA mode/state, and global counters. 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, system 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.
panLogCollector	Log Collector information such as the logging rate, log database storage duration (in days), and RAID disk usage.

 © Palo Alto Networks, Inc.

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 firewall 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).

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 firewalls, Panorama, and WF-500 appliances 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**. 2. Enter a **Name** for the profile. 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. 4. For the **Active Timeout**, specify the frequency in minutes at which the firewall exports records (default is 5). 5. Select the **PAN-OS Field Types** check box if you want the firewall to export App-ID and User-ID fields. 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). 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. 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. 2. In the **NetFlow Profile** drop-down, select the NetFlow server profile and click **OK**. 3. Click **Commit**.
Step 3	Monitor the firewall traffic in a NetFlow collector.	Refer to the documentation for your NetFlow collector. 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 IPv4 enterprise IPv4 with NAT standard 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 IPv4 enterprise IPv4 with NAT standard 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 IPv6 enterprise IPv6 with NAT standard IPv6 with NAT enterprise
28	IPV6_DST_ADDR	IPv6 destination address.	IPv6 standard IPv6 enterprise IPv6 with NAT standard IPv6 with NAT enterprise
32	ICMP_TYPE	Internet Control Message Protocol (ICMP) packet type. This is reported as: ICMP Type * 256 + ICMP code	All templates
61	DIRECTION	Flow direction: • 0 = ingress • 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

 © Palo Alto Networks, Inc.

Value	Field	Description	Templates
233	firewallEvent	Indicates a firewall event: • 0 = Ignore (invalid) • 1 = Flow created • 2 = Flow deleted • 3 = Flow denied • 4 = Flow alert • 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 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 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 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 IPv4 with NAT enterprise
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 IPv6 with NAT enterprise

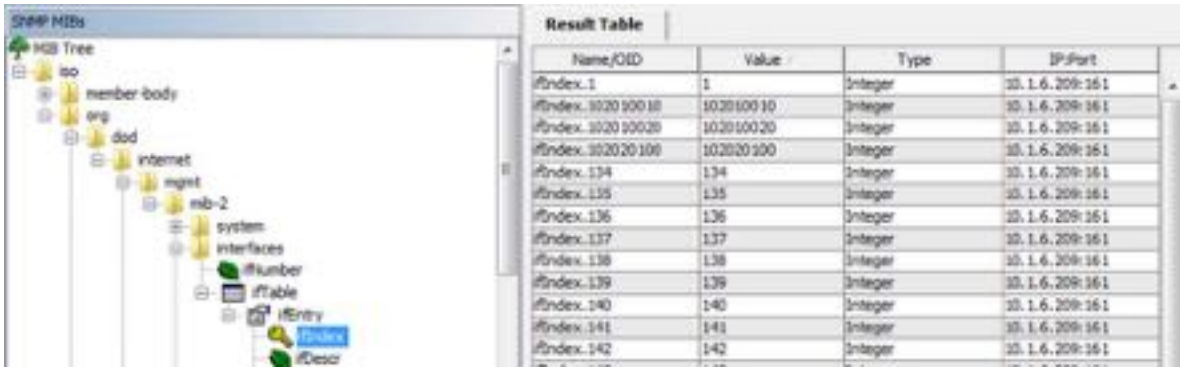
Value	Field	Description	Templates
282	postNATDestinationIPv6Address	The definition of this information element is identical to the definition of information element destinationIPv6Address, except that it reports a modified value that the firewall produced during NAT64 network address translation after 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 IPv6 with NAT enterprise
346	privateEnterpriseNumber	This is a unique private enterprise number that identifies Palo Alto Networks: 25461.	IPv4 enterprise IPv4 with NAT enterprise IPv6 enterprise 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 IPv4 with NAT enterprise IPv6 enterprise IPv6 with NAT enterprise
56702	User-ID	A username that User-ID identified. The name can be up to 64 bytes.	IPv4 enterprise IPv4 with NAT enterprise IPv6 enterprise 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:
VM-Series, PA-200, PA-500, PA-2000 Series, PA-3000 Series, PA-4000 Series, PA-5000 Series
The PA-4000 Series platform supports SNMP but not NetFlow.
	MGT port + physical port offset
	• MGT port—This is a constant that depends on the platform:
	  • 2 for hardware-based firewalls (for example, the PA-5000 Series firewall)
	  • 1 for the VM-Series firewall
	• Physical port offset—This is the physical port number.
		PA-5000 Series firewall, Eth1/4 = 2 (MGT port) + 4 (physical port) = **6**

Chassis based: PA-7000 Series firewalls This platform supports SNMP but not NetFlow.	(Max. ports * slot) + physical port offset + MGT port • Maximum ports—This is a constant of 64. • Slot—This is the chassis slot number of the network interface card. • Physical port offset—This is the physical port number. • MGT port—This is a constant of 5 for PA-7000 Series firewalls.	PA-7000 Series firewall, Eth3/9 = [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 identity, as opposed to an IP address, is an integral component of an effective security infrastructure. Knowing who is using each of the applications on your network, and who may have transmitted a threat or is transferring files, can strengthen security policies and reduce incident response times. User-ID™, a standard feature on the Palo Alto Networks firewall, enables you to leverage user information stored in a wide range of repositories. The following topics provide more details about User-ID and how to configure it:

- User-ID Overview
- User-ID Concepts
- Enable User-ID
- Map Users to Groups
- Map IP Addresses to Users
- 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

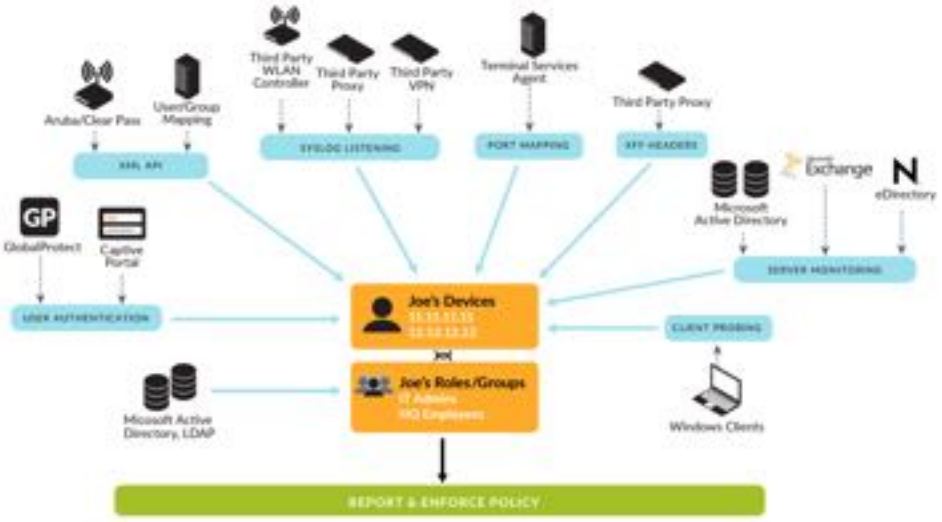
User-ID Overview

User-ID™ enables you to identify all users on your network using a variety of techniques to ensure that you can identify users in all locations using a variety of access methods and operating systems, including Microsoft Windows, Apple iOS, Mac OS, Android, and Linux®/UNIX. Knowing who your users are instead of just their IP addresses enables:

- **Visibility**—Improved visibility into application usage based on users gives you a more relevant picture of network activity. The power of User-ID becomes evident when you notice a strange or unfamiliar application on your network. Using either ACC or the log viewer, your security team can discern what the application is, who the user is, the bandwidth and session consumption, along with the source and destination of the application traffic, as well as any associated threats.
- **Policy control**—Tying user information to Security policy rules improves safe enablement of applications traversing the network and ensures that only those users who have a business need for an application have access. For example, some applications, such as SaaS applications that enable access to Human Resources services (such as Workday or Service Now) must be available to any known user on your network. However, for more sensitive applications you can reduce your attack surface by ensuring that only users who need these applications can access them. For example, while IT support personnel may legitimately need access to remote desktop applications, the majority of your users do not.
- **Logging, reporting, forensics**—If a security incident occurs, forensics analysis and reporting based on user information rather than just IP addresses provides a more complete picture of the incident. For example, you can use the pre-defined User/Group Activity to see a summary of the web activity of individual users or user groups, or the SaaS Application Usage report to see which users are transferring the most data over unsanctioned SaaS applications.

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 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 to ensure that you are safely enabling access to applications for all users, in all locations, all the time.

Figure: User-ID



 © Palo Alto Networks, Inc.

To enable user- and group-based policy enforcement, the firewall requires a list of all available users and their corresponding group memberships so that you can select groups when defining your policy rules. The firewall collects Group Mapping information by connecting directly to your LDAP directory server, or using XML API integration with your directory server.

See User-ID Concepts for information on how User-ID works and Enable User-ID for instructions on setting up User-ID.



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.

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. If you are using a directory server that is not natively supported by the firewall, you can integrate the group mapping function using the XML API. You can then create a group mapping configuration to Map Users to Groups and Enable User- and Group-Based Policy.

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. 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.

User Mapping

Knowing user and groups names 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
- Port Mapping
- Syslog
- XFF Headers
- Captive Portal
- GlobalProtect
- XML API
- Client Probing

 © Palo Alto Networks, Inc.

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. 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.

XFF Headers

User-ID can read the IPv4 or IPv6 addresses of users from the X-Forwarded-For (XFF) header in HTTP client requests when the firewall is deployed between the Internet and a proxy server that would otherwise hide the user IP addresses. User-ID matches the true user IP addresses with usernames. See Configure the firewall to obtain the user IP address from the X-Forwarded-For (XFF) header.

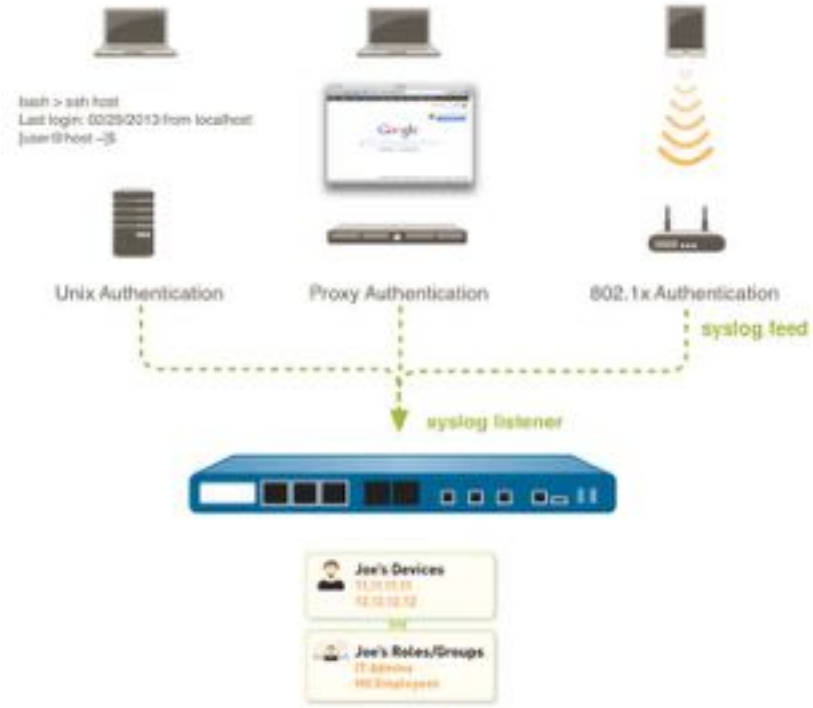
Captive Portal

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.

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



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.

 © Palo Alto Networks, Inc.

XML API

Captive Portal and the other standard user mapping methods might not work for certain types of user access. For example, the standard methods cannot add mappings of users connecting from a third-party VPN solution or users connecting to a 802.1x-enabled wireless network. For such cases, you can use the PAN-OS XML API to capture login events and send them to the PAN-OS integrated User-ID agent . See Send User Mappings to User-ID Using the XML API 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) and/or NetBIOS probing at regular intervals to verify that an existing user mapping is still valid or to obtain the username for an IP address that is not yet mapped.



NetBIOS probing is only supported on the Windows-based User-ID agent; it is not supported on the PAN-OS integrated User-ID agent.

Client probing was designed for legacy networks where most users were on Windows workstations on the internal network, but is not ideal for today's more modern networks that support a roaming and mobile user base on a variety of devices and operating systems. Additionally, client probing can generate a large amount of network traffic (based on the total number of mapped IP addresses) and can pose a security threat when misconfigured. Therefore, client probing is no longer a recommended method for user mapping. Instead collect user mapping information from more isolated and trusted sources, such as domain controllers and through integrations with Syslog or the XML API, which allow you to safely capture user mapping information from any device type or operating system. If you have sensitive applications that require you to know exactly who a user is, configure Captive Portal to ensure that you are only allowing access to authorized users.



Because WMI probing trusts data reported back from the endpoint, it is not a recommended method of obtaining User-ID information in a high-security network. If you are using the User-ID agent to parse AD security event logs, syslog messages, or the XML API to obtain User-ID mappings, Palo Alto Networks recommends disabling WMI probing.

If you do choose to use WMI probing, do not enable it on external, untrusted interfaces, as this would cause the agent to send WMI probes containing sensitive information such as the username, domain name, and password hash of the User-ID agent service account outside of your network. This information could potentially be exploited by an attacker to penetrate the network to gain further access.

If you do choose to enable probing in your trusted zones, 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.

Enable User-ID

The user identity, as opposed to an IP address, is an integral component of an effective security infrastructure. Knowing who is using each of the applications on your network, and who may have transmitted a threat or is transferring files, can strengthen your security policy and reduce incident response times. User-ID enables you to leverage user information stored in a wide range of repositories for visibility, user- and group-based policy control, and improved logging, reporting, and forensics:



On PA-5060 and PA-7000 Series firewalls that have the multiple virtual systems capability disabled, you can base policies on up to 3,200 distinct user groups. If these platforms have multiple virtual systems, the limit is 640 groups. All other firewall platforms support up to 640 groups per virtual system or per firewall (if it doesn't have multiple virtual systems).

Use the following workflow to configure User-ID.

Configure User-ID		
Step 1	Enable User-ID on the source zones that contain the users who will send requests that require user-based access controls.  Enable User-ID on trusted zones only. If you enable User-ID and client probing on an external untrusted zone (such as the internet), probes could be sent outside your protected network, resulting in an information disclosure of the User-ID agent service account name, domain name, and encrypted password hash, which could allow an attacker to gain unauthorized access to protected resources.	1. Select **Network > Zones** and click the Name of the zone. 2. Select the **Enable User Identification** check box and click **OK**.
Step 2	Create a Dedicated Service Account for the User-ID Agent.  Create a service account with the minimum set of permissions required to support the User-ID options you enable to reduce your attack surface in the event that the service account is compromised.	This is required if you plan to use the Windows-based User-ID agent or the PAN-OS integrated User-ID agent to monitor domain controllers, Exchange servers, Windows clients for user login and logout events.
Step 3	Map Users to Groups.	This enables the firewall to connect to your LDAP directory and retrieve Group Mapping information so that you will be able to select usernames and group names when creating policy.

 © Palo Alto Networks, Inc.

Configure User-ID (Continued)	
Step 4 Map IP Addresses to Users.  As a best practice, do not enable client probing as a user mapping method on high-security networks. Client probing can generate a large amount of network traffic and can pose a security threat when misconfigured.	The way you do this depends on where your users are located and what types of systems they are using, and what systems on your network are collecting login and logout events for your users. You must configure one or more User-ID agents to enable User Mapping: • Configure User Mapping Using the Windows User-ID Agent. • Configure User Mapping Using the PAN-OS Integrated User-ID Agent. • Configure User-ID to Receive User Mappings from a Syslog Sender. • Configure User Mapping for Terminal Server Users. • Send User Mappings to User-ID Using the XML API.
Step 5 Specify the networks to include and exclude from user mapping.  As a best practice, always specify which networks to include and exclude from User-ID. This allows you to ensure that only your trusted assets are probed and that unwanted user mappings are not created unexpectedly.	Configure each agent that you configured for user mapping as follows: • Specify the subnetworks the Windows User-ID agent should include in or exclude from User-ID. • Specify the subnetworks the PAN-OS integrated User-ID agent should include in or exclude from user mapping.
Step 6 Enable user- and group-based policy enforcement.  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.	After configuring User-ID, you will be able to choose a username or group name when defining the source or destination of a security rule: 1. Select **Policies > Security** and **Add** a new rule or click an existing rule name to edit. 2. Select the **User** tab and specify which users and groups to match in the rule in one of the following ways: • If you want to select specific users/groups as matching criteria, 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. • If you want to match any user who has or has not 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. 3. 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 a Basic Security Policy.

Configure User-ID (Continued)	
Step 7 Create the Security policy rules to safely enable User-ID within your trusted zones and prevent User-ID traffic from egressing your network.	Follow the Best Practice Internet Gateway Security Policy to ensure that the User-ID application (paloalto-userid-agent) is only allowed in the zones where your agents (both your Windows agents and your PAN-OS integrated agents) are monitoring services and distributing mappings to firewalls. Specifically: • Allow the paloalto-userid-agent application between the zones where your agents reside and the zones where the monitored servers reside (or even better, between the specific systems that host the agent and the monitored servers). • Allow the paloalto-userid-agent application between the agents and the firewalls that need the user mappings and between firewalls that are redistributing user mappings and the firewalls they are redistributing the information to. • Deny the paloalto-userid-agent application to any external zone, such as your internet zone.
Step 8 Configure Captive Portal. Because Captive Portal authenticates users rather than relying on user mappings, it is useful for ensuring that you know exactly who is accessing your most sensitive applications and data. You can configure Captive Portal as the fall-back to identify users who have not yet been identified using another user mapping method before allowing access. As a best practice, choose Kerberos transparent authentication over NTLM authentication when configuring Captive Portal. Kerberos is a stronger, more robust authentication method than NTLM and it does not require the firewall to have an administrative account to join the domain.	1. Select **Policies > Captive Portal**. 2. **Add** a **Name** for the rule. 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 rule. See Set Up a Basic Security Policy for details. 4. Define the **Action** to take on traffic that matches the rule: • **no-captive-portal**—Allow traffic to pass without presenting a Captive Portal page for authentication. • **web-form**—Present a Captive Portal page for the user to explicitly enter authentication credentials or use client certificate authentication. • **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. 5. Click **OK** and **Commit**.

 © Palo Alto Networks, Inc.

Configure User-ID (Continued)		
Step 9	Configure the firewall to obtain the user IP address from the X-Forwarded-For (XFF) header. When the firewall is between the Internet and a proxy server, the IP address in the packet the firewall sees contains is for the proxy server rather than the user. To enable visibility of the user IP address instead, configure the firewall to use the XFF header for user mapping. With this option enabled, the firewall matches the IP addresses with usernames referenced in policy to enable control and visibility 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. 2. Select the **X-Forwarded-For Header in User-ID** check box.  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. 3. Click **OK** to save your changes.
Step 10	Verify the User-ID Configuration.	After you configure user mapping and group mapping, verify that it is working properly and that you can safely enable and monitor user and group access to the applications, resources, and services.

Map Users to Groups

Defining policy rules based on user group membership rather than individual users simplifies administration because you don't have to update the rules whenever group membership changes. Use the following procedure to enable the firewall to connect to your LDAP directory and retrieve Group Mapping information. You can then Enable User- and Group-Based Policy.



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. 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: 1. Select **Device > Server Profiles > LDAP**, click **Add**, and enter a **Profile Name**. 2. For each LDAP server, click **Add** and enter the server **Name**, IP address (**LDAP Server**), and **Port** (default is 389). 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. 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. 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`). 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**. 2. If the firewall has more than one virtual system (vsys), select a **Location** (vsys or **Shared**) for this configuration. 3. Click **Add** and enter a unique **Name** to identify the group mapping configuration. 4. Select the LDAP **Server Profile** you just created. 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. 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**. 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**. 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. 9. Make sure the **Enabled** check box is selected.

Map Users to Groups (Continued)	
Step 3 Limit which groups will be available in policy rules. 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.  Any custom groups you create will also be available in the Allow List of authentication profiles.	1. Add existing groups from the directory service: a. Select the **Group Include List** tab. b. In the Available Groups list, select the groups you want to appear in policy rules and click the Add icon. 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: a. Select the **Custom Group** tab and click **Add**. 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). 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.  To minimize the performance impact on the LDAP directory server, use only indexed attributes in the filter. 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

User-ID provides many different methods for mapping IP addresses to usernames. Before you begin configuring user mapping, consider where your users are logging in from, what services they are accessing, and what applications and data you need to control access to. This will inform which types of agents or integrations would best allow you to identify your users. For guidance, refer to Architecting User Identification Deployments.

Once you have your plan, you can begin configuring user mapping using one or more of the following methods as needed to enable user-based access and visibility to applications and resources:

- To map users as they log in to your Exchange servers, domain controllers, eDirectory servers, or Windows clients you must configure a User-ID agent:
  - Configure User Mapping Using the PAN-OS Integrated User-ID Agent
  - Configure User Mapping Using the Windows User-ID Agent
- 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 Services 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 PAN-OS 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.

  While you can configure either the Windows agent or the PAN-OS integrated User-ID agent on the firewall to listen for authentication syslog messages from the network services, because only the PAN-OS integrated agent supports syslog listening over TLS, it is the preferred configuration.
- If you have users with client systems that aren't logged in to 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 other methods, you can Send User Mappings to User-ID Using the XML API.
- A large-scale network can have hundreds of information sources that firewalls query for user and group mapping and can have numerous firewalls that enforce policies based on the mapping information. You can simplify User-ID administration for such a network by aggregating the mapping information before the User-ID agents collect it. You can also reduce the resources that the firewalls and information sources use in the querying process by configuring some firewalls to redistribute the mapping information. For details, see Deploy User-ID in a Large-Scale Network.

Create a Dedicated Service Account for the User-ID Agent

If you plan to use either the Windows-based User-ID agent or the PAN-OS integrated User-ID agent to map users as they log in to your Exchange servers, domain controllers, eDirectory servers, or Windows clients, you must create a dedicated service account for the User-ID agent on a domain controller in each domain that the agent will monitor.

The required permissions for the service account depend on what user mapping methods and settings you plan to use. To reduce the risk associated with compromise of the User-ID service account, always configure the account with the minimum set of permissions necessary for the agent to function properly.



User-ID provides many methods for safely collecting user mapping information. Some of the legacy features, which were designed for environments that only required mapping of users on Windows desktops attached to the local network, require privileged service accounts. In the event that the privileged service account is compromised, this would open your network to attack. As a best practice, avoid using these legacy features—such as client probing, NTLM authentication, and session monitoring—that require privileges that would pose a threat if compromised. The following workflow details all privileges required and provide guidance as to which User-ID features require privileges that could pose a threat so that you can decide how to best identify users without compromising your overall security posture.

Configure an Active Directory account for the User-ID Agent	
Step 1 Create an AD account for the User-ID agent. You must create a service account in each domain the agent will monitor.	1. Log in to the domain controller. 2. Right-click the Windows icon ( ), **Search** for `Active Directory Users and Computers`, and launch the application. 3. In the navigation pane, open the domain tree, right-click **Managed Service Accounts** and select **New > User**. 4. Enter the **First Name**, **Last Name**, and **User logon name** of the user and click **Next**. 5. Enter the **Password** and **Confirm Password**, and then click **Next** and **Finish**.

 © Palo Alto Networks, Inc.

Configure an Active Directory account for the User-ID Agent (Continued)	
Step 2 Add the account to the Builtin groups that have privileges for accessing the services and hosts the User-ID agent will monitor.	1. Right-click the service account you just added and **Add to a group**. 2. **Enter the object names to select** as follows to assign the account to groups. Separate each entry with a semicolon. • **Event Log Readers** or a custom group that has privileges for reading Security log events. These privileges are required if the User-ID agent will collect mapping information by monitoring Security logs. • (PAN-OS integrated agent only) **Distributed COM Users** group, which has privileges for launching, activating, and using Distributed Component Object Model (DCOM) objects. • (Not recommended) **Server Operators** group, which has privileges for opening sessions. The agent only requires these privileges if you plan to configure it to refresh existing mapping information by monitoring user sessions.  Because this group also has privileges for shutting down and restarting servers, assign the account to it only if monitoring user sessions is very important. • (PAN-OS integrated agent only) If you plan to configure NTLM authentication for Captive Portal, the firewall where you've configured the agent will need to join the domain. To enable this, enter the name of a group that has administrative privileges to join the domain, write to the validated service principal name, and create a computer object within the *computers* organization unit (ou=computers).  The PAN-OS integrated agent requires privileged operations to join the domain, which poses a security threat if the account is compromised. Consider configuring Kerberos authentication for Captive Portal instead of NTLM. Kerberos is a stronger, more secure authentication method and it does not require the firewall to join the domain. For a firewall with multiple virtual systems, only vsys1 can join the domain because of AD restrictions on virtual systems running on the same host. 3. **Check Names** to validate your entries and click **OK** twice.

Configure an Active Directory account for the User-ID Agent (Continued)	
Step 3	If you plan to use WMI probing, enable the account to read the CIMV2 namespace on the client systems. By default, accounts in the Server Operators group have this permission. Do not enable client probing on high-security networks. Client probing can generate a large amount of network traffic and can pose a security threat when misconfigured. Instead collect user mapping information from more isolated and trusted sources, such as domain controllers and through integrations with Syslog or the XML API, which have the added benefit of allowing you to safely capture user mapping information from any device type or operating system, instead of just Windows clients.	Perform this task on each client system that the User-ID agent will probe for user mapping information: 1. Right-click the Windows icon (), **Search** for `wmimgmt.msc`, and launch the WMI Management Console. 2. In the console tree, right-click **WMI Control** and select **Properties**. 3. Select **Security**, select **Root > CIMV2**, and click **Security**. 4. **Add** the name of the service account you created, **Check Names** to verify your entry, and click **OK**. You might have to change the **Locations** or click **Advanced** to query for account names. See the dialog help for details. 5. In the Permissions for <Username> section, **Allow** the **Enable Account** and **Read Security** permissions. 6. Click **OK** twice.
Step 4	Turn off account privileges that are not necessary. By ensuring that the User-ID service account has the minimum set of account privileges, you can reduce the attack surface should the account be compromised.	To ensure that the User-ID account has the minimum privileges necessary, deny the following privileges on the account: • **Deny interactive logon for the User-ID service account**—While the User-ID service account does need permission to read and parse Active Directory security event logs, it does not require the ability to logon to servers or domain systems interactively. You can restrict this privilege using Group Policies or by using a Managed Service account (refer to Microsoft TechNet for more information). • **Deny remote access for the User-ID service account**—This prevents an attacker from using the account to access your network from the outside the network.
Step 5	Next steps...	You are now ready to: • Configure User Mapping Using the Windows User-ID Agent. • Configure User Mapping Using the PAN-OS Integrated User-ID Agent.

 © 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 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, locate your User-ID agents near the servers it will monitor (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 user 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 supported server OS versions, refer to "Operating System (OS) Compatibility User-ID Agent" in the User-ID Agent Release Notes.

Install the Windows User-ID Agent		
Step 1	Create a dedicated Active Directory service account for the User-ID agent to access the services and hosts it will monitor to collect user mappings.	Create a Dedicated Service Account for the User-ID Agent.
Step 2	Decide where to install the User-ID agent. 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.  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. • Make sure the system that will host the User-ID agent is a member of the same domain as the servers it will monitor. • 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). • 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.

Install the Windows User-ID Agent (Continued)		
Step 3	Download the User-ID agent installer.  Install the User-ID agent version that is the same as the PAN-OS version running on the firewalls. If there is not a User-ID agent version that matches the PAN-OS version, install the latest version that is closest to the PAN-OS version. For example, if you are running PAN-OS 7.1 on your firewalls, install User-ID agent version 7.0.	1. Log in to the Palo Alto Networks Customer Support web site. 2. Select **Software Updates** from the Manage Devices section. 3. Scroll to the User Identification Agent section of the screen and **Download** the version of the User-ID agent you want to install. 4. Save the `UaInstall-x.x.x-xx.msi` file on the system(s) where you plan to install the agent.

Step 4	Run the installer as an administrator.	1. Open the Windows **Start** menu, right-click the **Command Prompt** program, and select **Run as administrator**.
		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:
		`C:\Users\administrator.acme>cd Desktop`
		`C:\Users\administrator.acme\Desktop>UaInstall-6.0.0-1.msi`
		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.
		4. When the installation completes, **Close** the setup window.

Step 5	Launch the User-ID Agent application.	Open the Windows **Start** menu and select **User-ID Agent**.
Step 6	(Optional) Change the service account that the User-ID agent uses to log in.	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: 1. Select **User Identification > Setup** and click **Edit**. 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. 3. Enter the **Password** for the specified account.

 © Palo Alto Networks, Inc.

Install the Windows User-ID Agent (Continued)	
Step 7 (Optional) Assign account permissions to the installation folder. 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: a. From the Windows Explorer, navigate to `C:\Program Files\Palo Alto Networks` and right-click the folder and select **Properties**. 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. 2. Give the service account permissions to the User-ID Agent registry sub-tree: a. Run `regedit32` and navigate to the Palo Alto Networks sub-tree in one of the following locations: – **32-bit systems**—`HKEY_LOCAL_MACHINE\Software\ Palo Alto Networks` – **64-bit systems**—`HKEY_LOCAL_MACHINE\Software\ WOW6432Node\Palo Alto Networks` b. Right-click the Palo Alto Networks node and select **Permissions**. c. Assign the User-ID service account **Full Control** and then click **OK** to save the setting. 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): a. Run the MMC and Launch the Active Directory Users and Computers snap-in. 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. 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. 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.



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.

Map IP Addresses to Users Using the Windows-based User-ID Agent	
Step 1 Define the servers the User-ID agent will monitor to collect IP address to user mapping information. The User-ID agent can monitor up to 100 servers, of which up to 50 can be syslog senders.  To collect all of the required mappings, the User-ID agent 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. Open the Windows **Start** menu and select **User-ID Agent**. 2. Select **User Identification > Discovery**. 3. In the **Servers** section of the screen, click **Add**. 4. Enter a **Name** and **Server Address** for the server to be monitored. The network address can be a FQDN or an IP address. 5. 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. 6. (Optional) To enable the firewall to automatically discover domain controllers on your network using DNS lookups, click **Auto Discover**.  Auto-discovery locates domain controllers in the local domain only; you must manually add Exchange servers, eDirectory servers, and syslog senders. 7. (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. Increase the value in this field to 5 seconds in environments with older Domain Controllers or high-latency links.  Ensure that the **Enable Server Session Read** setting is not selected. This setting requires that the User-ID agent have an Active Directory account with Server Operator privileges so that it can read all user sessions. Instead, use a Syslog or XML API integration to monitor sources that capture login and logout (XML API only) events for all device types and operating systems (instead of just Windows), such as wireless controllers and Network Access Controllers (NACs). 8. Click **OK** to save the settings.

 © Palo Alto Networks, Inc.

Map IP Addresses to Users Using the Windows-based User-ID Agent (Continued)	
Step 2 Specify the subnetworks the Windows User-ID agent should include in or exclude from User-ID. By default, the User-ID maps all users accessing the servers you are monitoring.  As a best practice, always specify which networks to include and exclude from User-ID to ensure that the agent is only communicating with internal resources and to prevent unauthorized users from being mapped. You should only enable User-ID on the subnetworks where users internal to your organization are logging in.	1. Select **User Identification > Discovery**. 2. **Add** an entry to the Include/Exclude list of configured networks and enter a **Name** for the entry and enter the IP address range of the subnetwork in as the **Network Address**. 3. Select whether to include or exclude the network: • **Include specified network**—Select this option if you want to limit user mapping to users logged in to the specified subnetwork only. For example, if you include 10.0.0.0/8, the agent maps the users on that subnetwork and excludes all others. If you want the agent to map users in other subnetworks, you must repeat these steps to add additional networks to the list. • **Exclude specified network**—Select this option only if you want the agent to exclude a subset of the subnetworks you added for inclusion. For example, if you include 10.0.0.0/8 and exclude 10.2.50.0/22, the agent will map users on all the subnetworks of 10.0.0.0/8 except 10.2.50.0/22, and will exclude all subnetworks outside of 10.0.0.0/8.  If you add subnetworks for exclusion without adding any for inclusion, the agent will not perform user mapping in any subnetwork. 4. Click **OK**.
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. 2. Select the **eDirectory** tab and then complete the following fields: • **Search Base**—The starting point or root context for agent queries, for example: `dc=domain1, dc=example, dc=com`. • **Bind Distinguished Name**—The account to use to bind to the directory, for example: `cn=admin, ou=IT, dc=domain1, dc=example, dc=com`. • **Bind Password**—The bind account password. The agent saves the encrypted password in the configuration file. • **Search Filter**—The search query for user entries (default is `objectClass=Person`). • **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. • **Use SSL**—Select the check box to use SSL for eDirectory binding. • **Verify Server Certificate**—Select the check box to verify the eDirectory server certificate when using SSL.

Map IP Addresses to Users Using the Windows-based User-ID Agent (Continued)		
Step 4	(Optional, not recommended) Configure client probing.  Do not enable client probing on high-security networks. Client probing can generate a large amount of network traffic and can pose a security threat when misconfigured.	1. On the **Client Probing** tab, select the **Enable WMI Probing** check box and/or the **Enable NetBIOS Probing** check box. 2. Make sure the Windows firewall will allow client probing by adding a remote administration exception to the Windows firewall for each probed client.  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. Although client probing is not recommended, if you plan to enable it, WMI probing is 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.

Step 6	(Optional) Define the set of users for which you do not need to provide IP address-to-username mappings, such as kiosk accounts.
	  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.
		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:
		`SPAdmin`
		`SPInstall`
		`TFSReport`
		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: 1. Select **Device > User Identification > User-ID Agents** and click **Add**. 2. Enter a **Name** for the User-ID agent. 3. Enter the IP address of the Windows **Host** on which the User-ID Agent is installed. 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. 5. Make sure that the configuration is **Enabled**, then click **OK**. 6. **Commit** the changes. 7. Verify that the **Connected status** displays as connected (a green light).

 © Palo Alto Networks, Inc.

Map IP Addresses to Users Using the Windows-based User-ID Agent (Continued)		
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**. 2. Verify that the agent status shows **Agent is running**. If the Agent is not running, click **Start**. 3. To verify that the User-ID agent can connect to monitored servers, make sure the Status for each Server is **Connected**. 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**. 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 User-ID agent on the firewall for IP address-to-username 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 service account for the User-ID agent to access the services and hosts it will monitor for collecting user mapping information.	Create a Dedicated Service Account for the User-ID Agent.
Step 2	Define the servers that the firewall will monitor to collect user mapping information. Within the total maximum of 100 monitored servers per firewall, you can define no more than 50 syslog senders for any single virtual system. 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**. 2. Click **Add** in the Server Monitoring section. 3. Enter a **Name** to identify the server. 4. Select the **Type** of server. 5. Enter the **Network Address** (an FQDN or IP address) of the server. 6. Make sure the server profile is **Enabled** and click **OK**. 7. (Optional) Click **Discover** if you want the firewall to automatically discover domain controllers on your network using DNS lookups. The auto-discovery feature is for domain controllers only; you must manually add any Exchange servers or eDirectory servers you want to monitor. 8. (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. If the query load is high, the observed delay between queries might significantly exceed the specified frequency. a. Edit the Palo Alto Networks User ID Agent Setup. b. Select the **Server Monitor** tab and specify the **Server Log Monitor Frequency** in seconds (default is 2, range is 1-3600). Increase the value in this field to 5 seconds in environments with older domain controllers or high-latency links. Ensure that the **Enable Session** setting is not selected. This setting requires that the User-ID agent have an Active Directory account with Server Operator privileges so that it can read all user sessions. Instead, use a Syslog or XML API integration to monitor sources that capture login and logout events for all device types and operating systems (instead of just Windows), such as wireless controllers and NACs. c. Click **OK** to save the changes.

 © Palo Alto Networks, Inc.

Map IP Addresses to Users Using the Integrated User-ID Agent (Continued)		
Step 3	Specify the subnetworks the PAN-OS integrated User-ID agent should include in or exclude from user mapping. By default, the User-ID maps all users accessing the servers you are monitoring. As a best practice, always specify which networks to include and, optionally, to exclude from User-ID to ensure that the agent is only communicating with internal resources and to prevent unauthorized users from being mapped. You should only enable user mapping on the subnetworks where users internal to your organization are logging in.	1. Select **Device > User Identification > User Mapping**. 2. **Add** an entry to the Include/Exclude Networks and enter a **Name** for the entry and make sure to keep the **Enabled** check box selected. 3. Enter the **Network Address** and then select whether to include or exclude it: • **Include**—Select this option if you want to limit user mapping to users logged in to the specified subnetwork only. For example, if you include 10.0.0.0/8, the agent maps the users on that subnetwork and excludes all others. If you want the agent to map users in other subnetworks, you must repeat these steps to add additional networks to the list. • **Exclude**—Select this option only if you want the agent to exclude a subset of the subnetworks you added for inclusion. For example, if you include 10.0.0.0/8 and exclude 10.2.50.0/22, the agent will map users on all the subnetworks of 10.0.0.0/8 except 10.2.50.0/22, and will exclude all subnetworks outside of 10.0.0.0/8. If you add subnetworks for exclusion without adding any for inclusion, the agent will not perform user mapping in any subnetwork. 4. Click **OK**.
Step 4	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. 2. Select the **WMI Authentication** tab and enter the **User Name** and **Password** for the account that the User-ID agent will use to probe the clients and monitor servers. Enter the username using the domain\username syntax.
Step 5	(Optional, not recommended) Configure WMI probing (the PAN-OS integrated User-ID agent does not support NetBIOS probing). Do not enable WMI probing on high-security networks. Client probing can generate a large amount of network traffic and can pose a security threat when misconfigured.	1. Select the **Client Probing** tab and select the **Enable Probing** check box. 2. (Optional) Modify the **Probe Interval** (in minutes) if necessary 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. If the request load is high, the observed delay between requests might significantly exceed the specified interval. 3. Click **OK**. 4. Make sure the Windows firewall will allow client probing by adding a remote administration exception to the Windows firewall for each probed client.

Map IP Addresses to Users Using the Integrated User-ID Agent (Continued)		
Step 6	(Optional) Define the set of users for which you don't require IP address-to-username mappings, such as kiosk accounts. You can also use the ignore user list to identify users whom you want to force to authenticate using Captive Portal.	Select the **Ignore User List** tab and **Add** each username to exclude from user mapping. 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`. You can add up to 5,000 entries to exclude from user mapping.
Step 7	Activate your configuration changes.	Click **OK** and **Commit**.
Step 8	Verify the configuration.	1. Access the firewall CLI. 2. Enter the following operational command: `> show user server-monitor state all` 3. 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.

Configure User-ID to Receive User Mappings from a Syslog Sender

The following topics describe how to configure the PAN-OS integrated User-ID agent or Windows-based User-ID agent 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.

 © Palo Alto Networks, Inc.

Collect User Mappings from Syslog Senders	
Step 1 Determine whether there is a predefined syslog filter for your particular syslog senders. Palo Alto Networks provides several predefined syslog filters, which are delivered as Application content updates and are therefore updated dynamically as new filters are developed. The predefined filters are global to the firewall, whereas custom filters apply to a single virtual system only.  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 Applications or Applications and Threats database is up to date: a. Select **Device > Dynamic Updates** and click **Check Now** to check for the latest updates. b. **Download** and **Install** any new update. 2. Determine which predefined filters are available: a. Select **Device > User Identification > User Mapping** and **Add** an entry to the Server Monitoring section. b. Select **Syslog Sender** as the server **Type**. c. 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 to define the servers. Otherwise, continue to Step 2.
Step 2 Define custom syslog filters to extract IP address-to-username mapping information from syslog messages. For the User-ID agent to parse syslog messages, they must meet the following criteria: • Each message must be a single-line text string. A new line (`\n`) or a carriage return plus a new line (`\r\n`) are the delimiters for line breaks. • The maximum size for individual messages is 2,048 bytes. • Messages sent over UDP must be contained in a single packet; messages sent over SSL can span multiple packets. • A single packet might contain multiple messages.	1. Review the syslog messages that the authenticating service generates to identify the login event syntax. This enables you to define the matching patterns by which the firewall identifies and extracts authentication events from the messages.  While reviewing syslog messages, also determine whether log entries include the domain name. If they don't, 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. 3. Select the **Syslog Filters** tab and **Add** a Syslog Parse profile. 4. Enter a name to identify the **Syslog Parse Profile**. 5. Specify the **Type** of parsing to extract user mapping information: • **Regex Identifier**—Step 3 describes how to specify regular expressions that describe search patterns to identify and extract user mapping information from syslog messages. • **Field Identifier**—Step 4 describes how to specify a string to match the authentication event, and prefix and delimiter strings to identify user mapping information in syslog messages.

Collect User Mappings from Syslog Senders (Continued)	
Step 3 If you selected **Regex Identifier** as the parsing **Type**, define regex matching patterns to identify authentication events and extract user mapping information. This example shows how to configure a Syslog Parse profile 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 message contains a standalone space or tab as a delimiter, use `\s` for a space and `\t` for a tab.	1. In the **Event Regex** field, enter the regex to match successful authentication events in syslog messages. For the sample 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. In the **Username Regex** field, enter the regex to identify the start of the username in authentication success messages. In the sample message, the regex `User:([a-zA-Z0-9\\\._]+)` matches the string `User:johndoe1` and extracts `johndoe1` as the username.  If the syslog messages do not contain domain information and you require domain names in your user mappings, enter the **Default Domain Name** when defining the monitored server entry in Step 5. 3. In the **Address Regex** field, enter the regex to identify the IP address portion of authentication success messages. In the sample message, the regular expression `Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})` matches the IPv4 address `Source:192.168.3.212`. 4. Click **OK** to save the profile.
Step 4 If you selected **Field Identifier** as the parsing **Type**, define string matching patterns to identify authentication events and extract user mapping information. This example shows how to configure a Syslog Parse profile that matches syslog messages with the following format: `[Tue Jul 5 13:15:04 2005 CDT] Administrator authentication success User:johndoe1 Source:192.168.3.212`	1. In the **Event String** field, enter a matching string to identify successful authentication events in syslog messages. In the sample message, the string `authentication success` identifies authentication events. 2. In the **Username Prefix** field, enter a matching string to identify the start of the username field in syslog messages. The field does not support regex expressions such as /s (for a space) or /t (for a tab). In the sample message, `User:` identifies the start of the username field. 3. Enter the **Username Delimiter** that indicates the end of the username field in syslog messages. Use `\s` to indicate a standalone space (as in the sample message) and `\t` to indicate a tab. 4. In the **Address Prefix** field, enter a matching string to identify the start of the IP address field in syslog messages. The field does not support regex expressions such as /s (for a space) or /t (for a tab). In the sample message, `Source:` identifies the start of the address field. 5. Enter the **Address Delimiter** that indicates the end of the IP address field in syslog messages. For example, enter `\n` to indicate the delimiter is a line break. 6. Click **OK** to save the profile.

 © Palo Alto Networks, Inc.

Collect User Mappings from Syslog Senders (Continued)	
Step 5 Define the servers that will send syslog messages to the firewall for user mapping purposes. Within the total maximum of 100 monitored servers per firewall, you can define no more than 50 syslog senders for any single virtual system. The firewall will discard any syslog messages received from servers that are not on this list.	1. Select **Device > User Identification > User Mapping** and **Add** an entry to the Server Monitoring section. 2. Enter a **Name** to identify the server. 3. Make sure the server profile is **Enabled** (default). 4. Select **Syslog Sender** as the server **Type**. 5. Enter the **Network Address** of the syslog server (IP address or FQDN). 6. Select the Syslog Parse profile you configured as a **Filter**. 7. Select **UDP** or **SSL** (default) as the **Connection Type**.  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. A syslog server using SSL to connect will show a Status of Connected only when there is an active SSL connection. Syslog servers using UDP will not show a Status value. 8. (Optional) If the syslogs that the authenticating firewall sends do not include domain information in the login event logs, enter the **Default Domain Name** to append to the user mappings. 9. Click **OK** to save the settings.

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.	1. Select **Network > Network Profiles > Interface Mgmt** and edit an existing Interface Management profile or **Add** a new profile. 2. Select **User-ID Syslog Listener-SSL** and/or **User-ID Syslog Listener-UDP**, depending on the protocols you defined for the syslog senders in the Server Monitoring list.  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. 3. Click **OK** to save the interface management profile.  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. The firewall discards connections or messages from servers that are not on the list. 4. If you created a new Interface Management profile, assign it to the interface used for user mapping: a. Select **Network > Interfaces** and edit the interface. b. Select **Advanced > Other info**, select the Interface **Management Profile** you just added, and click **OK**.
Step 7 Save the configuration.	Click **Commit** to save the configuration.

 © Palo Alto Networks, Inc.

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 syslog messages 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 Define custom syslog filters to extract IP address-to-username mapping information from syslog messages. For the User-ID agent to parse syslog messages, they must meet the following criteria: • Each message must be a single-line text string. A new line (`\n`) or a carriage return plus a new line (`\r\n`) are the delimiters for line breaks. • The maximum size for individual messages is 2,048 bytes. • Messages sent over UDP must be contained in a single packet; messages sent over SSL can span multiple packets. • A single packet might contain multiple messages.	1. Open the Windows **Start** menu and select **User-ID Agent**. 2. Review the syslog messages that the authenticating service generates to identify the login event syntax. This enables you to define the matching patterns by which the firewall identifies and extracts authentication events from the messages. While reviewing syslog messages, also determine whether log entries include the domain name. If they don't, consider defining a default domain name when adding the syslog sender to the monitored servers list . 3. Select **User Identification > Setup** and **Edit** the Setup section. 4. Select **Syslog** and **Add** a Syslog Parse profile. 5. Enter a **Profile Name** and **Description**. 6. Specify the **Type** of parsing to extract user mapping information: • **Regex**—Uses regular expressions that describe search patterns to identify and extract user mapping information from syslog messages. • **Field**—Uses a string to match the authentication event, and prefix and delimiter strings to identify user mapping information in syslog messages.
Step 2 If you selected **Regex** as the parsing **Type**, define regex matching patterns to identify authentication events and extract user mapping information. This example shows how to configure a Syslog Parse profile 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 message contains a standalone space or tab as a delimiter, use `\s` for a space and `\t` for a tab.	1. In the **Event Regex** field, enter the regex to match successful authentication events in syslog messages. For the sample 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. In the **Username Regex** field, enter the regex to identify the start of the username in authentication success messages. In the sample message, the regex `User:([a-zA-Z0-9\\\._]+)` matches the string `User:johndoe1` and extracts `johndoe1` as the username. If the syslog messages do not contain domain information and you require domain names in your user mappings, enter the **Default Domain Name** when defining the monitored server entry. 3. In the **Address Regex** field, enter the regex to identify the IP address portion of authentication success messages. In the sample message, the regular expression `Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})` matches the IPv4 address `Source:192.168.3.212`. 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 3	If you selected **Field** as the parsing **Type**, define string matching patterns to identify authentication events and extract user mapping information. This example shows how to configure a Syslog Parse profile that matches syslog messages with the following format: `[Tue Jul 5 13:15:04 2005 CDT] Administrator authentication success User:johndoe1 Source:192.168.3.212`	1. In the **Event String** field, enter a matching string to identify successful authentication events in syslog messages. In the sample message, the string `authentication success` identifies authentication events. 2. In the **Username Prefix** field, enter a matching string to identify the start of the username field in syslog messages. The field does not support regex expressions such as /s (for a space) or /t (for a tab). In the sample message, `User:` identifies the start of the username field. 3. Enter the **Username Delimiter** that indicates the end of the username field in syslog messages. Use `\s` to indicate a standalone space (as in the sample message) and `\t` to indicate a tab. 4. In the **Address Prefix** field, enter a matching string to identify the start of the IP address field in syslog messages. The field does not support regex expressions such as /s (for a space) or /t (for a tab). In the sample message, `Source:` identifies the start of the address field. 5. Enter the **Address Delimiter** that indicates the end of the IP address field in syslog messages. For example, enter `\n` to indicate the delimiter is a line break. 6. Click **OK** to save the profile.
Step 4	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 5	Define the servers that will send syslog messages to the User-ID agent. Within the total maximum of 100 servers of all types that the User-ID agent can monitor, up to 50 can be 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. **Add** an entry to the Servers section. 3. Enter a server **Name** and **Server Address**. 4. Select **Syslog Sender** as the **Server Type**. 5. Select a **Filter** you defined in Step 1. 6. (Optional) If the syslog messages that the authenticating firewall 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 6	Save the configuration.	Click **Commit** to save the configuration.

Configure the Windows User-ID Agent to Collect User Mappings from Syslog Senders (Continued)

Step 7 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
```

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

 © Palo Alto Networks, Inc.

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.  As a best practice, choose Kerberos transparent authentication over NTLM authentication. Kerberos is a stronger, more robust authentication method than NTLM and it does not require the firewall to have an administrative account to join the domain. 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. If you use the Windows-based User-ID agent, NTLM responses go directly to the domain controller where you installed the agent. 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. 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.
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 (or an authentication sequence). 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. If you use Kerberos SSO or NTLM authentication, you must use Redirect mode because the browser will provide credentials only to trusted sites.

 © Palo Alto Networks, Inc.

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.

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. 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**. 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 communicate with directory servers. You must enable **Response Pages** and **User-ID** in the Interface Management profile. 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: • **Kerberos**, **RADIUS**, or **TACACS+**—Configure a service route for one of these services only if you will use it for external authentication. • **UID Agent**—Configure this service only if you will enable NT LAN Manager (NTLM) authentication or if you will Enable User- and Group-Based Policy. 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. 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: `admin@PA-200> ping host dc1.acme.com`
Step 3	Create a Kerberos keytab for the redirect host. 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). 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. 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: 1. Select **Device > Certificate Management > Certificates > Device Certificates**. 2. Create a Self-Signed Root CA Certificate or import a CA certificate (see Import a Certificate and Private Key). 3. Generate a Certificate to use for Captive Portal. Be sure to configure the following fields: • **Common Name**—Enter the DNS name of the intranet host for the Layer 3 interface. • **Signed By**—Select the CA certificate you just created or imported. • 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. 4. Configure an SSL/TLS Service Profile. Assign the Captive Portal certificate you just created to the profile. 5. Configure clients to trust the certificate: a. Export the CA certificate you created or imported. 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).
Step 5 Configure an authentication server profile. Required for external authentication. If you enable Kerberos SSO or NTLM authentication, the firewall uses the external service only if those methods fail. As a best practice, choose Kerberos transparent authentication over NTLM authentication. Kerberos is a stronger, more robust authentication method than NTLM and it does not require the firewall to have an administrative account to join the domain.	• Configure a RADIUS Server Profile. • Configure a TACACS+ Server Profile • Configure an LDAP Server Profile • Configure a Kerberos Server Profile 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. 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.

 © Palo Alto Networks, Inc.

Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued)	
Step 6 Add an authentication profile 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. 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: 1. If the authentication **Type** is an external service (**RADIUS**, **TACACS+**, **LDAP**, or **Kerberos**), select the authentication **Server Profile** you created. 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. 3. Select **Advanced** and **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 completing the Allow List, click **OK**. 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, and assign the sequence to the Captive Portal configuration.
Step 7 (Optional) Configure Client Certificate Authentication. 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. 2. Export the CA certificate in PEM format to a system that the firewall can access. 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**. 4. Configure a Certificate Profile. • In the **Username Field** drop-down, select the certificate field that contains the user identity information. • In the **CA Certificates** list, click **Add** and select the CA certificate you just imported.
Step 8 (Optional) Enable NT LAN Manager (NTLM) authentication. As a best practice, choose Kerberos transparent authentication over NTLM authentication. Kerberos is a stronger, more robust authentication method than NTLM and it does not require the firewall to have an administrative account to join the domain. If you do configure NTLM, the PAN-OS integrated User-ID agent must be able to successfully resolve the DNS name of your domain controller to join the domain.	1. If you haven't already done so, Create a Dedicated Service Account for the User-ID Agent. 2. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User ID Agent Setup section. 3. On the **NTLM** tab, select the **Enable NTLM authentication processing** check box. 4. Enter the **NTLM Domain** against which the User-ID agent on the firewall will check NTLM credentials. 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. 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. 6. Click **OK**.

Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued)	
Step 9 Configure the Captive Portal settings.	1. Select **Device > User Identification > Captive Portal Settings** and edit the settings. 2. Make sure the **Enable Captive Portal** check box is selected. 3. Select the **SSL/TLS Service Profile** you created for redirect requests over TLS. 4. Select the **Mode** (in this example, **Redirect**). 5. (Redirect mode only) Specify the **Redirect Host** name that resolves to the IP address of the Layer 3 interface for redirected requests. 6. Select the authentication method to use if NTLM fails (or if you don't use NTLM): • To use Kerberos SSO, an external server, or the local database, select the **Authentication Profile** or authentication sequence you created. • To use client certificate authentication, select the **Certificate Profile** you created. 7. Click **OK** and **Commit** to save the Captive Portal configuration.

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 PAN-OS 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 Services Agent for User Mapping
- Retrieve User Mappings from a Terminal Server Using the PAN-OS XML API

Configure the Palo Alto Networks Terminal Services Agent for User Mapping

Use the following procedure to install and configure the TS agent on the terminal server. To map all your users, you must install the TS agent on all terminal servers that your users log in to.



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.

 © Palo Alto Networks, Inc.

Configure the Palo Alto Networks Terminal Services Agent for User Mapping	
Step 1 Download the TS agent installer.	1. Log in to the Palo Alto Networks Customer Support web site. 2. Select **Software Updates** from the Manage Devices section. 3. Scroll to the **Terminal Services Agent** section and **Download** the version of the agent you want to install. 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 systems where you plan to install the agent.

Step 2 Run the installer as an administrator.	1. Open the Windows **Start** menu, right-click the **Command Prompt** program, and select **Run as administrator**.
	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:
	`C:\Users\administrator.acme>cd Desktop`
	`C:\Users\administrator.acme\Desktop>TaInstall-6.0.0-1.msi`
	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.
	4. When the installation completes, **Close** the setup window.
	  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.

Configure the Palo Alto Networks Terminal Services Agent for User Mapping (Continued)		
Step 3	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. Open the Windows **Start** menu and select **Terminal Server Agent** to launch the Terminal Services agent application. 2. Select **Configure** in the side menu. 3. Enter 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**. 4. (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`. 5. 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). 6. Specify the **Port Allocation Maximum Size Per User**, which is the maximum number of ports the Terminal Services agent can allocate to an individual user. 7. 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 4	Configure the firewalls to connect to the Terminal Services agent.	Complete the following steps on each firewall you want to connect to the Terminal Services agent to receive user mappings: 1. Select **Device > User Identification > Terminal Server Agents** and click **Add**. 2. Enter a **Name** for the Terminal Services agent. 3. Enter the IP address of the Windows **Host** on which the Terminal Services 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 Services 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 Services 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).

 © Palo Alto Networks, Inc.

Configure the Palo Alto Networks Terminal Services Agent for User Mapping (Continued)		
Step 5	Verify that the Terminal Services agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent.	1. Open the Windows **Start** menu and select **Terminal Server Agent**. 2. Verify that the firewalls can connect by making sure the **Connection Status** of each firewall in the Connection List is **Connected**. 3. Verify that the Terminal Services agent is successfully mapping port ranges to usernames by selecting **Monitor** in the side menu and making sure that the mapping table is populated.
Step 6	(Windows 2012 R2 servers only) Disable Enhanced Protected Mode in Microsoft Internet Explorer for each user who uses that browser. This task is not necessary for other browsers such as Google Chrome or Mozilla Firefox. To disable Enhanced Protected Mode for all users, use Local Security Policy.	Perform these steps on the Windows Server: 1. Start Internet Explorer. 2. Select **Internet options > Advanced** and scroll down to the Security section. 3. Clear **Enable Enhanced Protected Mode**. 4. Click **OK**. In Internet Explorer, Palo Alto Networks recommends that you do not disable Protected Mode, which differs from Enhanced Protected Mode.

Retrieve User Mappings from a Terminal Server Using the PAN-OS XML API

The PAN-OS XML API 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 PAN-OS 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 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 PAN-OS XML API to send user mappings from a non-Windows terminal server to the firewall.

Use the PAN-OS 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>`

 © Palo Alto Networks, Inc.

Use the PAN-OS 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 Services 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: 1,000
		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 PAN-OS 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 1,000 port blocks.

Use the PAN-OS 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 PAN-OS 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"
```

 © Palo Alto Networks, Inc.

Use the PAN-OS 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:
		To verify if the terminal server is connecting to the firewall over XML:
		`admin@PA-5050> show user xml-api multiusersystem`
		`Host Vsys Users Blocks`
		`----------`
		`10.5.204.43 vsys1 5 2`
		To verify that the firewall is receiving mappings from a terminal server over XML:
		`admin@PA-5050> show user ip-port-user-mapping all`
		`Global max host index 1, host hash count 1`
		`XML API Multi-user System 10.5.204.43`
		`Vsys 1, Flag 3`
		`Port range: 20000 - 39999`
		`Port size: start 200; max 2000`
		`Block count 100, port count 20000`
		`20000-20199: acme\administrator`
		`Total host: 1`

Send User Mappings to User-ID Using the XML API

User-ID provides many out-of-the box methods for obtaining user mapping information. However, you might have applications or devices that capture user information but cannot natively integrate with User-ID. For example, you might have a custom, internally developed application or a device that no standard user mapping method supports. In such cases, you can use the PAN-OS XML API to create custom scripts that send the information to the PAN-OS integrated User-ID agent or directly to the firewall. The PAN-OS XML API 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 POST and GET requests.

To enable an external system to send user mapping information to the PAN-OS integrated User-ID agent, create scripts that extract user login and logout events and use the events as input to the PAN-OS XML API request. Then define the mechanisms for submitting the XML API requests to the firewall (using cURL, for example) and use the API key of the firewall for secure communication. For more details, refer to the PAN-OS XML API Usage Guide.

Enable User- and Group-Based Policy

After you Enable User-ID, you will be able to configure Security Policy that applies to specific users and groups. User-based policy controls can also include application information (including which category and subcategory it belongs in, its underlying technology, or what the application characteristics are). You can define policy rules to safely enable applications based on users or groups of users, in either outbound or inbound directions.

Examples of user-based policies include:

- Enable only the IT department to use tools such as SSH, telnet, and FTP on standard ports.
- Allow the Help Desk Services group to use Slack.
- Allow all users to read Facebook, but block the use of Facebook apps, and restrict posting to employees in marketing.

 © Palo Alto Networks, Inc.

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. 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). 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: 1. Select **Device > User Identification > Group Mapping Settings** and **Add** a group mapping configuration with a unique **Name**. 2. Select an LDAP **Server Profile** and ensure the **Enabled** check box is enabled. 3. Select the **Custom Group** tab and **Add** a custom group with network_services as a **Name**. 4. Specify an **LDAP Filter** that matches an LDAP attribute of corp_user and click **OK**. 5. Click **OK** and **Commit**. 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.

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: 1. Configure a security rule that allows the corp_employees group to access email. 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. 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. 1. Create an `ignore_user_list.txt` file. 2. Open the file and add admin_user. If you later add more usernames, each must be on a separate line. 3. Save the file to the User-ID agent folder on the domain server where the agent is installed. If you use the PAN-OS integrated User-ID agent, see Configure User Mapping Using the PAN-OS Integrated User-ID Agent for instructions on how to configure the ignore list.
Step 4	Configure endpoint authentication for the restricted services. 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). 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. 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. The firewall is now ready to enforce rules for a user with multiple usernames.

 © Palo Alto Networks, Inc.

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 operational command: `> show user group-mapping statistics`

Step 2 Verify that user mapping is working.	If you are using the PAN-OS integrated User-ID agent, you can verify this from the CLI using the following command:
	`> show user ip-user-mapping-mp all`
	`IP Vsys From User Timeout (sec)`
	`192.168.201.1 vsys1 UIA acme\george 210`
	`192.168.201.11 vsys1 UIA acme\duane 210`
	`192.168.201.50 vsys1 UIA acme\betsy 210`
	`192.168.201.10 vsys1 UIA acme\administrator 210`
	`192.168.201.100 vsys1 AD acme\administrator 748`
	`Total: 5 users`
	`*: 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.
	• You can also use the `test security-policy-match` operational 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:
	`> test security-policy-match application worldofwarcraft source-user acme\duane source any destination any destination-port any protocol 6`
	`"deny worldofwarcraft" {`
	`from corporate;`
	`source any;`
	`source-region any;`
	`to internet;`
	`destination any;`
	`destination-region any;`
	`user acme\duane;`
	`category any;`
	`application/service worldofwarcraft;`
	`action deny;`
	`terminal no;`
	`}`

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. The Captive Portal web form should display and prompt you for login credentials.
		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` operationalcommand 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 the log files display usernames.	Select a logs page (for example, **Monitor > Logs > Traffic**) and verify that the Source User column displays usernames.
Step 6	Verify that reports display usernames.	1. Select **Monitor > Reports**. 2. Select a report type that includes usernames. For example, the Denied Applications report, Source User column, should display a list of the users who attempted to access the applications.

 © Palo Alto Networks, Inc.

Deploy User-ID in a Large-Scale Network

A large-scale network can have hundreds of information sources that firewalls query to map IP addresses to usernames and to map usernames to user groups. You can simplify User-ID administration for such a network by aggregating the user mapping and group mapping information before the User-ID agents collect it, thereby reducing the number of required agents.

A large-scale network can also have numerous firewalls that use the mapping information to enforce policies. You can reduce the resources that the firewalls and information sources use in the querying process by configuring some firewalls to acquire mapping information through redistribution instead of direct querying. Redistribution also enables the firewalls to enforce user-based policies when users rely on local sources for authentication (for example, regional directory services) but need access to remote resources (for example, global data center applications).

- Deploy User-ID for Numerous Mapping Information Sources
- Configure Firewalls to Redistribute User Mapping Information

Deploy User-ID for Numerous Mapping Information Sources

You can use Windows Log Forwarding and Global Catalog servers to simplify user mapping and group mapping in a large-scale network of Microsoft Active Directory (AD) domain controllers or Exchange servers. These methods simplify User-ID administration by aggregating the mapping information before the User-ID agents collect it, thereby reducing the number of required agents.

- Windows Log Forwarding and Global Catalog Servers
- Plan a Large-Scale User-ID Deployment
- Configure Windows Log Forwarding
- Configure User-ID for Numerous Mapping Information Sources

Windows Log Forwarding and Global Catalog Servers

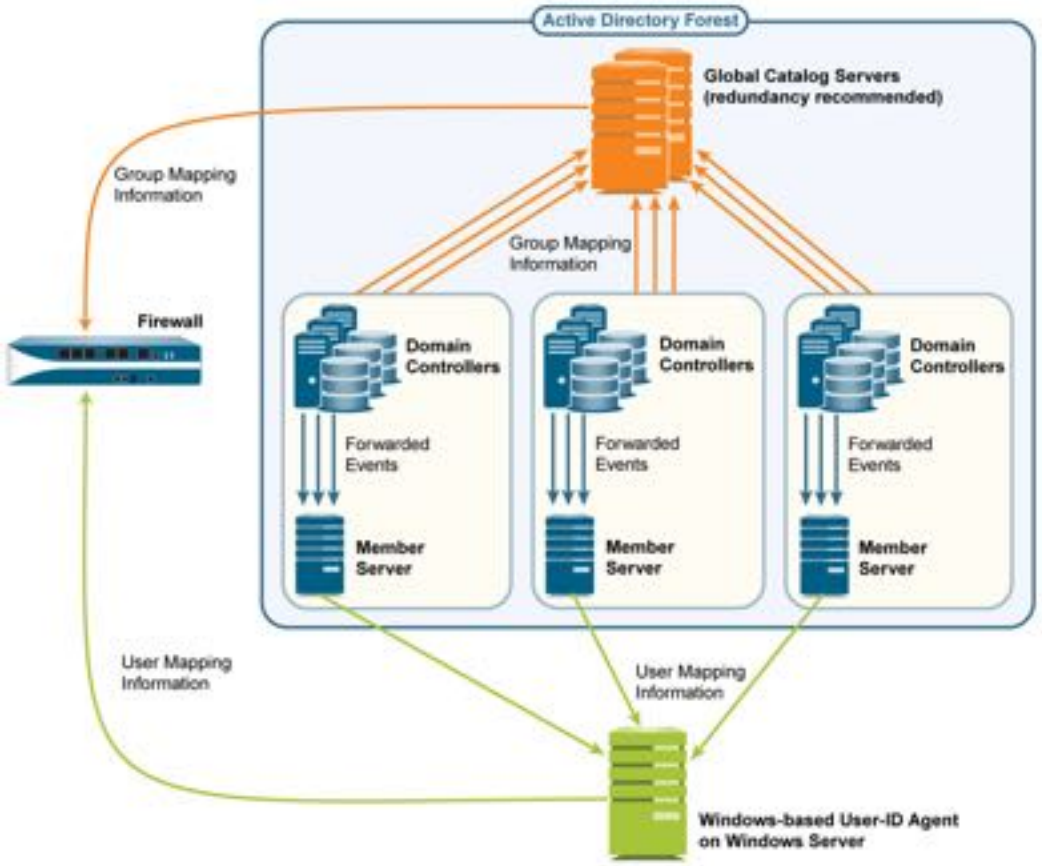
Because each User-ID agent can monitor up to 100 servers, the firewall needs multiple User-ID agents to monitor a network with hundreds of AD domain controllers or Exchange servers. Creating and managing numerous User-ID agents involves considerable administrative overhead, especially in expanding networks 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 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 a Large-Scale User-ID Deployment to determine if this deployment suits your network.



Plan a Large-Scale User-ID Deployment

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—Must support the processing load associated with forwarding the events.
  - Member Servers—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: • Windows Server 2003—The event IDs for the required events are 672 (Authentication Ticket Granted), 673 (Service Ticket Granted), and 674 (Ticket Granted Renewed). • 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).  You must forward events to the security logs location on the member servers, not to the default forwarded logs location.  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 Numerous Mapping Information Sources

Configure User-ID for Numerous Mapping Information Sources		
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.
Step 3	Configure the User-ID agent to collect user mapping information from the member servers.	1. Start the Windows-based User-ID agent. 2. Select **User Identification > Discovery** and perform the following steps for each member server that will receive events from domain controllers: a. In the Servers section, click **Add** and enter a **Name** to identify the member server. b. In the **Server Address** field, enter the FQDN or IP address of the member server. c. For the **Server Type**, select **Microsoft Active Directory**. d. Click **OK** to save the server entry. 3. Configure the remaining User-ID agent settings: see Configure the User-ID Agent for User Mapping.

Configure User-ID for Numerous Mapping Information Sources (Continued)	
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. To improve availability, use at least two Global Catalog servers for redundancy. 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. 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. 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`). 4. For the **Type**, select **active-directory**. 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. User-ID uses this information to map DNS domain names to NetBIOS domain names. This mapping ensures consistent domain/username references in policy rules. 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: • **LDAP Server**—Enter the IP address of the domain controller that contains the domain mapping information. • **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. • **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`).
Step 6 Create a group mapping configuration for each LDAP server profile you created.	1. Select **Device > User Identification > Group Mapping Settings**. 2. Click **Add** and enter a **Name** to identify the group mapping configuration. 3. Select the LDAP **Server Profile** and ensure the **Enabled** check box is selected. 4. Configure the remaining fields as necessary: see Map Users to Groups. 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. 5. Click **OK** and **Commit**.

 © Palo Alto Networks, Inc.

Configure Firewalls to Redistribute User Mapping Information

Every firewall that enforces user-based policy requires user mapping information. However, a large-scale network where numerous firewalls directly query the mapping information sources requires both the firewalls and sources to use considerable resources. To improve resource efficiency, you can configure some firewalls to acquire mapping information through redistribution instead of direct querying. Redistribution also enables the firewalls to enforce user-based policies when users rely on local sources for authentication (for example, regional directory services) but need access to remote resources (for example, global data center applications).

▲ Firewall Deployment for User-ID Redistribution

▲ Configure User-ID Redistribution

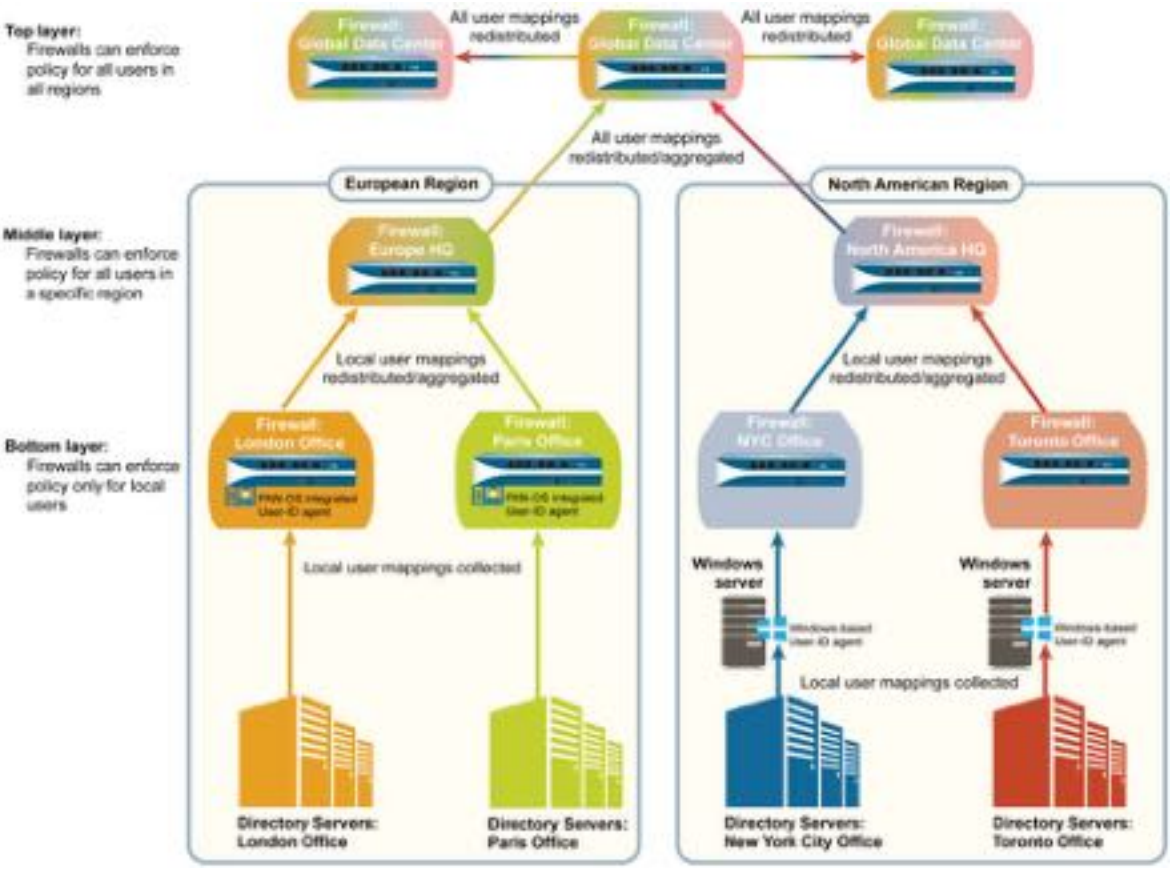
Firewall Deployment for User-ID Redistribution

You can organize the redistribution sequence in layers, where each layer has one or more firewalls. In the bottom layer, PAN-OS integrated User-ID agents running on firewalls and Windows-based User-ID agents running on Windows servers perform the IP address-to-username mapping. Each higher layer has firewalls that receive the mapping information from up to 100 User-ID agents in the layer beneath it. The top-layer firewalls aggregate the mapping information from all layers. This deployment provides the option to configure global policies for all users (in top-layer firewalls) and region- or function-specific policies for a subset of users in the corresponding domains (in lower-layer firewalls).

Figure: User-ID-Redistribution shows a deployment with three layers of firewalls that redistribute mapping information from local information sources (directory servers, in this example) to regional offices and then to a global data center. The data center firewall that aggregates all the mapping information shares it with other data center firewalls so that they can all enforce global policy. Only the bottom layer firewalls use PAN-OS integrated User-ID agents and Windows-based User-ID agents to query the directory servers.

The information sources from which User-ID agents collect mapping information do not count towards the maximum of ten *hops* in the sequence. However, Windows-based User-ID agents that forward mapping information to firewalls do count. Therefore, in this example, redistribution from the European region to all the data center firewalls requires only three hops, while redistribution from the North American region requires four hops. Also in this example, the top layer has two hops: the first to aggregate mapping information in one data center firewall and the second to share the information with other data center firewalls.

Figure: User-ID-Redistribution



Configure User-ID Redistribution

Configure User-ID Redistribution	
Step 1 Plan the redistribution architecture.	❑ Decide which User-ID agents and methods to use for mapping IP addresses to usernames. You can redistribute user mapping information collected through any method except Terminal Services (TS) agents. You cannot redistribute Group Mapping or HIP match information. ❑ Determine the most efficient Firewall Deployment for User-ID Redistribution. Some factors to consider are: • Which firewalls will enforce global policies for all users and which firewalls will enforce region- or function-specific policies for a subset of users? • How many hops does the redistribution sequence require to aggregate mapping information for firewalls in different functional or regional layers to enforce policy? • How can you minimize the number of firewalls that query the information sources? The fewer the number of querying firewalls, the lower the processing load is on both the firewalls and sources.

 © Palo Alto Networks, Inc.

Configure User-ID Redistribution (Continued)		
Step 2	Configure the User-ID agents to perform the user mapping.	• Configure User Mapping Using the PAN-OS Integrated User-ID Agent. • Configure User Mapping Using the Windows User-ID Agent.
Step 3	Enable each bottom-layer firewall to forward mapping information to firewalls in the layer above.	1. Configure the firewall to function as a User-ID agent. a. Select **Device > User Identification > User Mapping**. b. (Firewalls with multiple virtual systems only) Select the **Location**. You must configure the User-ID settings for each virtual system.  You can redistribute mapping information among virtual systems on different firewalls or on the same firewall. In both cases, each virtual system counts as one hop in the redistribution sequence. c. Edit the Palo Alto Networks User-ID Agent Setup and select **Redistribution**. d. Enter a **Collector Name** to identify this firewall as a User-ID agent. e. Enter and confirm a **Pre-Shared Key** to secure communication between this firewall and the higher-layer firewalls. On a multi-vsys firewall, each vsys requires a unique pre-shared key. f. Click **OK**. 2. Configure an Interface Management profile with the **User-ID** service enabled and assign the profile to the interface you want the firewall to use when responding to mapping information queries from firewalls in the layer above. 3. (Optional) Configure policies that are specific to the user accounts for which you want this firewall to collect mapping information. 4. **Commit** your changes.

Configure User-ID Redistribution (Continued)	
Step 4 Enable each middle layer firewall to receive mapping information from the layer below and forward it to the layer above. You must also perform this task for any firewall that redistributes mapping information to other firewalls in the same layer. For example, Figure: User-ID-Redistribution shows one data center firewall that redistributes to other data center firewalls. Each firewall can receive mapping information from up to 100 User-ID agents. Figure: User-ID-Redistribution shows only one middle layer of firewalls but you can deploy as many layers as the redistribution limit of ten hops allows.	1. Configure the firewall to receive mapping information from firewalls acting as User-ID agents in the layer below. a. Select **Device > User Identification > User-ID Agents** and click **Add**. b. Enter a **Name** to identify the lower-layer firewall. c. Enter the **Host** name or IP address of the interface that you configured on the lower-layer firewall to respond to mapping information queries. d. Enter the **Port** number (default is 5007) on which the lower-layer firewall will listen for User-ID queries. e. Enter the **Collector Name** you specified when configuring the lower-layer firewall to act as a User-ID agent. f. Enter and confirm the **Collector Pre-Shared Key** you specified on the lower-layer firewall. g. Ensure the configuration is **Enabled** (default) and click **OK**. h. Check the Connected column to confirm the firewall you just added as a User-ID agent is connected ( ). 2. Configure a service route for the firewall to use for sending mapping information queries to firewalls in the layer below. a. Select **Device > Setup > Services**. b. (Firewalls with multiple virtual systems only) Select **Global** (for a firewall-wide service route) or **Virtual Systems** (for a virtual system-specific service route). For details, refer to Customize Service Routes to Services for Virtual Systems. c. Click **Service Route Configuration**, select **Customize**, and select **IPv4** or **IPv6** depending on your network protocols. Configure the service route for both protocols if your network uses both. d. Select **UID Agent** and then select the **Source Interface** and **Source Address**. e. Click **OK** twice to save the service route. 3. Enable the firewall to forward the mapping information to firewalls in the layer above. a. Configure the firewall to function as a User-ID agent. b. Configure an Interface Management profile with the **User-ID** service enabled and assign the profile to the interface you want the firewall to use when responding to mapping information queries from firewalls in the layer above. 4. (Optional) Configure policies specific to user accounts for which you want this firewall to aggregate mapping information from lower layers. 5. **Commit** your changes.

 © Palo Alto Networks, Inc.

Configure User-ID Redistribution (Continued)		
Step 5	Enable each top-layer firewall to receive mapping information from all other layers. You must also perform this task for any firewall that is an end point in the redistribution sequence within a layer. In the example of Figure: User-ID-Redistribution, you would perform this task for the two data center firewalls that receive mapping information from another data center firewall.	1. Configure the firewall to receive mapping information from firewalls acting as User-ID agents in the layer below. 2. Configure a service route for the firewall to use for sending mapping information queries to firewalls in the layer below. 3. (Optional) Configure policies that are global to all user accounts. 4. **Commit** your changes.

Step 6	Verify that the top-layer firewalls are aggregating mapping information from all other layers.
	This step samples a single user mapping that is collected in a bottom-layer firewall and forwarded to a top-layer firewall. Repeat the step for several user mappings and several firewalls to ensure your configuration is successful.
		1. Access the CLI of a bottom-layer firewall and run the following operational command:
		`> show user ip-user-mapping all`
		2. Record the IP address associated with any username.
		3. Access the CLI of a top-layer firewall and run the following command, where `<address>` is the IP address you recorded in the previous step:
		`> show user ip-user-mapping ip <address>`
		If the firewall successfully received the user mapping from the bottom-layer firewall, it displays output similar to the following and displays the same username as you recorded in the bottom-layer firewall.
		`IP address: 192.0.2.0 (vsys1)`
		`User: corpdomain\username1`
		`From: AD`
		`Idle Timeout: 2643s`
		`Max. TTL: 2643s`
		`Groups that the user belongs to (used in policy)`

 © 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)

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.

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

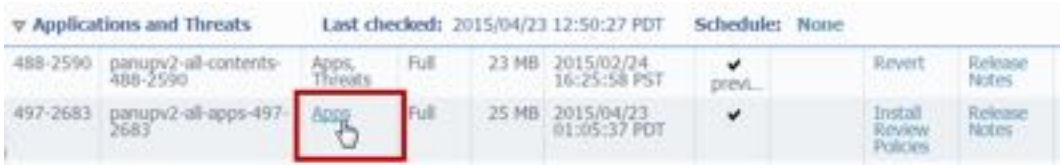
Review new App-ID signatures introduced in a Applications and/or Threats content update. For each new application signature introduced, you can preview the App-ID details, including a description of the application identified by the App-ID, other existing App-IDs that the new signature is dependent on (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:

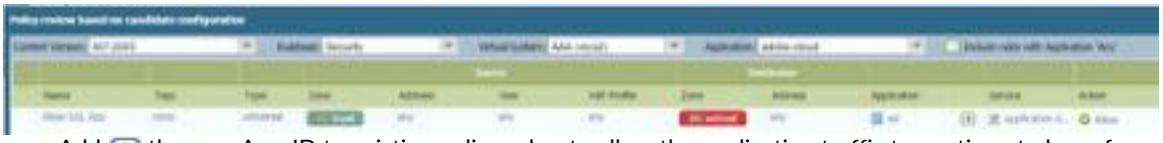
- **Depends on**—Lists the application signatures that this App-ID relies on to uniquely identify the application. If one of the application signatures listed in the **Depends On** field is disabled, the dependent App-ID is also disabled.
- **Previously Identified As**—Lists the App-IDs that matched 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. • 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. • 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. • 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**. • 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.
❑ 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:  . The new App-ID is added to the existing rules as a disabled App-ID. 4. 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. 5. **Install** the latest content release version. 6. **Commit** your changes to seamlessly update policy enforcement for new App-IDs.

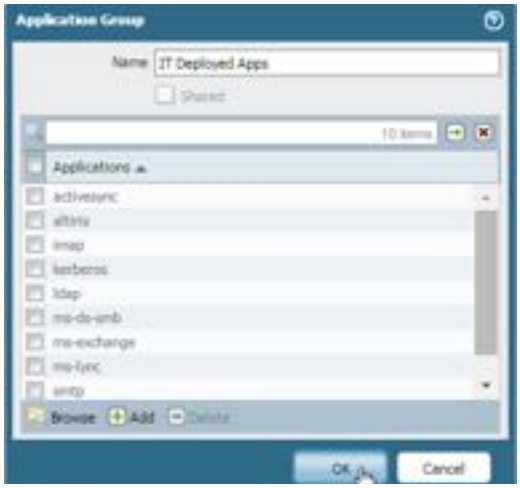
Use Application Objects in Policy

- 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 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.

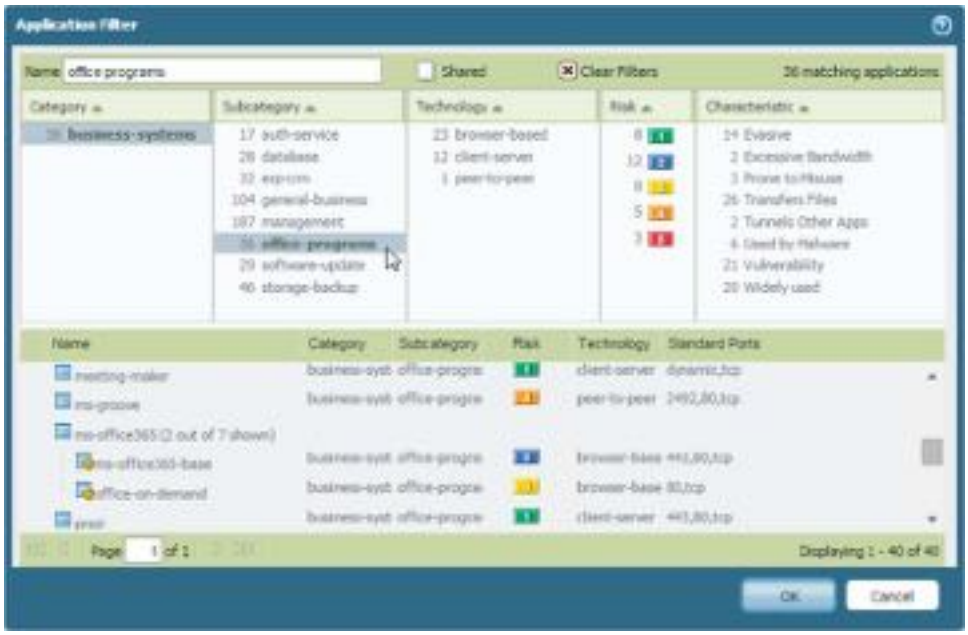
Create an Application Group	
Step 1	Select **Objects > Application Groups**.
Step 2	**Add** a group and give it a descriptive **Name**.
Step 3	(Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.
Step 4	**Add** the applications you want in the group and then click **OK**. 
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.



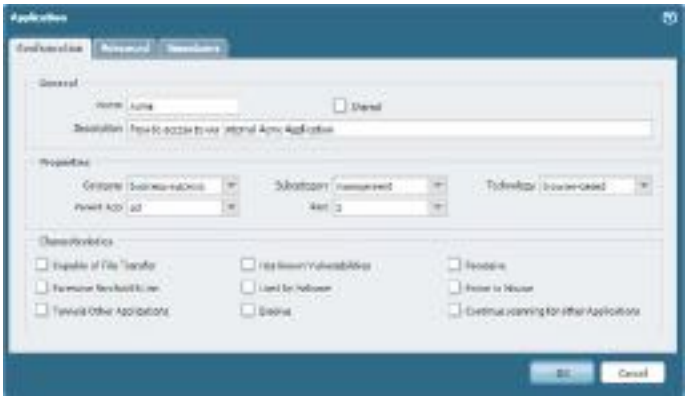
In order to collect the right data to create a custom application signature, you'll need a good understanding of packet captures and how datagrams are formed. If the signature is created too broadly, you might inadvertently include other similar traffic; if it is defined too narrowly, the traffic will evade detection if it does not strictly match the pattern.

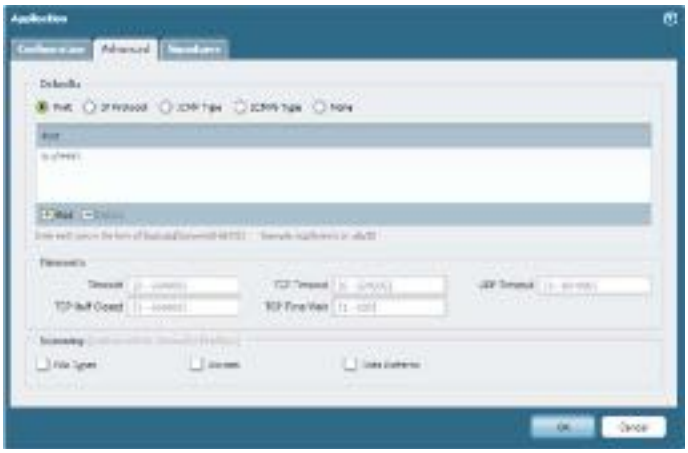
Custom applications are stored in a separate database on the firewall and this database is not impacted by the weekly App-ID updates.

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 release version 609: FTP, HTTP, IMAP, POP3, SMB, and SMTP.

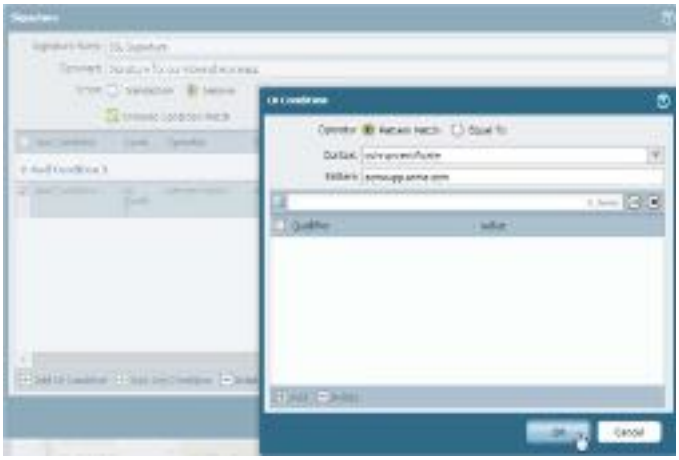
The following is a basic example of how to create a custom application.

 © 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. 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). • 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. • 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**. 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. 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. 4. Define the application Properties and Characteristics. 

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: • Specify the default ports or protocol that the application uses. • Specify the session timeout values. If you don't specify timeout values, the default timeout values will be used. • Indicate any type of additional scanning you plan to perform on the application traffic. 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. 

 © 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. 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. 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 595).

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

Application	Implicitly Supports
google-desktop	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

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.

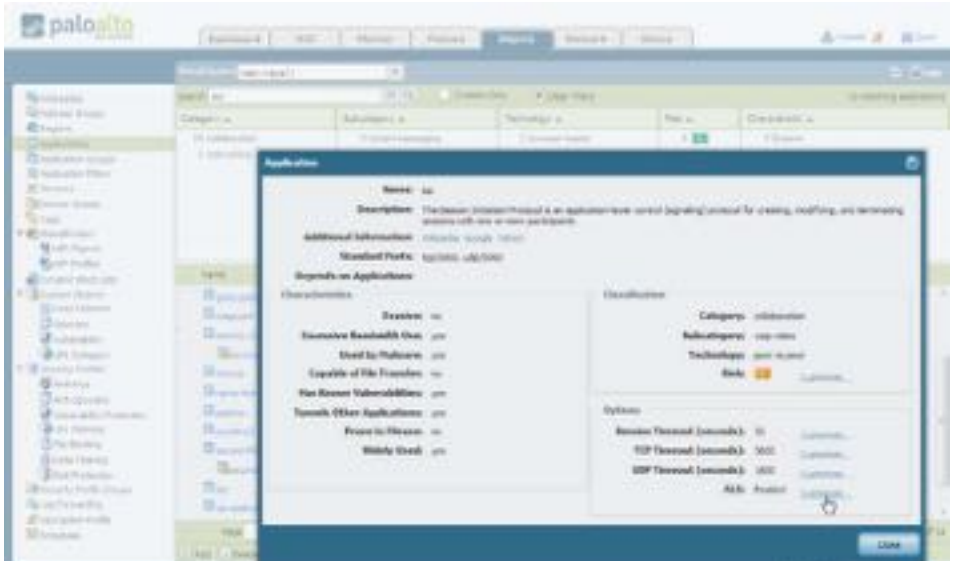
 © Palo Alto Networks, Inc.

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 DNS Proxy
- Enable Passive DNS Collection for Improved Threat Intelligence
- Use DNS Queries to Identify Infected Hosts on the Network
- DoS Protection Against Flooding of New Sessions
- Content Delivery Network Infrastructure for Dynamic Updates
- Threat Prevention Resources

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

Use the following workflow 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 subscription bundles the antivirus, anti-spyware, and the vulnerability protection features in one license. To verify that you have an active Threat Prevention subscription, 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. 2. In the **Actions** column, click **Download** to install the latest Antivirus and Applications and Threats signatures.

 © Palo Alto Networks, Inc.

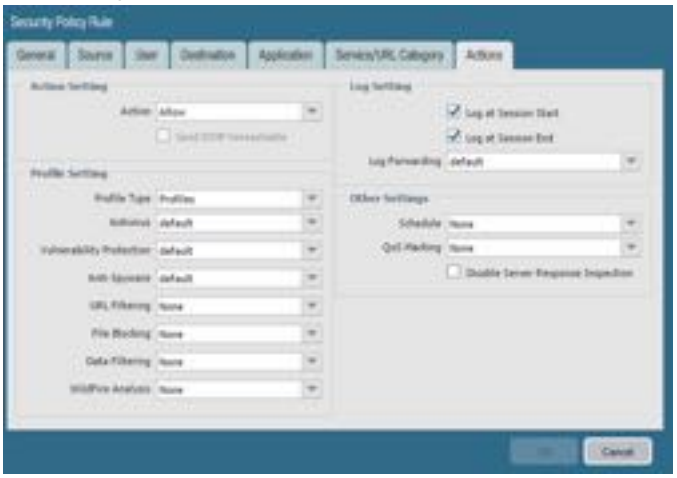
Set up Antivirus/Anti-Spyware/Vulnerability Protection (Continued)	
Step 3 Schedule signature updates.	1. From **Device > Dynamic Updates**, click the text to the right of **Schedule** to automatically retrieve signature updates for **Antivirus** and **Applications and Threats**. 2. Specify the frequency 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 firewall; you need to configure the schedule on each firewall.

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 firewalls and both firewalls will download/install independently. If you are using a data port for downloads, the passive firewall will not perform downloads while it is in the passive state. In this case you would set a schedule on both firewalls and then select the **Sync To Peer** option. This will ensure that whichever firewall is active, the updates will occur and will then push to the passive firewall.
- **Active/Active HA**—If the MGT port is used for antivirus signature downloads on both firewalls, then schedule the download/install on both firewalls, but do not select the **Sync To Peer** option. If you are using a data port, schedule the signature downloads on both firewalls and select **Sync To Peer**. This will ensure that if one firewall in the active/active configuration goes into the active-secondary state, the active firewall will download/install the signature and will then push it to the active-secondary firewall.

Set up Antivirus/Anti-Spyware/Vulnerability Protection (Continued)	
Step 4 Attach the security profiles to a Security policy rule.	1. Select **Policies > Security**, select the desired policy to modify it and then click the **Actions** tab. 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**. The default Anti-Spyware rule enables DNS Sinkholing. 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. 
Step 5 Save the configuration.	Click **Commit**.

 © Palo Alto Networks, Inc.

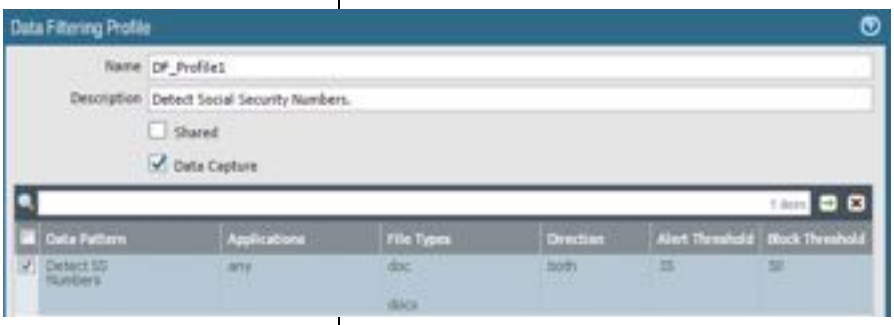
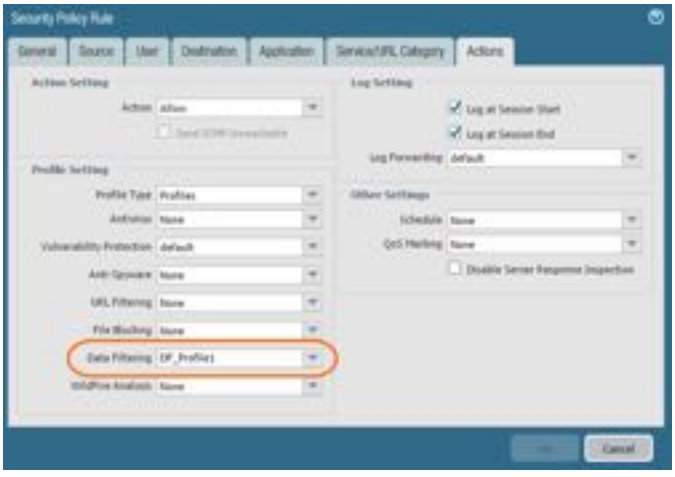
Set Up Data Filtering

Use the following workflow to configure a Data Filtering profile. This example shows a Data Filtering profile for detecting Social Security Numbers and a custom pattern 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**. 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*. 3. (Optional) If you want to collect data that is blocked by the filter, select the **Data Capture** check box.  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. 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**. 2. Click **Manage Data Protection** in the Content-ID Features section. 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. 
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. 2. Set **File Types** to **doc** and **docx** to only scan doc and docx files.

 © Palo Alto Networks, Inc.

Data Filtering Configuration Example (Continued)	
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. 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. 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. 
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. 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**. 
Step 7 **Commit** the configuration.	

Data Filtering Configuration Example (Continued)		
Step 8	Test the data filtering configuration. 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. 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. 2. Create a Microsoft Word document with one instance of the term confidential and five Social Security numbers with dashes. 3. Upload the file to a website. Use an HTTP site unless you have decryption configured, in which case you can use HTTPS. 4. Select **Monitoring > Logs > Data Filtering** logs. 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.

Set Up File Blocking

The following workflow shows how to set up a basic File Blocking profile. This example shows how to set up a profile that prompts users to continue before downloading .exe files from websites.

Configure File Blocking		
Step 1	Create the file blocking profile.	1. Select **Objects > Security Profiles > File Blocking** and click **Add**. 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. 2. Enter a **Name**, such as *BlockEXE*. 3. Set the **Applications** for filtering, for example **web-browsing**. 4. Set **File Types** to **exe**. 5. Set the **Direction** to **download**. 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. 7. Click **OK** to save the profile.

 © Palo Alto Networks, Inc.

Configure File Blocking (Continued)	
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 a Basic Security Policy. 2. Click the **Actions** tab within the policy rule. 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*. 4. **Commit** the configuration. 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.
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: 	
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.	

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 Install Content and Software 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:

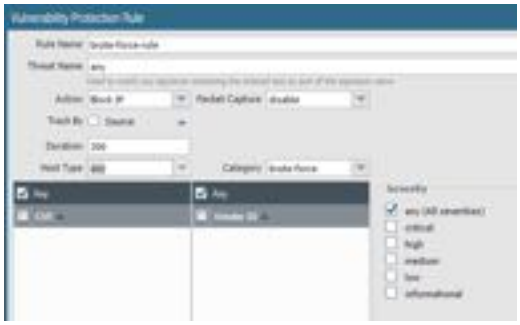
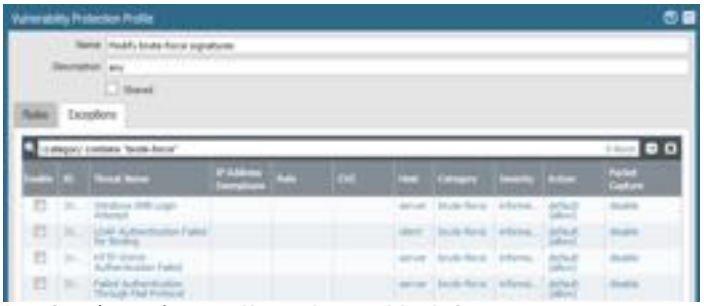
- Create a rule to modify the default action for all signatures in the brute force category. You can define the action to allow, alert, block, reset, or drop the traffic.
- Define an exception for a specific signature. For example, you can search for a CVE and define an exception for it.

  For a parent signature, you can modify both the trigger conditions and the action; for a child signature you can modify the action only.



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**. 2. Click **Add** and enter a **Name** for the Vulnerability Protection profile.

Customize the Threshold and Action for a Signature	
Step 2 Create a rule that defines the action for all signatures in a category.	1. Select **Rules**, click **Add** and enter a **Name** for the rule. 2. Set the **Action**. In this example, it is set to **Block IP**. 3. Set **Category** to **brute-force**. 4. (Optional) If blocking, specify whether to block based on **Host Type** server or client, the default is any. 5. See Step 3 to customize the action for a specific signature. 6. See Step 4 to customize the trigger threshold for a parent signature.  7. Click **OK** to save the rule and the profile.
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**.

 © Palo Alto Networks, Inc.

Customize the Threshold and Action for a Signature		
Step 4	Customize the trigger conditions for a parent signature. A parent signature that can be edited is marked with this icon: . 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. 2. To modify the trigger threshold specify the **Number of Hits** *per x* **seconds**. 3. Specify whether to aggregate the number of hits by **source**, **destination** or by **source and destination**. 4. Click **OK**.
Step 5	Attach this new profile to a security rule.	1. Select **Security > Policies**. 2. Modify an existing security policy rule or **Add** a new rule. 3. Select **Actions**. 4. In the Profile Setting section, set the **Profile Type** to **Profiles**. 5. Select the newly-created **Vulnerability Protection** profile. 6. Click **OK** to save changes to the security policy rule.
Step 6	Save your changes.	1. Click **Commit**.

Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions

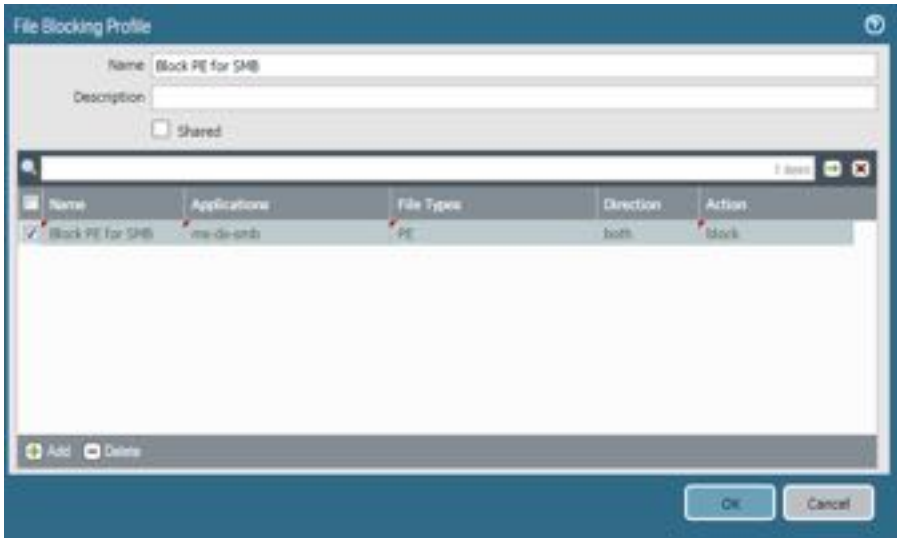
To monitor and protect your network from most Layer 4 and Layer 7 attacks, here are a few recommendations.

- Upgrade to the most current PAN-OS software version and content release version to ensure that you have the latest security updates. For evasion prevention, upgrade to PAN-OS 7.1.1 and Applications and Threats content release version 579. See Install Content and Software Updates.
- Set up the firewall to act as a DNS proxy and enable evasion signatures:
  - Enable DNS Proxy.

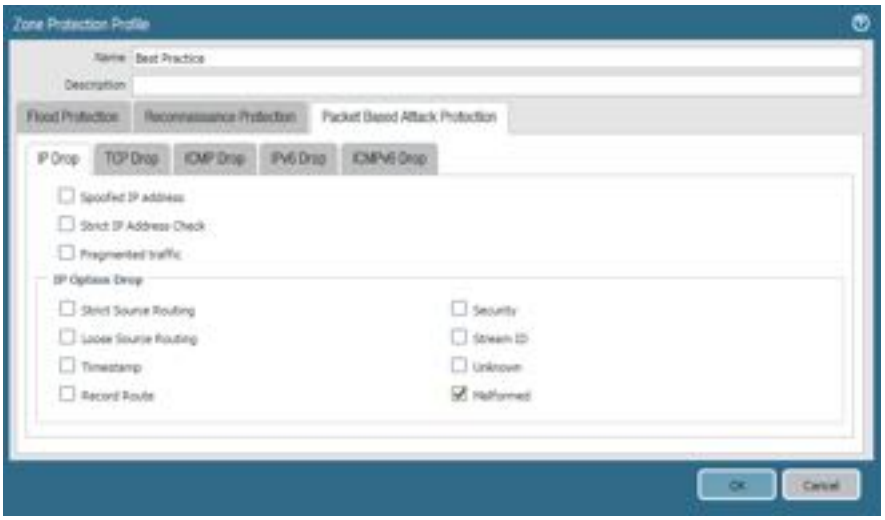
    When acting as a DNS proxy, the firewall resolves DNS requests and caches hostname-to-IP-address mappings in order to quickly and efficiently resolves future DNS queries.
  - Enable evasion signatures.

    Evasion signatures that detect crafted HTTP or TLS requests can alert when a client connects to a domain other than the domain specified in the original DNS request. Make sure that DNS proxy is configured if you choose to enable evasion signatures. Without DNS proxy enabled, evasion signatures can trigger when a DNS server in DNS load balancing configuration returns different IP addresses (for servers hosting identical resources) to the firewall and client in response to the same DNS request.
- 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 set 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 policy rules 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 with severity low and higher.
  - 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).

 © Palo Alto Networks, Inc.


File Blocking Profile
Name
Block PE for SMB
Description
Shared
Name
Applications
File Types
Direction
Action
Block PE for SMB
both
block
Add
Delete
OK
Cancel

- Create a Zone Protection profile that is configured to protect against packet-based attacks (**Network > Network Profiles > Zone Protection**):
  - Select the option to drop **Malformed** IP packets (**Packet Based Attack Protection > IP Drop**).



  - Remove TCP timestamps on SYN packets before the firewall forwards the packet. When you select the **Remove 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. (**Packet Based Attack Protection > TCP Drop**).
  - Select the option to drop **Mismatched overlapping TCP segment**. 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. Selecting this option causes PAN-OS 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.



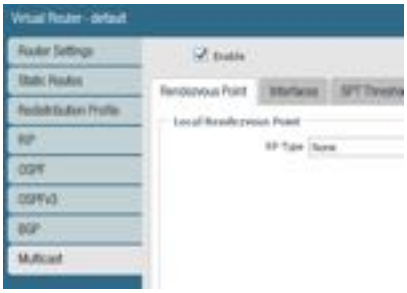
 © Palo Alto Networks, Inc.

- 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**).



- Configure the firewall to **Clear** the **Urgent Data Flag** in the TCP header (**Device > Setup > Session > TCP Settings**).

  Many hosts use the urgent data flag in the TCP header to promote a packet for immediate processing, removing it from the processing queue and expediting it through the TCP/IP stack. This process is called out-of-band processing. However, the implementation of the urgent data flag varies from host to host. Configuring the firewall to clear this flag eliminates ambiguity in how the packet is processed on the firewall and the host, allowing the firewall sees the same stream in the protocol stack as the host for which the packet is destined. When the firewall clears this flag, it includes it in the payload and prevents the packet from being processed urgently.

- Enable the **Drop segments without flag** option (**Device > Setup > Session > TCP Settings**).

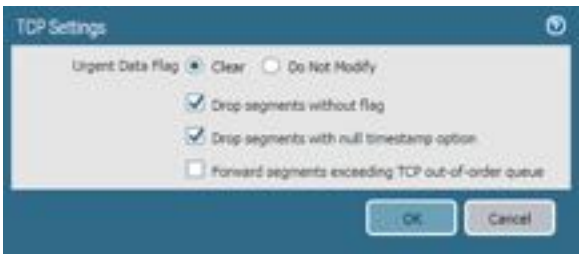
  Illegal TCP segments without any flags set can be used to evade content inspection. When you enable this option, the firewall will drop packets that have no flags set in the TCP header.

- Enable the **Drop segments with null timestamp** option (**Device > Setup > Session > TCP Settings**).

  The TCP timestamp records when the segment was sent and allows the firewall to verify that the timestamp is valid for that session, preventing TCP sequence number wrapping. The TCP timestamp is also used to calculate round trip time. When a TCP Timestamp is set to 0 (null) it could confuse either end of the connection, resulting in an evasion. The firewall drops packets with null timestamps with this setting enabled.

- Disable the **Forward segments exceeding TCP out-of-order queue** option (**Device > Setup > Session > TCP Settings**).

  By default, the firewall forwards segments that exceed the TCP out-of-order queue limit of 64 per session. By disabling this option, the firewall instead drops segments that exceed the out-of-order queue limit.

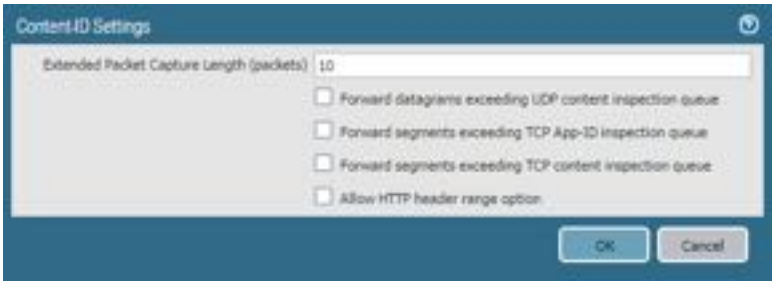


- Disable the **Forward segments exceeding TCP App-ID inspection queue** option (**Device > Setup > Content-ID > Content-ID Settings**).

  By default, when the App-ID inspection queue is full the firewall skips App-ID inspection—classifying the application as unknown-tcp—and forwards the segments. By disabling this option, the firewall instead drops segments when the App-ID inspection queue is full.
- Disable the **Forward datagrams exceeding UDP content inspection queue** and **Forward segments exceeding TCP content inspection queue** options (**Device > Setup > Content-ID > Content-ID Settings**).

  By default, when the TCP or UDP content inspection queue is full the firewall skips Content-ID inspection for TCP segments or UDP datagrams that exceed the queue limit of 64. By disabling these options, the firewall instead drops TCP segments and UDP datagrams when the corresponding TCP or UDP content inspection queue is full.
- Disable the **Allow HTTP Header Range Option** (**Device > Setup > Content-ID > Content-ID Settings**).

  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. Disabling this option prevents this from happening.



 © Palo Alto Networks, Inc.

Enable DNS Proxy

Domain name system (DNS) servers translate user-friendly domains to the associated IP addresses which locate and identify the corresponding resources. A Palo Alto Networks firewall intermediate to clients and servers can act as a DNS proxy to resolve domain name queries.

The DNS proxy feature enables the firewall to:

- Quickly, efficiently, and locally resolve domain name queries based on static and cached DNS entries.
- Reach out to specific DNS servers to resolve certain types of DNS requests (for example, the firewall can resolve corporate domains based on a corporate DNS server hostname-to-IP-address mappings, and resolve other domains using a public or ISP DNS server).

Enable the Firewall to Act as a DNS Proxy		
Step 1	Specify the interfaces on which you want the firewall to listen for DNS requests.	1. Select **Network > DNS Proxy** and **Add** a new object. 2. Verify that **Enable** is selected and **Name** the object. 3. **Add** one or more **Interface** on which the firewall listens for DNS requests. 4. (Virtual Systems Only) Allow the DNS proxy object to be shared across all virtual systems, or set the **Location** to apply the DNS proxy object settings to a specific virtual system.
Step 2	Define the DNS server with which the firewall should communicate to resolve DNS requests.	If you are enabling DNS proxy on a virtual system, you must select **New** in the Server Profile drop-down first, and then continue with either of the following options. Specify DNS Servers 1. Set **Inheritance Source** to **none**. 2. Enter a the **Primary** DNS server IP address or address object. 3. Enter the **Secondary** DNS server IP address or address object. Use Inherited DNS Servers Select an **Inheritance Source** from which the firewall can use existing DNS server settings for the DNS proxy object. Only interfaces configured to be DHCP client interfaces and PPPoE client interfaces are available as inheritance sources for DNS server settings. In this case, the DNS server settings the client interface dynamically receives from a DHCP server are also used to populate the **Primary** and **Secondary** DNS server settings (just continue to set both of these fields to **inherited**).

Enable the Firewall to Act as a DNS Proxy (Continued)		
Step 3	Enable the firewall to reach out to certain DNS servers to resolve specific domains. For example, the firewall can forward corporate domains to a corporate DNS server for domain name resolution.	1. Select **DNS Proxy Rules**, **Add** a rule, and give the rule a descriptive **Name**. 2. **Turn on caching of domains resolved by this mapping** to enable the firewall to save recently resolved DNS queries in order to quickly resolve future matching queries. 3. Add one or more **Domain Name**. 4. Enter the IP addresses or address objects for the **Primary** and **Secondary** DNS servers. The firewall communicates with these servers to resolve DNS requests for the listed domain names.  If you are enabling DNS proxy on a virtual system, you can instead configure a DNS Server Profile to define DNS settings for the virtual system, including the primary and secondary DNS server.
Step 4	Set up static FQDN-to-IP address entries that the firewall can resolve locally, without having to reach out to a DNS server.	1. Select **Static Entries**. 2. **Add** and **Name** a new static mapping entry. 3. Enter the **FQDN** that you want the firewall to resolve. 4. **Add** one or more IP **Address** to map to the domain you entered in the last step.
Step 5	Enable caching for resolved hostname-to-IP-address mappings, and customize additional DNS settings.	Select **Advanced** and configure settings to: • Store recently resolved hostname-to-IP-address mappings. Select **Cache** and continue to specify the number of entries for the cache to hold and the number of hours after which all cached DNS entries are removed. • Enable DNS queries using TCP. • Specify settings for UDP query retries.

 © Palo Alto Networks, Inc.

Enable the Firewall to Act as a DNS Proxy (Continued)	
Step 6 Enable evasion signatures.  When DNS proxy is enabled, evasion signatures that detect crafted HTTP or TLS requests can alert to instances where a client connects to a domain other than the domain specified in the original DNS query.	1. Install the Applications and Threats content version 579 or later: a. Select **Device > Dynamic Updates**. b. **Check Now** to get the latest Applications and Threats content update. c. Download and Install Applications and Threats content version 579. 2. Define how traffic matched to evasion signatures should be enforced: a. Select **Objects > Security Profiles > Anti-Spyware** and **Add** or modify an Anti-spyware profile. b. Select **Exceptions** and select **Show all signatures**. c. Filter signatures based on the keyword `evasion`. d. For all evasion signatures, set the **Action** to any setting other than allow or the default action (the default action is for evasion signatures is allow). For example, set the action to alert on or block. e. Click **OK** to save the updated Anti-spyware profile. f. Attach the Anti-spyware profile to a security policy rule: Select **Policies > Security**, select the desired policy to modify and then click the **Actions** tab. In Profile Settings, click the drop-down next to **Anti-Spyware** and select the anti-spyware profile you just modified to enforce evasion signatures.
Step 7 **Commit** your changes.	
Learn more about DNS features...	❑ Use DNS queries to identify infected hosts on the network. ❑ Enable passive DNS collection for better threat intelligence. ❑ To work with DNS features and virtual systems, see these DNS use cases for virtual systems and learn how to configure a DNS proxy object and DNS server profiles for virtual systems.

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.  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 or to a custom domain so that you can identify hosts on your network that have been infected with malware. By default, DNS queries to any domain included in the Palo Alto Networks DNS signatures list is sinkholed to a Palo Alto Networks server IP address. The following topics provide details on how to enable DNS sinkholing for custom domains and how to identify infected hosts.

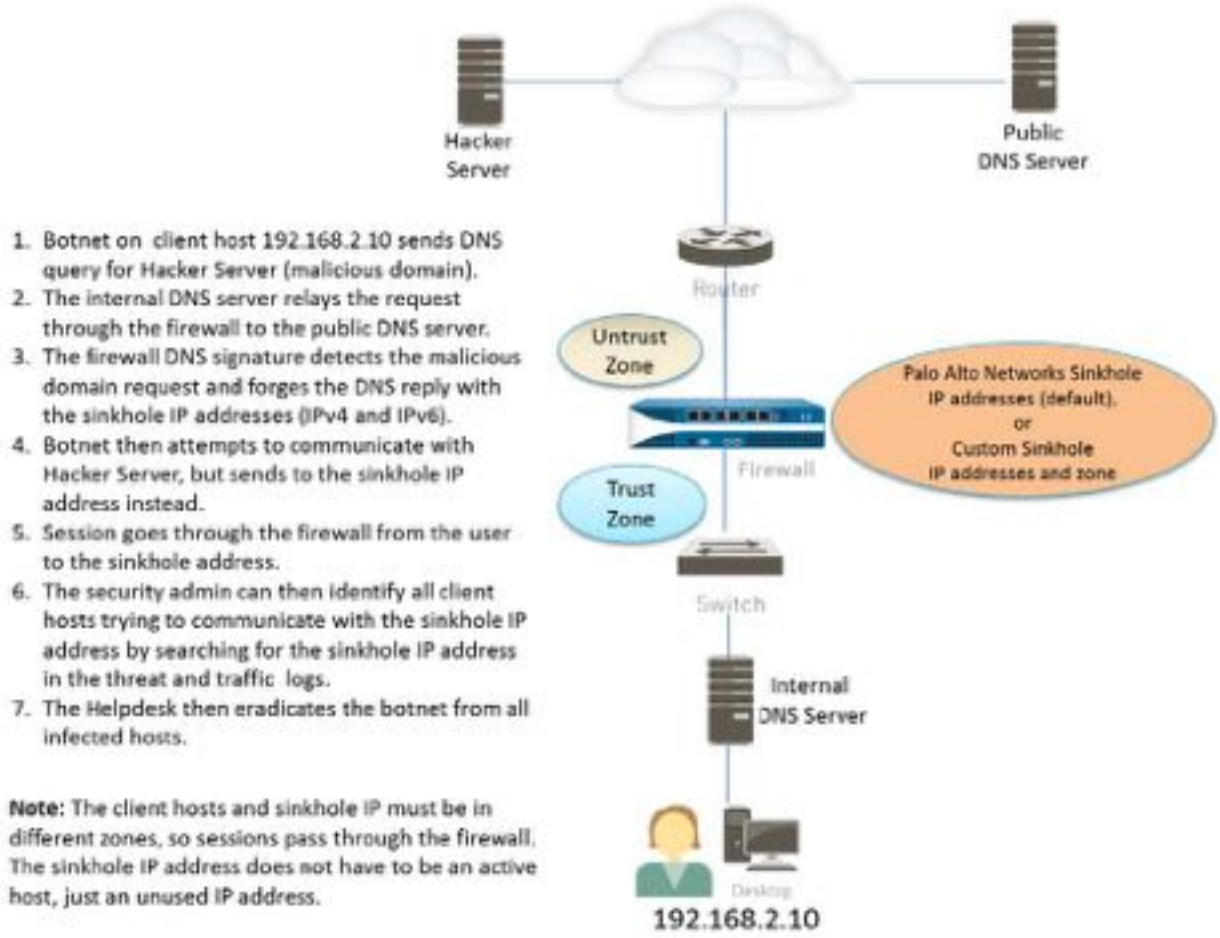
- DNS Sinkholing
- Configure DNS Sinkholing for a List of Custom Domains
- Configure the Sinkhole IP Address to a Local Server on Your Network
- Identify Infected Hosts

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 default Palo Alto Networks sinkhole IP address, or to a user-defined IP address as illustrated in Configure DNS Sinkholing for a List of Custom Domains. Infected hosts can then be easily identified in the traffic logs because any host that attempts to connect to the sinkhole IP address is most likely infected with malware.

If you want to enable DNS sinkholing for Palo Alto Networks DNS signatures, attach the default Anti-Spyware profile to a security policy rule (see Set Up Antivirus, Anti-Spyware, and Vulnerability Protection). DNS queries to any domain included in the Palo Alto Networks DNS signatures will be resolved to the default Palo Alto Networks sinkhole IP address. The IP addresses currently are IPv4—71.19.152.112 and a loopback address IPv6 address—::1. These address are subject to change and can be updated with content updates.

Figure: DNS Sinkholing Example



 © Palo Alto Networks, Inc.

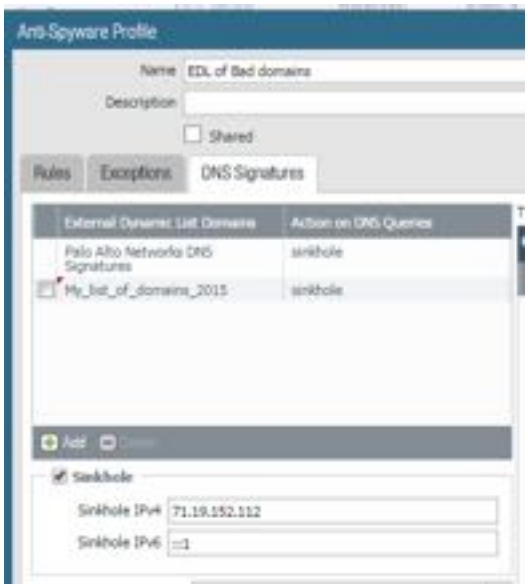
Configure DNS Sinkholing for a List of Custom Domains

To enable DNS Sinkholing for a custom list of domains, you must create an external dynamic list that includes the domains, enable the sinkhole action in an Anti-Spyware profile and attach the profile to a security policy rule. When a client attempts to access a malicious domain in the list, the firewall forges the destination IP address in the packet to the default Palo Alto Networks server or to a user-defined IP address for sinkholing.

For each custom domain included in the external dynamic list, the firewall generates DNS-based spyware signatures. The signature is named Custom Malicious DNS Query <domain name>, and is of type spyware with medium severity; each signature is a 24-byte hash of the domain name.

Each firewall platform supports a maximum of 50,000 domain names total in one or more External Dynamic List but no maximum limit is enforced for any one list.

Configure DNS Sinkholing for a Custom List of Domains	
Step 1 Enable DNS sinkholing for the custom list of domains in an external dynamic list.	1. Select **Objects > Security Profiles > Anti-Spyware**. 2. Modify an existing profile, or select one of the existing default profiles and clone it. 3. **Name** the profile and select the **DNS Signatures** tab. 4. **Add** an **External Dynamic List**. When you configure the external dynamic list from the Anti-Spyware profile, the **Type** is preset to **Domain List**.  If you have already created an external dynamic list of type: Domain List, you can select it from here. The drop-down does not display external dynamic lists of type URL or IP Address that you may have created. 5. Configure access to the external dynamic list. a. Enter a descriptive **Name** for the list. b. Enter the **Source** for the list you just created on the web server. The source must include the full path to access the list. For example, https://1.2.3.4/EDL_IP_2015. c. Populate the list with domain names. See Formatting Guidelines for an External Dynamic List. d. Click **Test Source URL** to verify that the firewall can connect to the list on the web server.  If the web server is unreachable after the connection is established, the firewall or Panorama uses the last successfully retrieved list for enforcing policy until the connection is restored with the web server. e. (Optional) Specify the **Repeat** frequency at which the firewall retrieves the list. By default, the list is retrieved once every hour. f. Click **OK**. 6. (Optional) In the **Packet Capture** drop-down, select **single-packet** to capture the first packet of the session or **extended-capture** to set between 1-50 packets. You can then use the packet captures for further analysis.

Configure DNS Sinkholing for a Custom List of Domains (Continued)	
Step 2 Verify the sinkholing settings on the Anti-Spyware profile.	7. On the **DNS Signatures** tab, verify that the **Action on DNS Queries** is **sinkhole**. 8. In the Sinkhole section, verify that **Sinkhole** is enabled. For your convenience, the default Sinkhole IP address is set to access a Palo Alto Networks server. Palo Alto Networks can automatically refresh this IP address through content updates. If you want to modify the **Sinkhole IPv4** or **Sinkhole IPv6** address to a local server on your network or to a loopback address, see Configure the Sinkhole IP Address to a Local Server on Your Network.  9. Click **OK** to save the Anti-Spyware profile.
Step 3 Attach the Anti-Spyware profile to a Security policy rule.	1. Select **Policies > Security**. 2. On the **Actions** tab, select the **Log at Session Start** check box to enable logging. 3. In the Profile Setting section, click the **Profile Type** drop-down to view all **Profiles**. From the **Anti-Spyware** drop-down and select the new profile. 4. Click **OK** to save the policy rule.
Step 4 Test that the policy action is enforced.	1. Access a domain in the external dynamic list. 2. To monitor the activity on the firewall: a. Select **ACC** and add a URL Domain as a global filter to view the Threat Activity and Blocked Activity for the domain you accessed. b. Select **Monitor > Logs > Threat** and filter by `(action eq sinkhole)` to view logs on sinkholed domains.

 © Palo Alto Networks, Inc.

Configure DNS Sinkholing for a Custom List of Domains (Continued)		
Step 5	Verify whether entries in the external dynamic list are ignored or skipped. In a list of type URL, the firewall skips entries that are not URLs as invalid and ignores entries that exceed the maximum limit for the platform.	Use the following CLI command on the firewall to review the details about the list. **request system external-list show type domain name <list_name>** For example: request system external-list show type domain name My_List_of_Domains_2015 vsys1/EBLDomain: Next update at : Thu May 21 10:15:39 2015 Source :https://1.2.3.4/My_List_of_Domains_2015 Referenced : Yes Valid : Yes Number of entries : 3 domains: www.example.com baddomain.com qqq.abcedfg.com
Step 6	(Optional) Retrieve the external dynamic list on-demand.	To force the firewall to retrieve the updated list on-demand instead of at the next refresh interval (the **Repeat** frequency you defined for the external dynamic list), use the following CLI command: request system external-list refresh type domain name <list_name>

Configure the Sinkhole IP Address to a Local Server on Your Network

By default, sinkholing is enabled for all Palo Alto Networks DNS signatures, and the sinkhole IP address is set to access a Palo Alto Networks server. Use the instructions in this section if you want to set the sinkhole IP address to a local server on your network.

You must obtain both an IPv4 and IPv6 address to use as the sinkhole IP addresses because malicious software may perform DNS queries using one or both of these protocols. 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 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:

IPv4 DNS sinkhole address—10.15.0.20

IPv6 DNS sinkhole address—fd97:3dec:4d27:e37c:5:5:5:5

Configure Sinkholing to a Local Server on Your Network		
Step 1	Configure the sinkhole interface and zone. 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.  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. 2. In the **Interface Type** drop-down, select **Layer3**. 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. 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. 5. Click **OK** to save. 6. To add a zone for the sinkhole, select **Network > Zones** and click **Add**. 7. Enter zone **Name**. 8. In the **Type** drop-down select **Layer3**. 9. In the **Interfaces** section, click **Add** and add the interface you just configured. 10. Click **OK**.
Step 2	Enable DNS sinkholing.	By default, sinkholing is enabled for all Palo Alto Networks DNS signatures. To change the sinkhole address to your local server, see step 8 in Configure DNS Sinkholing for a List of Custom Domains.
Step 3	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. Editing the security rule(s) that allows traffic from client hosts in the trust zone to the untrust zone ensures that you are identifying traffic from infected hosts. 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**. 2. Select an existing rule that allows traffic from the client host zone to the untrust zone. 3. On the **Destination** tab, **Add** the Sinkhole zone. This allows client host traffic to flow to the sinkhole zone. 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. 5. In the **Profile Setting** section, select the **Anti-Spyware** profile in which you enabled DNS sinkholing. 6. Click **OK** to save the security rule and then **Commit**.

 © Palo Alto Networks, Inc.

Configure Sinkholing to a Local Server on Your Network	

Step 4 To confirm 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.
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:
	`C:\>ping <sinkhole address>`
	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:
	`C:\>ping 10.15.0.20`
	`Pinging 10.15.0.20 with 32 bytes of data:`
	`Request timed out.`
	`Request timed out.`
	`Ping statistics for 10.15.0.20:`
	`Packets: Sent = 4, Received = 0, Lost = 4 (100% loss)`
	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.
	  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.

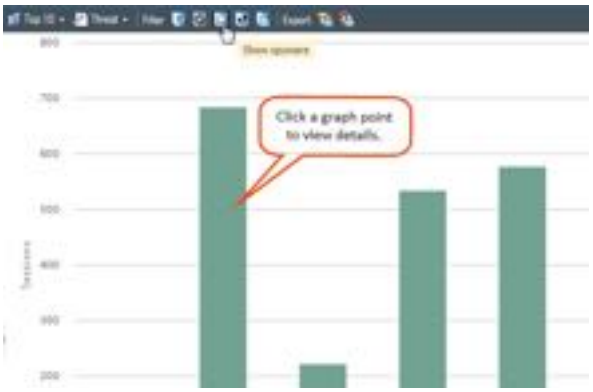
Configure Sinkholing to a Local Server on Your Network	

Step 5 Test that DNS sinkholing is configured properly.
You are simulating the action that an infected client host would perform when a malicious application attempts to call home.
	1. Find a malicious domain that is included in the firewall's current Antivirus signature database to test sinkholing.
	a. Select **Device > Dynamic Updates** and in the **Antivirus** section click the **Release Notes** link for the currently installed antivirus database. You can also find the antivirus release notes that list the incremental signature updates under Dynamic Updates on the Palo Alto Networks support site.
	b. 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 display the domain name. For example, Antivirus release 1117-1560, includes an item in the left column named "tbsbana" and the right column lists "net".
	The following shows the content in the release note for this line item:
	`conficker:tbsbana1 variants: net`
	2. From the client host, open a command prompt.
	3. Perform an NSLOOKUP to a URL that you identified as a known malicious domain.
	For example, using the URL `track.bidtrk.com`:
	`C:\>nslookup track.bidtrk.com`
	`Server: my-local-dns.local`
	`Address: 10.0.0.222`
	`Non-authoritative answer:`
	`Name: track.bidtrk.com.org`
	`Addresses: fd97:3dec:4d27:e37c:5:5:5:5`
	`10.15.0.20`
	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.
	4. Select **Monitor > Logs > Threat** and locate the corresponding threat log entry to verify that the correct action was taken on the NSLOOKUP request.
	5. 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	
• Use App Scope to identify infected client hosts.	1. Select **Monitor > App Scope** and select **Threat Monitor**. 2. Click the **Show spyware** button along the top of the display page. 3. Select a time range. 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. 

DNS Sinkhole Verification and Reporting (Continued)

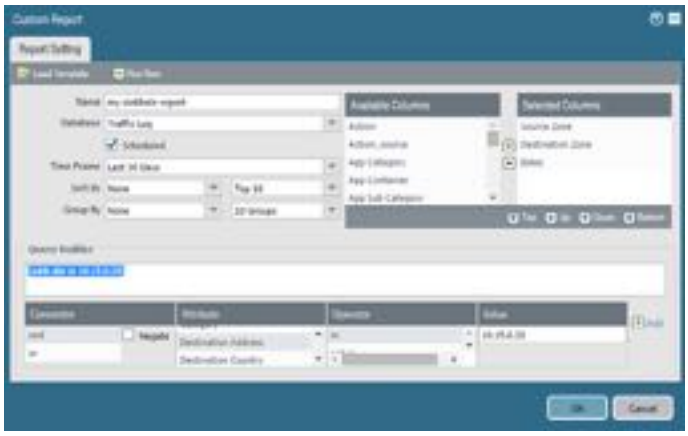
- Configure a custom report to identify all client hosts that have sent traffic to the sinkhole IP address, which is 10.15.0.20 in this example.

  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.

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.



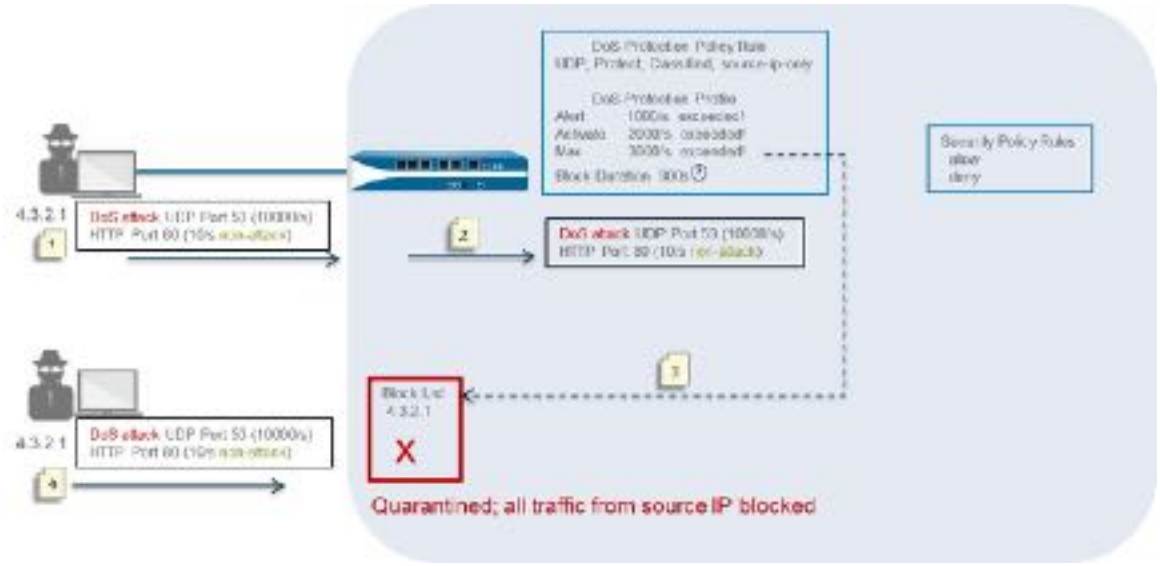
This feature defends only against DoS attacks of new sessions, that is, traffic that has not been offloaded to hardware. An offloaded attack is not protected by this feature. However, this topic describes how you can create a Security policy rule to reset the client; the attacker reinitiates the attack with numerous connections per second and is blocked by the defenses illustrated in this topic.

- Multiple-Session DoS Attack
- Single-Session DoS Attack
- 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

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 Activate Rate, the firewall takes the action specified in the DoS Protection profile.

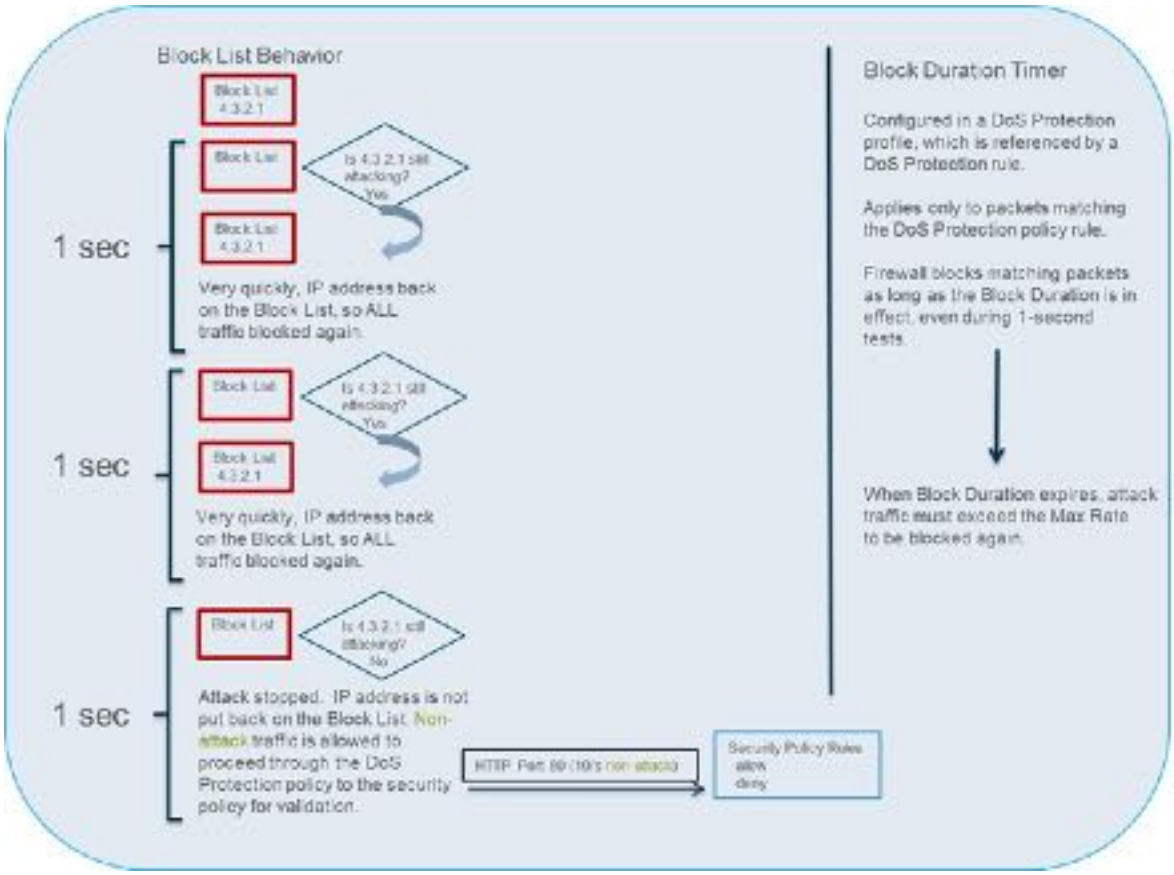
The following figure and table describe how the Security policy rules, DoS Protection policy rules and profile work together in an example.



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. 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. 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: • the threshold is exceeded, • a Block Duration is specified, and • Classified is set to include source IP address, 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 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.

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) 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. This step is one of the steps typically performed to stop an existing attack. See Use the CLI to End a Single Attacking Session.	• Components of a Security Policy Rule • Create a Security Policy Rule

 © Palo Alto Networks, Inc.

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 all types of flood protection: • **SYN Flood** • **UDP Flood** • **ICMP Flood** • **ICMPv6 Flood** • **Other IP Flood** 4. When you enable **SYN Flood**, select the **Action** that occurs when the **Activate Rate** threshold is exceeded: **Random Early Drop** or **SYN Cookies**. 5. (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 0-2,000,000; default is 10,000.) • **Activate Rate (packets/s)**—Specify the threshold rate (pps) above which a DoS response is activated. When the **Activate Rate** threshold is reached, **Random Early Drop** occurs. (Range is 0-2,000,000; default is 10,000.) • **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. (Range is 2-2,000,000; default is 40,000.)  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. 6. 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 1-21,600; 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. 7. Click **OK**.

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. 2. On the **Source** tab, choose the **Type** to be a **Zone** or **Interface**, and then **Add** the zone(s) or interface(s). 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. 4. (Optional) For **Source User**, select **any** or specify a user. 5. (Optional) Select **Negate** to match any sources except those you specify. 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. 7. (Optional) For **Destination Address**, select **Any** or enter the IP address of the device you want to protect. 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. 9. On the **Option/Protection** tab, for **Action**, select **Protect**. 10. Select **Classified**. 11. For **Profile**, select the name of the **DoS Protection** profile you created. 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. • 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. • 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. 13. Click **OK**.
Step 4	Save the configuration.	Click **Commit**.

 © Palo Alto Networks, Inc.

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. 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. Alternatively, if you know the session ID, you can execute the `clear session id <value>` command to end that session only. 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.

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. 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 buffer with TCP packets (protocol 6) coming from source IP address 192.168.2.35.

- **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.
- **GRP-ID**—Indicates an internal stage of processing packets.
- **COUNT**—Indicates how many packets are in that GRP-ID for that session.
- **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
```

 © Palo Alto Networks, Inc.

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 not offloaded, you can Use the CLI to End a Single Attacking Session. At a minimum, you can Configure DoS Protection Against Flooding of New Sessions.
- On a hardware platform that has a field-programmable gate array (FPGA), the firewall offloads traffic to the FPGA when possible to increase performance. If the traffic is offloaded to hardware, clearing the session does not help because then it is the software that must handle the barrage of packets. You should instead Discard a Session Without a Commit.

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-500> show session id 60088184

Session       60088184

        c2s flow:
                source:      1.1.42.15 [trust]
                dst:         1.2.27.99
                proto:       6
                sport:       55093           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:      55093
                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
```

Discard a Session Without a Commit

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 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: `admin@PA-7050> request session-discard [timeout <seconds>] [reason <reason-string>] id <session-id>` The default timeout is 3600 seconds.
Step 2	Verify that sessions have been discarded. `admin@PA-7050> show session all filter state discard`

 © 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 firewalls. The firewalls access the web resources in the CDN to perform various App-ID and Content-ID functions. For enabling and scheduling the content updates, see Install Content and Software 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 • downloads.paloaltonetworks.com:443 As a best practice, set the update server to updates.paloaltonetworks.com. This allows the Palo Alto Networks firewall 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 Resolves to the primary URL s0000.urlcloud.paloaltonetworks.com and is then redirected to the regional server that is closest: • s0100.urlcloud.paloaltonetworks.com • s0200.urlcloud.paloaltonetworks.com • s0300.urlcloud.paloaltonetworks.com • 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 • 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
	• beta-s1.wildfire.paloaltonetworks.com:443/80
	  Beta sites are only accessed by a firewall running a Beta release version.
	• mail.wildfire.paloaltonetworks.com:25
	• wildfire.paloaltonetworks.com:443/80
		• mail.wildfire.paloaltonetworks.com:25 or the IP address 54.241.16.83
		• wildfire.paloaltonetworks.com:443/80 or 54.241.8.199
		• staticupdates.paloaltonetworks.com or the IP address 199.167.52.15 (the Palo Alto Networks update server delivers content updates to the firewall)
		• The regional URL/IP addresses for WildFire submission queues are:
		  • ca-s1.wildfire.paloaltonetworks.com:44 or 54.241.34.71
		  • va-s1.wildfire.paloaltonetworks.com:443or 174.129.24.252
		  • eu-s1.wildfire.paloaltonetworks.com:443 or 54.246.95.247
		  • sg-s1.wildfire.paloaltonetworks.com:443or 54.251.33.241
		  • jp-s1.wildfire.paloaltonetworks.com:443 or 54.238.53.161
		  • portal3.wildfire.paloaltonetworks.com:443/80 or 54.241.8.199
		  • ca-s3.wildfire.paloaltonetworks.com:443 or 54.241.34.71
		  • va-s3.wildfire.paloaltonetworks.com:443 or 23.21.208.35
		  • eu-s3.wildfire.paloaltonetworks.com:443 or 54.246.95.247
		  • sg-s3.wildfire.paloaltonetworks.com:443 or 54.251.33.241
		  • jp-s3.wildfire.paloaltonetworks.com:443 or 54.238.53.161
		  • wildfire.paloaltonetworks.com.jp:443/80 or 180.37.183.53
		  • wf1.wildfire.paloaltonetowrks.jp:443 or 180.37.180.37
		  • wf2.wildfire.paloaltonetworks.jp:443 or 180.37.181.18
		  • 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
- Define Traffic to Decrypt
- 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

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 firewall). 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, WildFire Submissions, 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:

- 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 you can attach a decryption profile to a policy rule 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 with a private key.

With a decryption policy configured, a 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 firewall trusts, use the **Device > Certificate Management > Certificates > Default Trusted Certificate Authorities** tab on the firewall web interface.

Table: Palo Alto Networks Firewall Keys and Certificates describes the different keys and certificates used by Palo Alto Networks firewalls for decryption. As a best practice, use different keys and certificates for each usage.

Table: Palo Alto Networks Firewall 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 4 in the Configure SSL Forward Proxy task.
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.

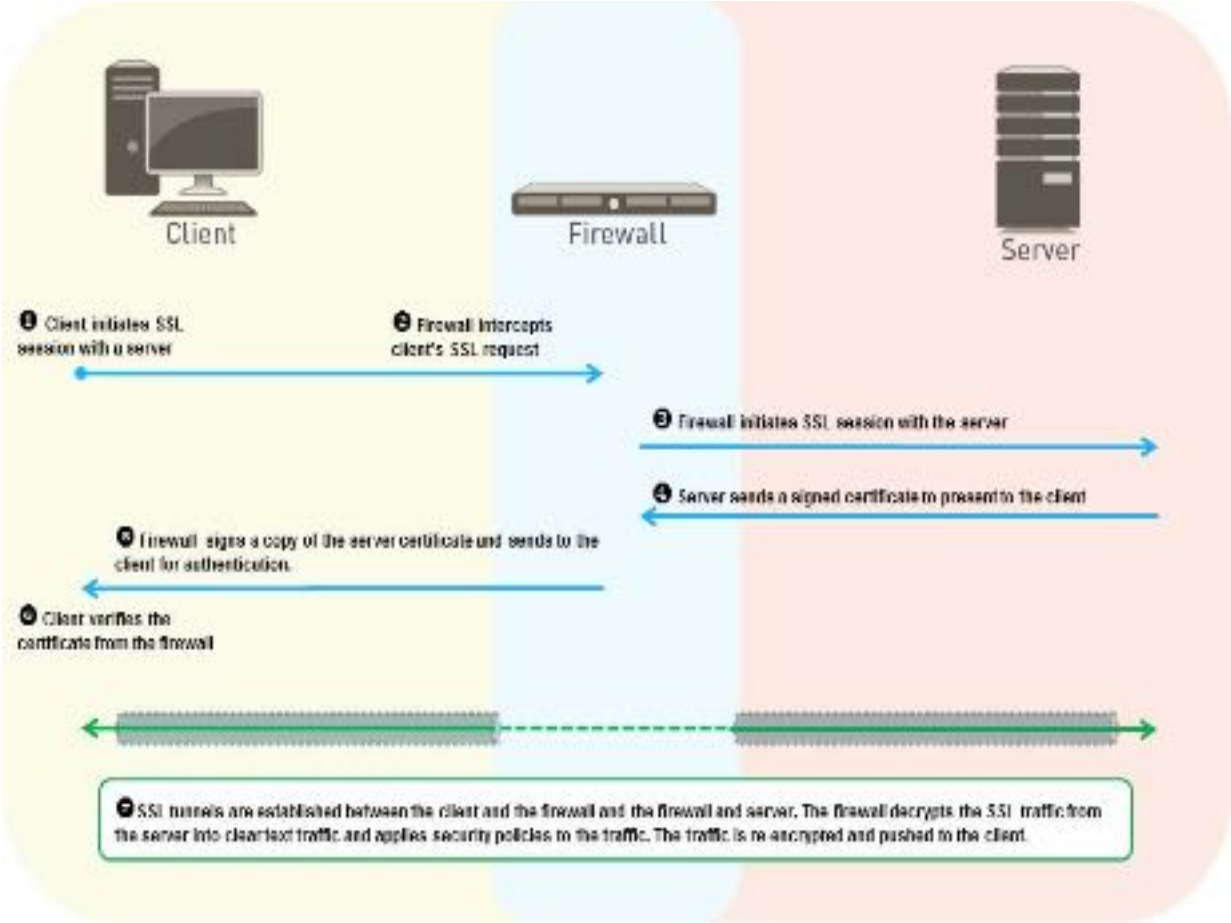
With SSL Forward Proxy decryption, the firewall resides between the internal client and outside server. The firewall uses certificates to establish itself as a trusted third party to the session between the client and the server (For details on certificates, see Keys and Certificates for Decryption Policies). 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 returns 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 signs it with the firewall Forward Trust certificate and sends the certificate to the client. If the server certificate is signed by a CA that the firewall does not trust, the firewall creates a copy of the server certificate, 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 decryption and security profiles 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.**

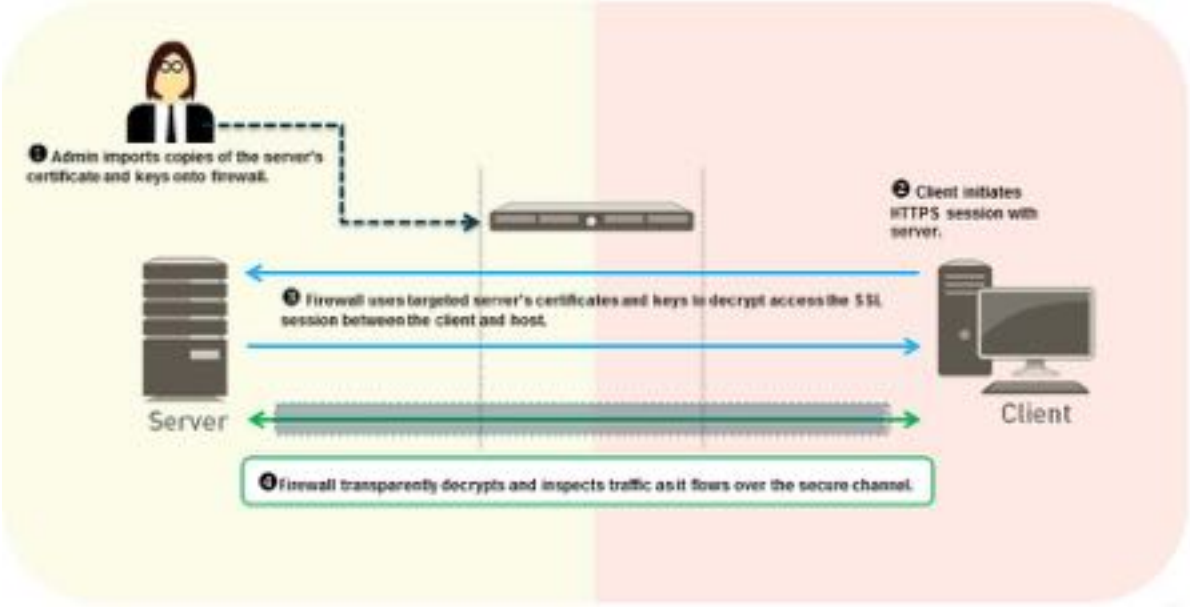
SSL Inbound Inspection

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.

SSH Proxy

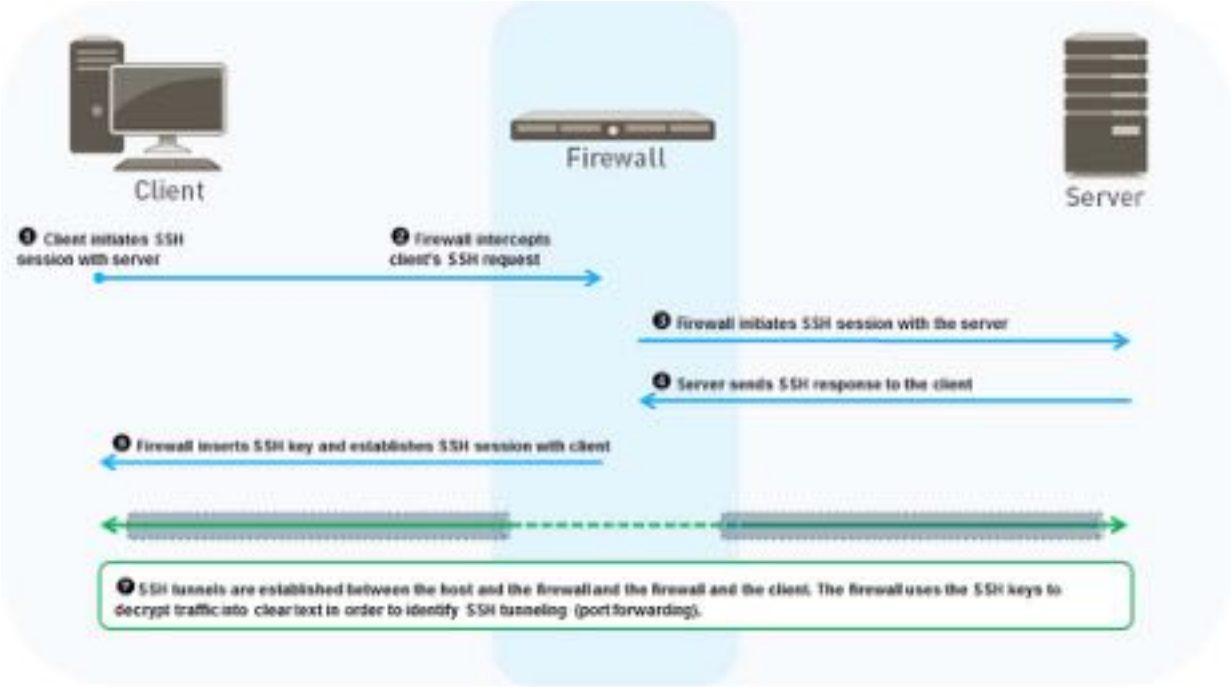
SSH Proxy provides the capability for the firewall to decrypt inbound and outbound SSH connections passing through the firewall, in order to ensure that SSH is not being used to tunnel unwanted applications and content. SSH decryption does not require any certificates and the key used for SSH decryption 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 firewall. 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.

 © Palo Alto Networks, Inc.

Figure: SSH Proxy Decryption



See Configure SSH Proxy for details on configuring an SSH Proxy policy.

Decryption Exceptions

Applications that do not function properly when the firewall decrypts them are automatically excluded from SSL decryption. For a current list of applications the firewall excludes from SSL decryption by default, see List of Applications Excluded from SSL Decryption.

You can also Configure Decryption Exceptions to exclude applications, URL categories, and targeted server traffic from decryption:

- Exclude 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 choose URL categories to exclude traffic that is categorized as financial or health-related from decryption.
- Exclude server traffic 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 for which you do not want to decrypt traffic, such as the web services for your HR systems, exclude those servers from decryption by importing the server certificate onto the firewall and modifying the certificate to be an **SSL Exclude Certificate**.

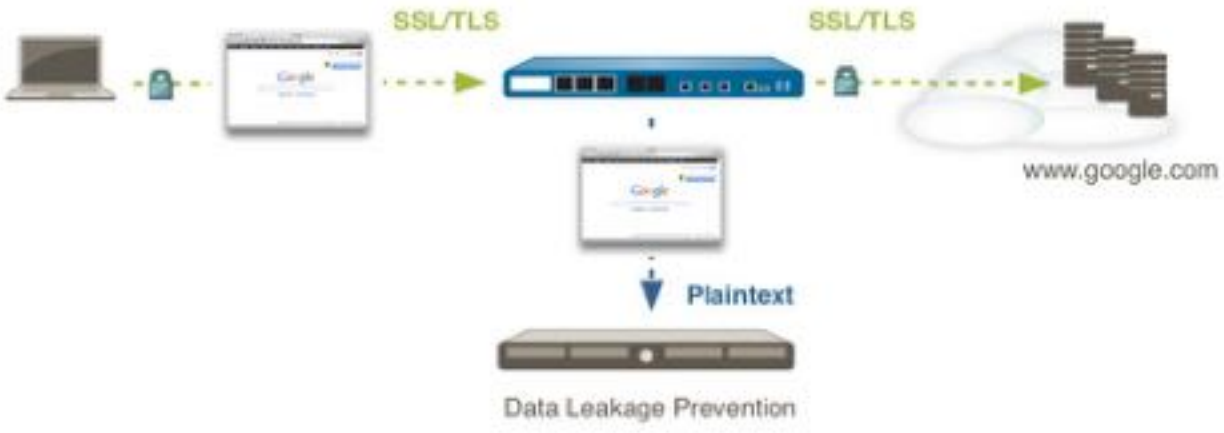
Decryption Mirroring

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 counsel 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



 © Palo Alto Networks, Inc.

Define Traffic to Decrypt

A decryption policy rule allows you to define traffic that you want the firewall to decrypt, or to define traffic that you want the firewall to exclude from decryption. You can attach a decryption profile rule to a decryption policy rule to more granularly control matching traffic.

▲ Create a Decryption Profile

▲ Create a Decryption Policy Rule

Create a Decryption Profile

A decryption profile allows you to perform checks on both decrypted traffic and traffic that you have excluded from decryption. Create a decryption profile to:

- Block sessions using unsupported protocols, cipher suits, or sessions that require client authentication.
- Block sessions based on certificate status, where the certificate is expired, is signed by an untrusted CA, has extensions restricting the certificate use, has an unknown certificate status, or the certificate status can't be retrieved during a configured timeout period.
- Block sessions if the resources to perform decryption are not available or if a hardware security module is not available to sign certificates.

After you create a decryption profile, you can attach it to a decryption policy rule; the firewall then enforces the decryption profile settings on traffic matched to the decryption policy rule.

Palo Alto Networks firewalls include a default decryption profile that you can use to enforce the basic recommended protocol versions and cipher suites for decrypted traffic.

Configure a Decryption Profile Rule	
Step 1	Select **Objects > Decryption Profile**, **Add** or modify a decryption profile rule, and give the rule a descriptive **Name**.
Step 2	(Optional) Allow the profile rule to be **Shared** across every virtual system on a firewall or every Panorama device group.
Step 3	(Decryption Mirroring Only) To Configure Decryption Port Mirroring, enable an Ethernet **Interface** for the firewall to use to copy and forward decrypted traffic. Decryption mirroring requires a decryption port mirror license.

Configure a Decryption Profile Rule (Continued)		
Step 4	(Optional) Block and control SSL tunneled and/or inbound traffic undergoing SSL Forward Proxy decryption or SSL Inbound Inspection.	Select **SSL Decryption**: • Select **SSL Forward Proxy** to configure settings to verify certificates, enforce protocol versions and cipher suites, and perform failure checks on SSL decrypted traffic. These settings are active only when this profile is attached to a decryption policy rule that is set to perform SSL Forward Proxy decryption. • Select **SSL Inbound Inspection** to configure settings enforce protocol versions and cipher suites and to perform failure checks on inbound SSL traffic. These settings are active only when this profile is attached to a decryption policy rule that is set to perform SSL Inbound Inspection. • Select **SSL Protocol Settings** to configure minimum and maximum protocol versions and key exchange, encryption, and authentication algorithms to enforce for SSL traffic. These settings are active when this profile is attached to decryption policy rules that are set to perform either SSL Forward Proxy decryption or SSL Inbound Inspection.
Step 5	(Optional) Block and control traffic (for example, a URL category) for which you have disabled decryption.	Select **No Decryption** and configure settings to validate certificates for traffic that is excluded from decryption. These setting are active only when the decryption profile is attached to a decryption policy rule that disables decryption for certain traffic.
Step 6	(Optional) Block and control SSH traffic undergoing SSH Proxy decryption.	Select **SSH Proxy** and configure settings to enforce supported protocol versions and These settings are active only when the decryption profile is attached to a decryption policy rule that decrypts SSH traffic.
Step 7	Add the decryption profile rule to a decryption policy rule. Traffic that the policy rules matches to is enforced based on the additional profile rule settings.	1. Select **Policies > Decryption** and Create a Decryption Policy Rule or modify an existing rule. 2. Select **Options** and select a **Decryption Profile** to block and control various aspects of the traffic matched to the rule. The profile rule settings that are applied to matching traffic depend on the policy rule Action (Decrypt or No Decrypt) and the policy rule Type (SSL Forward Proxy, SSL Inbound Inspection, or SSH Proxy). This allows you to use the default decryption profile, standard decryption profile customized for your organization, with different types of decryption policy rules. 3. Click **OK**.
Step 8	**Commit** the configuration.	

 © Palo Alto Networks, Inc.

Create a Decryption Policy Rule

Create a decryption policy rule to define traffic for the firewall to decrypt and the type of decryption you want the firewall to perform: SSL Forward Proxy, SSL Inbound Inspection, or SSH Proxy decryption. You can also use a decryption policy rule to define Decryption Exceptions.

Configure a Decryption Policy Rule		
Step 1	Select **Policies > Decryption** and **Add** a new decryption policy rule.	
Step 2	Give the policy rule a descriptive **Name**.	

Step 3	Configure the decryption rule to match to traffic based on network and policy objects:
	• **Firewall security zones**—Select **Source** and/or **Destination** and match to traffic based on the **Source Zone** and/or the **Destination Zone**.
	• **IP addresses, address objects, and/or address groups**—Select **Source** and/or **Destination** to match to traffic based on **Source Address** and/or the **Destination Address**. Alternatively, select **Negate** to exclude the source address list from decryption.
	• **Users**—Select **Source** and set the **Source User** for whom to decrypt traffic. You can decrypt specific user or group traffic, or decrypt traffic for certain types of users, such as unknown users or pre-logon users (users that are connected to GlobalProtect but are not yet logged in).
	• **Ports and protocols**—Select **Service/URL Category** to set the rule to match to traffic based on service. By default, the policy rule is set to decrypt **Any** traffic on TCP and UDP ports. You can **Add** a service or a service group, and optionally set the rule to **application-default** to match to applications only on the application default ports.
	  The application-default setting is useful to Configure Decryption Exceptions. You can exclude applications running on their default ports from decryption, while continuing to decrypt the same applications when they are detected on non-standard ports
	• **URLs and URL categories**—Select Service/URL Category and decrypt traffic based on:
	  • An externally-hosted list of URLs that the firewall retrieves for policy-enforcement (see **Objects > External Dynamic Lists**).
	  • Custom URL categories (see **Objects > Custom Objects > URL Category**).
	  • Palo Alto Networks URL categories. This option is useful to Configure Decryption Exceptions. For example, you could create a custom URL category to group sites that you do not want to decrypt, or you could exclude financial or healthcare-related sites from decryption based on the Palo Alto Networks URL categories.	

Step 4	Set the action the policy rule enforces on matching traffic: the rule can either decrypt matching traffic or exclude matching traffic from decryption.	Select **Options** and set the policy rule **Action**:
		Decrypt matching traffic:
		1. Select **Decrypt** .
		2. Set the **Type** of decryption for the firewall to perform on matching traffic:
		  • SSL Forward Proxy
		  • SSH Proxy
		  • SSL Inbound Inspection. If you want to enable SSL Inbound Inspection, also select the **Certificate** for the destination internal server for the inbound SSL traffic.
		Exclude matching traffic from decryption:
		Select **No Decrypt**.

Step 5	(Optional) Select a **Decryption Profile** to apply the profile settings to decrypted traffic. (To Create a Decryption Profile, select **Objects > Decryption Profile**).	

Configure a Decryption Policy Rule		
Step 6	Click **OK** to save the policy.	
Step 7	Choose your next step...	Fully enable the firewall to decrypt traffic: • Configure SSL Forward Proxy • Configure SSL Inbound Inspection • Configure SSH Proxy • Configure Decryption Exceptions

 © Palo Alto Networks, Inc.

Configure SSL Forward Proxy

To enable the firewall to perform SSL Forward Proxy decryption, you must set up the certificates required to establish the firewall as a trusted third party to the session between the client and the server. The firewall can use self-signed certificates or certificates signed by an enterprise certificate authority (CA) as *forward trust certificates* to authenticate the SSL session with the client.

- (Recommended) **Enterprise CA-signed Certificates**
  An enterprise CA can issue a signing certificate which the firewall can use to sign the certificates for sites requiring SSL decryption. When the firewall trusts the CA that signed the certificate of the destination server, the firewall can then send a copy of the destination server certificate to the client signed by the enterprise CA.
- **Self-signed Certificates**
  When a client connects to a server with a certificate that is signed by a CA that the firewall trusts, the firewall can sign a copy of the server certificate to present to the client and establish the SSL session. You can use self-signed certificates for SSL Forward Proxy decryption if your organization does not have an enterprise CA or if you intend to only perform decryption for a limited number of clients.

Additionally, set up a *forward untrust certificate* for the firewall to present to clients when the server certificate is signed by a CA that the firewall does not trust. This ensures that clients are prompted with a certificate warning when attempting to access sites with untrusted certificates.

After setting up the forward trust and forward untrust certificates required for SSL Forward Proxy decryption, add a decryption policy rule to define the traffic you want the firewall to decrypt. SSL tunneled traffic matched to the decryption policy rule is decrypted to clear text traffic. The clear text traffic is blocked and restricted based on the decryption profile attached to the policy and the firewall security policy. Traffic is re-encrypted as it exits the firewall.

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 for the firewall to present to clients when the server certificate is signed by a trusted CA: • (Recommended) Use an enterprise CA-signed certificate as the forward trust certificate. • Use a self-signed certificate as the forward trust certificate.	

Configure SSL Forward Proxy (Continued)	
• (Recommended) Use an enterprise CA-signed certificate as the forward trust certificate.	1. Generate a Certificate Signing Request (CSR) for the enterprise CA to sign and validate: a. Select **Device > Certificate Management > Certificates** and click **Generate**. b. Enter a **Certificate Name**, such as *my-fwd-proxy*. c. In the **Signed By** drop-down, select **External Authority (CSR)**. d. (Optional) If your enterprise CA requires it, add **Certificate Attributes** to further identify the firewall details, such as Country or Department. e. Click **OK** to save the CSR. The pending certificate is now displayed on the **Device Certificates** tab. 2. 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**. 3. Provide the certificate file to your enterprise CA. When you receive the enterprise CA-signed certificate from your enterprise CA, save the enterprise CA-signed certificate for import onto the firewall. 4. Import the enterprise CA-signed certificate 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. 5. 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. 6. Click **OK** to save the enterprise CA-signed forward trust certificate.

 © Palo Alto Networks, Inc.

Configure SSL Forward Proxy (Continued)	
• Use a self-signed certificate as the forward trust certificate.	1. Generate a new certificate: a. Select **Device > Certificate Management > Certificates**. b. Click **Generate** at the bottom of the window. c. Enter a **Certificate Name**, such as *my-fwd-trust*. d. 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. e. Leave the **Signed By** field blank. f. 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. g. **Generate** the certificate. 2. Click the new certificate *my-fwd-trust* to modify it and enable the certificate to be a **Forward Trust Certificate**. 3. Click **OK** to save the self-signed forward trust certificate.
Step 3 Distribute the forward trust certificate to client system certificate stores. – If you do not install the forward trust certificate on client systems, users will see certificate warnings for each SSL site they visit. – If you are using an enterprise-CA signed certificate as the forward trust certificate for SSL Forward Proxy decryption, and the client systems already have the enterprise CA added to the local trusted root CA list, you can skip this step.	On a firewall configured as a GlobalProtect portal: This option is supported with Windows and Mac client OS versions, and requires GlobalProtect agent 3.0.0 or later to be installed on the client systems. 1. Select **Network > GlobalProtect > Portals** and then select an existing portal configuration or **Add** a new one. 2. Select **Agent** and then select an existing agent configuration or **Add** a new one. 3. **Add** the SSL Forward Proxy forward trust certificate to the Trusted Root CA section. 4. **Install in Local Root Certificate Store** so that the GlobalProtect portal automatically distributes the certificate and installs it in the certificate store on GlobalProtect client systems. 5. Click **OK** twice. Without GlobalProtect: 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. 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).

Configure SSL Forward Proxy (Continued)		
Step 4	Configure the forward untrust certificate.	1. Click **Generate** at the bottom of the certificates page. 2. Enter a **Certificate Name**, such as *my-fwd-untrust*. 3. Set the **Common Name**, for example **192.168.2.1**. Leave **Signed By** blank. 4. Click the **Certificate Authority** check box to enable the firewall to issue the certificate. 5. Click **Generate** to generate the certificate. 6. Click **OK** to save. 7. Click the new *my-ssl-fw-untrust* certificate to modify it and enable the **Forward Untrust Certificate** option. 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. 8. Click **OK** to save.
Step 5	(Optional) Set the key size of the SSL Forward Proxy certificates that the firewall presents to clients. By default, the firewall determines the key size to use based on the key size of the destination server certificate.	Configure the Key Size for SSL Forward Proxy Server Certificates.
Step 6	Create a Decryption Policy Rule to define traffic for the firewall to decrypt.	1. Select **Policies > Decryption**, Add or modify an existing rule, and define traffic to be decrypted. 2. Select **Options** and: • Set the rule **Action** to **Decrypt** matching traffic. • Set the rule **Type** to **SSL Forward Proxy**. • (Optional) Select a **Decryption Profile** to block and control various aspects of the decrypted traffic (for example, Create a Decryption Profile to perform certificate checks and enforce strong cipher suites and protocol versions). 3. Click **OK** to save.
Step 7	Enable the firewall to forward decrypted SSL traffic for WildFire analysis. This option requires an active WildFire license and is a WildFire best practice.	
Step 8	**Commit** the configuration.	
Step 9	Choose your next step...	• Enable Users to Opt Out of SSL Decryption. • Configure Decryption Exceptions to disable decryption for certain types of traffic.

 © Palo Alto Networks, Inc.

Configure SSL Inbound Inspection

Use SSL Inbound Inspection to decrypt and inspect inbound SSL traffic destined for a network server (you can perform SSL Inbound Inspection for any server if you have the server certificate). With an SSL Inbound Inspection decryption policy enabled, all SSL traffic identified by the policy is decrypted to clear text traffic and inspected. The clear text traffic is blocked and restricted based on the decryption profile attached to the policy and any configured Antivirus, Vulnerability, Anti-Spyware, URL-Filtering and File Blocking profiles. You can also enable the firewall to o forward decrypted SSL traffic for WildFire analysis and signature generation. Traffic is re-encrypted as it exits the firewall.

Configuring SSL Inbound Inspection includes installing the targeted server certificate on the firewall and creating an SSL Inbound Inspection decryption policy.

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. To import the targeted server certificate onto the firewall: 1. On the **Device Certificates** tab, select **Import**. 2. Enter a descriptive **Certificate Name**. 3. Browse for and select the targeted server **Certificate File**. 4. Click **OK**.
Step 3	Create a Decryption Policy Rule to define traffic for the firewall to decrypt.	1. Select **Policies > Decryption**, Add or modify an existing rule, and define traffic to be decrypted. 2. Select **Options** and: • Set the rule **Action** to **Decrypt** matching traffic. • Set the rule **Type** to **SSL Inbound Inspection**. • Select the **Certificate** for the internal server that is the destination of the inbound SSL traffic. • (Optional) Select a **Decryption Profile** to block and control various aspects of the decrypted traffic (for example, Create a Decryption Profile to terminate sessions if system resources are not available to process decryption). 3. Click **OK** to save.
Step 4	Enable the firewall to forward decrypted SSL traffic for WildFire analysis. This option requires an active WildFire license and is a WildFire best practice.	
Step 5	**Commit** the configuration.	

Configure SSL Inbound Inspection	
Step 6 Choose your next step...	• Enable Users to Opt Out of SSL Decryption. • Configure Decryption Exceptions to disable decryption for certain types of traffic.

 © 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.

With SSH decryption 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 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	Create a Decryption Policy Rule to define traffic for the firewall to decrypt.	1. Select **Policies > Decryption**, Add or modify an existing rule, and define traffic to be decrypted. 2. Select **Options** and: • Set the rule **Action** to **Decrypt** matching traffic. • Set the rule **Type** to **SSH Proxy**. • (Optional) Select a **Decryption Profile** to block and control various aspects of the decrypted traffic (for example, Create a Decryption Profile to terminate sessions if system resources are not available to process decryption). 3. Click **OK** to save.
Step 3	**Commit** the configuration.	
Step 4	(Optional) Continue to Configure Decryption Exceptions to disable decryption for certain types of traffic.	

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 traffic from decryption, create a decryption policy rule and set the policy action to **No Decrypt**. Exclude traffic from decryption based on application, 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	Exclude traffic from decryption based match criteria. This example shows how to exclude traffic categorized as financial or health-related from SSL Forward Proxy decryption.	1. Select **Policies > Decryption** and **modify or Create a Decryption Policy rule.** 2. Define the traffic that you want to exclude from decryption. In this example: a. Give the rule a descriptive **Name**, such as *No-Decrypt-Finance-Health*. b. Set the **Source** and **Destination** to **Any** to apply the *No-Decrypt-Finance-Health* rule to all SSL traffic destined for an external server. c. Select **URL Category** and **Add** the URL categories financial-services and health-and-medicine. 3. Select **Options** and set the rule to **No Decrypt**. 4. (Optional) You can still use a decryption profile to validate certificates for sessions the firewall does not decrypt. Attach a decryption profile to the rule that is set to **Block sessions with expired certificates** and/or **Block sessions with untrusted issuers**. 5. Click **OK** to save the *No-Decrypt-Finance-Health* decryption rule.
Step 2	Place the decryption exclusion rule at the top of your decryption policy. 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 the traffic matched to the rule remains encrypted, even if the traffic is later matched to other decryption rules.	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).

 © Palo Alto Networks, Inc.

Exclude Traffic from a Decryption Policy	
Step 3	**Commit** the configuration.

Exclude a Server from Decryption

You can exclude 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: 1. On the **Device > Certificate Management > Certificates > Device Certificates** tab, select **Import**. 2. Enter a descriptive **Certificate Name**. 3. Browse for and select the targeted server **Certificate File**. 4. Click **OK**.
Step 2	Select the targeted server certificate on the **Device Certificates** tab and enable it to be an **SSL Exclude Certificate**. When the targeted server certificate is designated as an SSL Exclude Certificate, the firewall does not decrypt the server traffic even if the traffic matches decryption policy rule.

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.

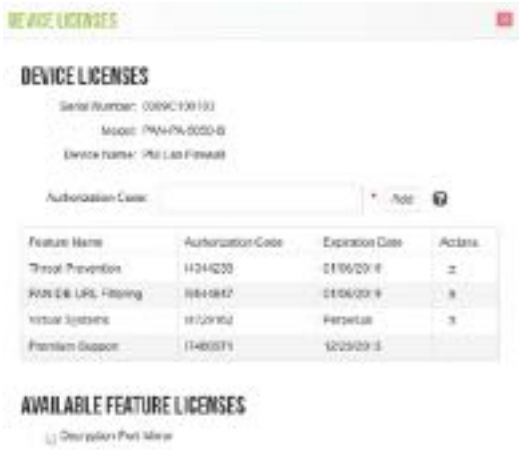
Enable Users to Opt Out of SSL Decryption	
Step 1 (Optional) Customize the SSL Decryption Opt-out Page.	1. Select **Device > Response Pages**. 2. Select the **SSL Decryption Opt-out Page** link. 3. Select the **Predefined** page and click **Export**. 4. Using the HTML text editor of your choice, edit the page. 5. If you want to add an image, host the image on a web server that is accessible from your end user systems. 6. Add a line to the HTML to point to the image. For example: `<img src="http://cdn.slidesharecdn.com/Acme-logo-96x96.jpg?1382722588"/>` 7. Save the edited page with a new filename. Make sure that the page retains its UTF-8 encoding. 8. Back on the firewall, select **Device > Response Pages**. 9. Select the **SSL Decryption Opt-out Page** link. 10. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file. 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. 12. Click **OK** to import the file. 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. 2. Select the **Enable SSL Opt-out Page** and click **OK**. 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. Verify that the SSL Decryption Opt-out response page displays. 

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 firewall on which you want to enable decryption port mirroring.	1. Log in to the Palo Alto Networks Customer Support web site and navigate to the **Assets** tab. 2. Select the entry for the firewall you want to license and select **Actions**. 3. Select **Decryption Port Mirror**. A legal notice displays. 4. If you are clear about the potential legal implications and requirements, click **I understand and wish to proceed**. 5. Click **Activate**. 
Step 2 Install the Decryption Port Mirror license on the firewall.	1. From the firewall web interface, select **Device > Licenses**. 2. Click **Retrieve license keys from license server**. 3. Verify that the license has been activated on the firewall.  4. Reboot the firewall (**Device > Setup > Operations**). This feature is not 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:** 1. Select **Device > Setup > Content - ID**. 2. Select the **Allow forwarding of decrypted content** check box. 3. Click **OK** to save. **On a firewall with multiple virtual systems:** 1. Select **Device > Virtual System.** 2. Select a Virtual System to edit or create a new Virtual System by selecting **Add**. 3. Select the **Allow forwarding of decrypted content** check box. 4. Click **OK** to save.
Step 4	Enable an Ethernet interface to be used for decryption mirroring.	1. Select **Network > Interfaces > Ethernet**. 2. Select the Ethernet interface that you want to configure for decryption port mirroring. 3. Select **Decrypt Mirror** as the **Interface Type**. This interface type will appear only if the Decryption Port Mirror license is installed. 4. Click **OK** to save.
Step 5	Enable mirroring of decrypted traffic.	1. Select **Objects > Decryption Profile**. 2. Select an **Interface** to be used for **Decryption Mirroring**. The **Interface** drop-down contains all Ethernet interfaces that have been defined as the type: **Decrypt Mirror**. 3. Specify whether to mirror decrypted traffic before or after policy enforcement. 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). 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**. 2. Click **Add** to configure a decryption policy or select an existing decryption policy to edit. 3. In the **Options** tab, select **Decrypt** and the **Decryption Profile** created in Step 4. 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
- Use an External Dynamic List in a URL Filtering Profile
- 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

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

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 (C2) 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. that 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 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.

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.

 © Palo Alto Networks, Inc.

Differences	PAN-DB Public Cloud	PAN-DB Private Cloud
URL Categorization Requests	Submit URL categorization change requests using the following options: • Palo Alto Networks Test A Site website. • URL filtering profile setup page on the firewall. • 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. 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.

URL Filtering Concepts

- URL Categories
- URL Filtering Profile
- URL Filtering Profile Actions
- Block and Allow Lists
- External Dynamic List for URLs
- Safe Search Enforcement
- Container Pages
- HTTP Header Logging
- URL Filtering Response Pages
- URL Category as Policy Match Criteria

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:

 © Palo Alto Networks, Inc.

Category	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 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. 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. 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 collects the list of URLs from the unknown category and processes them to determine the URL category. These URLs are processed automatically, everyday, provided the websites has machine readable content that is in a supported format and language. Upon categorization, the updated category information is made available to all PAN-DB customers. See Configure URL Filtering.



You can submit URL categorization change requests using the Palo Alto Networks dedicated web portal (Test A Site), the URL filtering profile setup page on the firewall, or the URL filtering log on the firewall. Each change request is automatically processed everyday, provided the websites provides machine readable content that is in a supported format and language. Sometimes, the categorization change requires a member of the Palo Alto Networks engineering staff to perform a manual review. In such cases, the process may take a little longer.

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 for 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
- External Dynamic List for URLs
- Safe Search Enforcement
- Container Pages
- HTTP Header Logging

URL Filtering Profile Actions

The URL Filtering profile specifies an action for each 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.  The Continue page doesn't display properly on client systems configured to use a proxy server.

 © Palo Alto Networks, Inc.

Action	Description
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. The Override page doesn't display properly on client systems 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 profile, if you do not want the block action to apply to the other profile, you must set the action to **none**. 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

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 or External Dynamic List for URLs, 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 the allow or block list entries. 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, to prevent a user from accessing any website within the paloaltonetworks.com domain, add `*.paloaltonetworks.com` to the block list. This will block all paloaltoneworks.com URLs, even if the address includes a domain prefix (`http://`, `www`) or a sub-domain prefix (`mail.paloaltonetworks.com`) . The same applies to the sub-domain suffix. For example, if you want to block paloaltonetworks.com/en/US, you would add `paloaltonetworks.com/*` to the block list as well.

  Further, to block access to a domain suffix such as paloaltonetworks.com.au, you must add an entry with a slash (`/`) at the end. In this example, you would add `*.paloaltonetworks.com/` to the block list.
- The block and allow lists support wildcard patterns. The following characters are considered separators:

  `.`

  `/`

  `?`

  `&`

  `=`

  `;`

  `+`

  Every substring separated by a character listed above is considered a token. A token can be any number of ASCII characters that does not contain any separator character or an asterisks (*). 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 asterisks (*) is not the only character in the token:

- ww*.yahoo.com
- www.y*.com

External Dynamic List for URLs

To protect your network from new sources of threat or malware, you can use External Dynamic List in URL Filtering profiles to block or allow, or to define granular actions such as continue, alert, or override for URLs, before you attach the profile to a Security policy rule. Unlike the allow list, block list, or a custom URL category on the firewall, an external dynamic list gives you the ability to update the list without a configuration change or commit on the firewall. The firewall dynamically imports the list at the configured interval and enforces policy for the URLs (IP addresses or domains will be ignored) in the list. For URL formatting guidelines, see Block and Allow Lists.

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.

 © Palo Alto Networks, Inc.

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: **Safe Search Enforcement for Google Searches on Individual Computers** 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. Appending `safe=active` to a Google search query URL also enables the strictest safe search settings. **Safe Search Enforcement for Google and YouTube Searches using a Virtual IP Address** 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. 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. Appending `vm=r` to a Yahoo search query URL also enables the strictest safe search settings. 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. Appending `adlt=strict` to a Bing search query URL also enables the strictest safe search settings. 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.

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:

 © Palo Alto Networks, Inc.

- **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.

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.

Variable	Usage
<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

 © Palo Alto Networks, Inc.

that matches those categories and set the action to no-decrypt. By placing this rule above the rule to decrypt all traffic, 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. 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: • 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. • 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. 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

▲ 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.

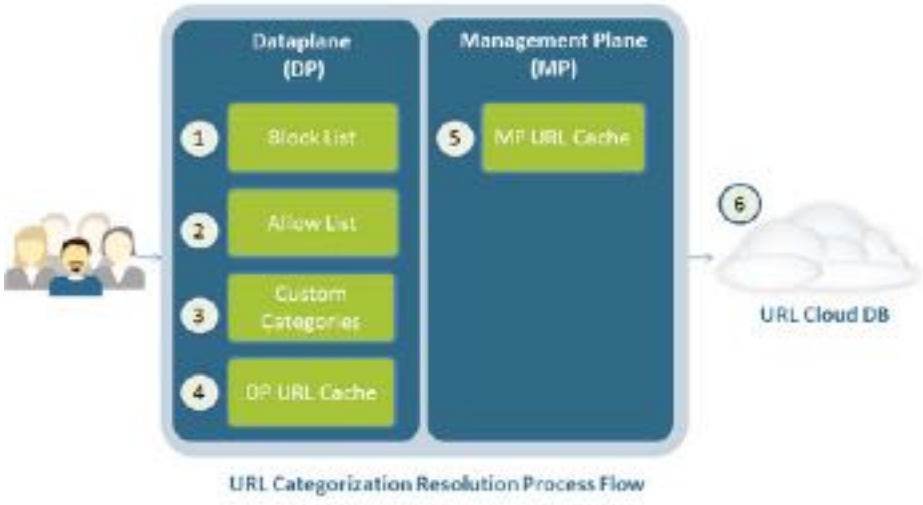
Component	Description
URL Filtering Seed Database	The initial seed database downloaded to the firewall is a small subset of the database that is maintained on the Palo Alto Networks URL cloud servers. The reason this is done is because the full database contains millions of URLs and many of these URLs 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. 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 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 local URL database on the firewall 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. 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 firewall (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 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 local drive on the firewall 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; this expiration period is not configurable.

PAN-DB URL Categorization Workflow

When a user requests a URL the firewall determines the URL category by comparing the URL with the following components (in order) until it finds a match:



URL Categorization Resolution Process Flow

If a requested URL matches an expired entry in the dataplane (DP) URL cache, the cache responds with the expired category, but also sends a URL categorization query to the management plane (MP) cache. This prevents unnecessary delays in the DP, assuming that the frequency of category change is low. Similarly, in the MP URL cache, if a URL query from the DP cache matches an expired entry in the MP cache, the MP responds to the DP with the expired category and will also send a URL categorization request to the PAN-DB cloud database. Upon getting the response from the cloud, the firewall sends the updated category to the DP.

As new URLs and categories are defined or if critical updates are needed, the cloud database is 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

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.  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: • **Retrieve license keys from license server** • **Activate feature using authorization code** • **Manually upload license key** 2. After installing the license, confirm that the PAN-DB URL Filtering section, **Date Expires** field, displays a valid date.

Enable PAN-DB URL Filtering (Continued)	
Step 2 Download the initial seed database and activate PAN-DB URL Filtering. 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**. 2. Choose a region (North America, Europe, APAC, Japan) and then click **OK** to start the download. 3. After the download completes, click **Activate**. 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. A Threat Prevention license is required to receive content updates, which covers Antivirus and Applications and Threats.	1. Select **Device > Dynamic Updates**. 2. In the Schedule field in the Applications and Threats section, click the **None** link to schedule periodic updates. 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. The Applications and Threats updates sometimes 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. 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: • **Activate feature using authorization code** • **Retrieve license keys from license server** • **Manually upload license key** 2. After installing the license, confirm that the BrightCloud URL Filtering section, **Date Expires** field, displays a valid date.

 © Palo Alto Networks, Inc.

Enable BrightCloud URL Filtering (Continued)		
Step 2	Install the BrightCloud database. The way you do this depends on whether or not the firewall has direct Internet access.	**Firewall with Direct Internet Access** 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. **Firewall without Direct Internet Access** 1. Download the BrightCloud database to a host that has Internet access. The firewall must have access to the host: a. On a host with Internet access, go to the Palo Alto Networks Customer Support web site, www.paloaltonetworks.com/support/tabs/overview.html, and log in. b. In the Resources section, click **Dynamic Updates**. c. In the BrightCloud Database section, click **Download** and save the file to the host. 2. Upload the database to the firewall: a. Log in to the firewall, select **Device > Dynamic Updates** and click **Upload**. b. For the **Type**, select **URL Filtering**. 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**. 3. Install the database: a. Select **Device > Dynamic Updates** and click **Install From File**. b. For the **Type**, select **URL Filtering**. The firewall automatically selects the file you just uploaded. 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.
		2. Enter the following commands to enable dynamic URL filtering:
		`configure`
		`set deviceconfig setting url dynamic-url yes`
		`commit`

Enable BrightCloud URL Filtering (Continued)	
Step 4 Schedule the firewall to download dynamic updates for Applications and Threats signatures and URL filtering. You can only schedule dynamic updates if the firewall has direct Internet access. 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. 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. A Threat Prevention license is required to receive Antivirus and Applications and Threats updates.	1. Select **Device > Dynamic Updates**. 2. In the Applications and Threats section, Schedule field, click the **None** link to schedule periodic updates. 3. In the URL Filtering section, Schedule field, click the **None** link to schedule periodic updates. If updates are already scheduled in a section, the link text displays the schedule settings.

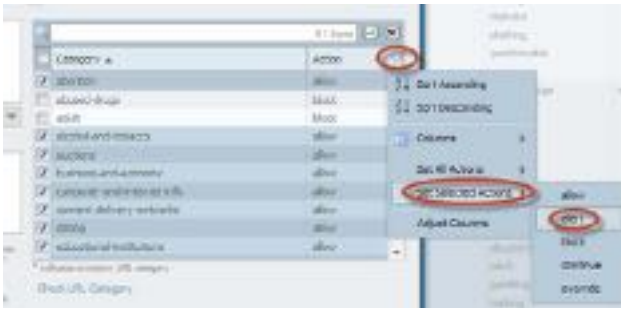
 © Palo Alto Networks, Inc.

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.

If you subscribe to third-party URL feeds and want to secure your users from emerging threats, see Use an External Dynamic List in a URL Filtering Profile.

Configure and Apply a Passive URL Filtering Profile		
Step 1	Create a new URL Filtering profile.	1. Select **Objects > Security Profiles > URL Filtering**. 2. Select the default profile and then click **Clone**. The new profile will be named **default-1**. 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. 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. 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**.  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. 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. 2. Select the **Actions** tab and in the **Profile Setting** section, click the drop-down for **URL Filtering** and select the new profile. 3. Click **OK** to save.

Configure and Apply a Passive URL Filtering Profile (Continued)		
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. 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.

Use an External Dynamic List in a URL Filtering Profile

An *external dynamic list* is a text file that is hosted on an external web server. You can use this list to import URLs and enforce policy on these URLs. When you update the list on the web server, the firewall retrieves the changes and applies policy to the modified list without requiring a commit on the firewall.

For more information, see External Dynamic List and Enforce Policy on Entries in an External Dynamic List.

Use an External Dynamic List with URLs in a URL Filtering Profile		

Step 1	Create the external dynamic list for URLs and host it on a web server.	Create a text file and enter the URLs in the file; each URL must be on a separate line. For example:
		`financialtimes.co.in`
		`www.wallaby.au/joey`
		`www.exyang.com/auto-tutorials/How-to-enter-Data-for-Success.aspx`
		`*.example.com/*`
		`abc?*/abc.com`
		`*&*.net`
		See Block and Allow Lists for formatting guidelines.

Step 2	Configure the firewall to access the external dynamic list.	1. Select **Objects > External Dynamic Lists**. 2. Click **Add** and enter a descriptive **Name** for the list. 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. 4. In the Type drop-down, select **URL List**. Ensure that the list does not include IP addresses or domain names; the firewall skips non-URL entries. 5. Enter the **Source** for the list you just created on the web server. The source must include the full path to access the list. For example, https://1.2.3.4/EDL_IP_2016. 6. Click **Test Source URL** to verify that the firewall can connect to the web server.  If the web server is unreachable after the connection is established, the firewall uses the last successfully retrieved list for enforcing policy until the connection is restored with the web server. 7. (Optional) Specify the **Repeat** frequency at which the firewall retrieves the list. By default, the firewall retrieves the list once every hour. 8. Click **OK**.

Use an External Dynamic List with URLs in a URL Filtering Profile (Continued)		
Step 3	Use the external dynamic list in a URL Filtering profile.	1. Select **Objects > Security Profiles > URL Filtering**. 2. **Add** or modify an existing URL Filtering profile. 3. **Name** the profile and, in the **Categories** tab, select the external dynamic list from the Category list. 4. Click Action to select a more granular action for the URLs in the external dynamic list. If a URL that is included in an external dynamic list is also included in a custom URL category, or Block and Allow Lists, the action specified in the custom category or the block and allow list will take precedence over the external dynamic list. 5. Click **OK**. 6. Attach the URL Filtering profile to a Security policy rule. a. Select **Policies > Security**. b. Select the **Actions** tab and, in the Profile Setting section, select the new profile in the **URL Filtering** drop-down. c. Click **OK** and **Commit**.
Step 4	Test that the policy action is enforced.	1. Attempt to access a URL that is included in the external dynamic list. 2. Verify that the action you defined is enforced in the browser. 3. To monitor the activity on the firewall: a. Select **ACC** and add a URL Domain as a global filter to view the Network Activity and Blocked Activity for the URL you accessed. b. Select **Monitor > Logs > URL Filtering** to access the detailed log view.

Step 5	Verify whether entries in the external dynamic list were ignored or skipped.
	In a list of type URL, the firewall skips non-URL entries as invalid and ignores entries that exceed the maximum limit for the platform.
		Use the following CLI command on a firewall to review the details for a list.
		`request system external-list show type url <list_name>`
		For example:
		`request system external-list show type url EBL_ISAC_Alert_List`

 © Palo Alto Networks, Inc.

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.

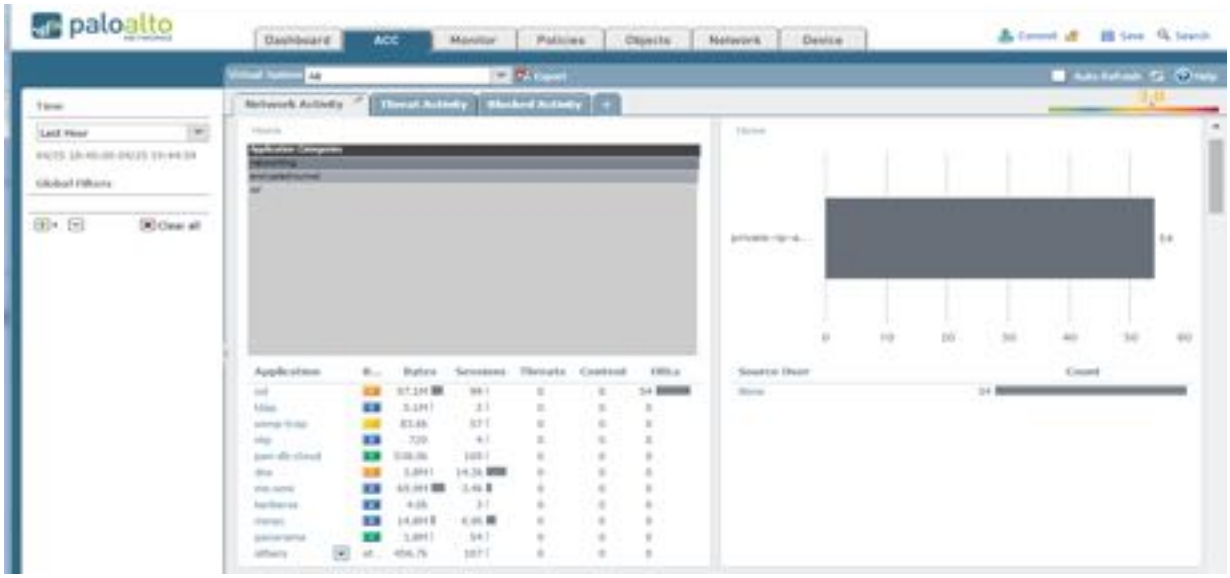
The following topics describe how to monitor web activity:

- Monitor Web Activity of Network Users
- View the User Activity Report
- Configure Custom URL Filtering Reports

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.



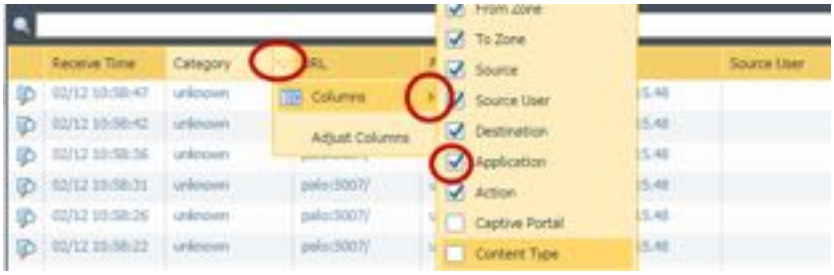
- **Block log**—In this log, the category malware was set to block, so the action is block-url and the user will see a response page indicating that the website was blocked.



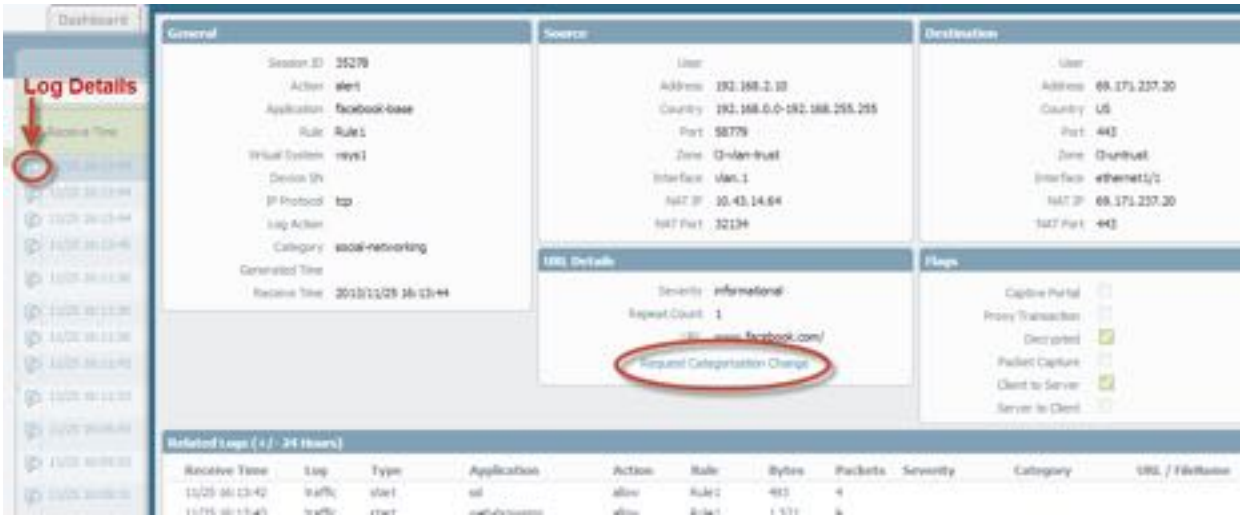
- **Alert log on encrypted website**—In this example, the category is social-networking and the application is facebook-base, which is required to access the Facebook website and other Facebook applications. Because faceboook.com is always encrypted using SSL, the traffic was decrypted by the firewall, 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.



View the User Activity Report

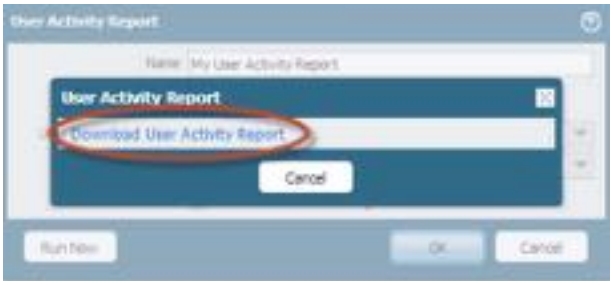
This report provides a quick method of viewing user or group activity and also provides an option to view browse time activity.

Generate a User Activity Report	
Step 1 Configure a User Activity Report.	1. Select **Monitor > PDF Reports > User Activity Report**. 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. 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. 3. Enter the Username/IP address for a user report or enter the group name for a user group report. 4. Select the time period. You can select an existing time period, or select **Custom**. 5. Select the **Include Detailed Browsing** check box, so browsing information is included in the report. 

Generate a User Activity Report (Continued)

Step 2 Run the user activity report and then download the report.

1. Click **Run Now**.
2. After the report is generated, click the **Download User Activity Report** link.



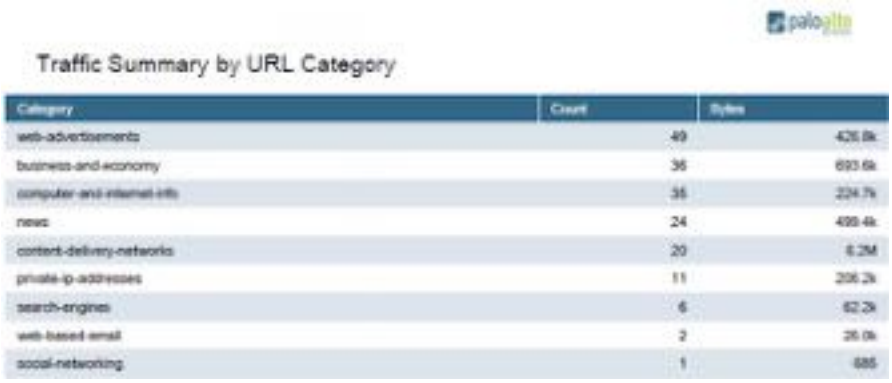
3. After the report is downloaded, click **Cancel** and then click **OK** to save the report.

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.

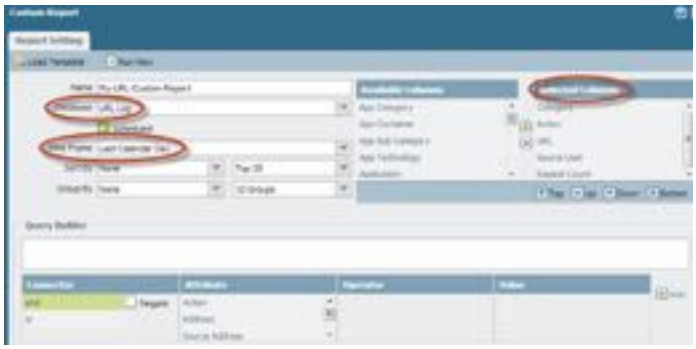

Traffic Summary by URL Category

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

 © Palo Alto Networks, Inc.

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**. 2. Enter a report **Name**, for example, My-URL-Custom-Report. 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. 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. 3. In the **Available Columns** list, select the fields to include the report. The following columns are typically used for a URL report: • Action • Category • Destination Country • Source User • URL 
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. 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**.

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.  Because the default URL filtering profile blocks risky and threat-prone content, it is a best practice to clone this profile to preserve these default settings, rather than creating a new profile.	1. Select **Objects > Security Profiles > URL Filtering**. Select the default profile and then click **Clone**. The new profile will be named default-1. 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: • If you do not care about traffic to a particular category (that is you neither want to block it nor log it), select **allow**. • For visibility into traffic to sites in a category, select **alert**. • To deny access to traffic that matches the category and to enable logging of the blocked traffic, select **block**. • To require users to click Continue to proceed to a questionable site, select **continue**. • 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. 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. 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. 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: • **block**—Block the URL. • **continue**—Prompt users click **Continue** to proceed to the web page. • **override**—The user will be a prompted for a password to continue to the website. • **alert**—Allow the user to access the website and add an alert log entry in the URL log. 2. For the **Allow** list, enter IP addresses or URLs that should always be allowed. Each row must be separated by a new line. 3. (Optional) Enable Safe Search Enforcement.

 © Palo Alto Networks, Inc.

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: • **User-Agent** • **Referer** • **X-Forwarded-For**
Step 6	Save the URL filtering profile.	1. Click **OK**. 2. (Optional) Customize the URL Filtering Response Pages. 3. Click **Commit**.  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.

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**. 2. Select the link for the URL filtering response page you want to modify. 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: • 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. • 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. 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**. 2. Select the link that corresponds to the URL Filtering response page you edited. 3. Click **Import** and then enter the path and filename in the **Import File** field or **Browse** to locate the file. 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. 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.

 © Palo Alto Networks, Inc.

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**. 2. In the **URL Admin Override** section, click **Add**. 3. In the **Location** field, select the virtual system to which this password applies. 4. Enter the **Password** and **Confirm Password**. 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. 6. Select the **Mode** for prompting the user for the password: • **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. • **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. 7. Click **OK**.
Step 2 (Optional) Set a custom override period.	1. Edit the URL Filtering section. 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. 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. 4. Click **OK**.

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: a. Select **Network > Interface Mgmt** and click **Add**. b. Enter a **Name** for the profile, select **Response Pages**, and then click **OK**. 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: 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. 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. 3. **Generate** the certificate. 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. 2. On the **Categories** tab, set the Action to **override** for each category that requires a password. 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. 2. Select the **Actions** tab and in the **Profile Setting** section, click the drop-down for **URL Filtering** and select the profile. 3. Click **OK** to save.
Step 7	Save the configuration.	Click **Commit**.

 © Palo Alto Networks, Inc.

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.

Enable Safe Search Enforcement	

Step 1 Enable Safe Search Enforcement in the URL Filtering profile.	1. Select **Objects > Security Profiles > URL Filtering**.
	2. Select an existing profile to modify, or clone the default profile to create a new profile.
	3. On the **Settings** tab, select the **Safe Search Enforcement** check box to enable it.
	4. (Optional) Restrict users to specific search engines:
	a. On the **Categories** tab, set the **search-engines** category to **block**.
	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:
	`www.google.com`
	`www.bing.com`
	5. Configure other settings as necessary to:
	• Define how to control access to web content.
	• Define websites that should always be blocked or allowed.
	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. 2. On the **Actions** tab, select the **URL Filtering** profile. 3. Click **OK** to save the security policy rule.

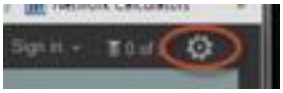
Step 3 Enable SSL Forward Proxy decryption.
Because most search engines encrypt their search results, you must enable SSL forward proxy decryption so that the firewall can inspect the search traffic and detect the safe search settings.
	1. Add a custom URL category for the search sites:
	a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.
	b. Enter a **Name** for the category, such as SearchEngineDecryption.
	c. **Add** the following to the Sites list:
	`www.bing.*`
	`www.google.*`
	`search.yahoo.*`
	d. Click **OK** to save the custom URL category object.
	2. Follow the steps to Configure SSL Forward Proxy.
	3. On the **Service/URL Category** tab in the Decryption policy rule, **Add** the custom URL category you just created and then click **OK**.

 © Palo Alto Networks, Inc.

Enable Safe Search Enforcement (Continued)		

Step 4	(Optional, but recommended) Block Bing search traffic running over SSL.
	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:
		a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.
		b. Enter a **Name** for the category, such as EnableBingSafeSearch.
		c. **Add** the following to the Sites list:
		`www.bing.com/images/*`
		`www.bing.com/videos/*`
		d. Click **OK** to save the custom URL category object.
		2. Create another URL filtering profile to block the custom category you just created:
		a. Select **Objects > Security Profiles > URL Filtering**.
		b. **Add** a new profile and give it a descriptive **Name**.
		c. Locate the custom category in the Category list and set it to **block**.
		d. Click **OK** to save the URL filtering profile.
		3. Add a security policy rule to block Bing SSL traffic:
		a. Select **Policies > Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.
		b. On the **Actions** tab, attach the URL filtering profile you just created to block the custom Bing category.
		c. On the **Service/URL Category** tab **Add** a **New Service** and give it a descriptive **Name**, such as bingssl.
		d. Select **TCP** as the **Protocol** and set the **Destination Port** to `443`.
		e. Click **OK** to save the rule.
		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**.

Enable Safe Search Enforcement (Continued)	
Step 6 Verify the Safe Search Enforcement configuration. 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.  2. Set the **SafeSearch** option to **Moderate** or **Off** and click **Save**. 3. Perform a Bing search and verify that the URL Filtering Safe Search Block page displays instead of the search results:  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**. 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**. 2. Check the **Applications and Threats** section to determine what update is currently running. 3. If the firewall is not running the required update or later, click **Check Now** to retrieve a list of available updates. 4. Locate the required update and click **Download**. 5. After the download completes, click **Install**.

 © Palo Alto Networks, Inc.

Enable Transparent Safe Search Enforcement (Continued)		

Step 2	Enable Safe Search Enforcement in the URL Filtering profile.	1. Select **Objects > Security Profiles > URL Filtering**.
		2. Select an existing profile to modify, or clone the default profile to create a new one.
		3. On the **Settings** tab, select the **Safe Search Enforcement** check box to enable it.
		4. (Optional) Allow access to specific search engines only:
		a. On the **Categories** tab, set the **search-engines** category to **block**.
		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:
		`www.google.com`
		`www.bing.com`
		5. Configure other settings as necessary to:
		• Define how to control access to web content.
		• Define websites that should always be blocked or allowed.
		6. Click **OK** to save the profile.

Step 3	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. 2. On the **Actions** tab, select the **URL Filtering** profile. 3. Click **OK** to save the security policy rule.

Enable Transparent Safe Search Enforcement (Continued)		

Step 4	(Optional, but recommended) Block Bing search traffic running over SSL.
	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:
		a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.
		b. Enter a **Name** for the category, such as EnableBingSafeSearch.
		c. **Add** the following to the Sites list:
		`www.bing.com/images/*`
		`www.bing.com/videos/*`
		d. Click **OK** to save the custom URL category object.
		2. Create another URL filtering profile to block the custom category you just created:
		a. Select **Objects > Security Profiles > URL Filtering**.
		b. **Add** a new profile and give it a descriptive **Name**.
		c. Locate the custom category you just created in the Category list and set it to **block**.
		d. Click **OK** to save the URL filtering profile.
		3. **Add** a security policy rule to block Bing SSL traffic:
		a. Select **Policies > Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.
		b. On the **Actions** tab, attach the URL filtering profile you just created to block the custom Bing category.
		c. On the **Service/URL Category** tab **Add** a **New Service** and give it a descriptive **Name**, such as bingssl.
		d. Select **TCP** as the **Protocol**, set the **Destination Port** to `443`.
		e. Click **OK** to save the rule.
		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	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 text here and then save the file.  Copy the transparent safe search script and paste it into the HTML editor, replacing the entire block page.
Step 6	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 7	Enable SSL Forward Proxy decryption.
	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:
		a. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.
		b. Enter a **Name** for the category, such as SearchEngineDecryption.
		c. **Add** the following to the Sites list:
		`www.bing.*`
		`www.google.*`
		`search.yahoo.*`
		d. Click **OK** to save the custom URL category object.
		2. Follow the steps to Configure SSL Forward Proxy.
		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 8	Save the configuration.	Click **Commit**.

Set Up the PAN-DB Private Cloud

To deploy one or more M-500 appliances as a PAN-DB private cloud within your network or data center, you must complete the following tasks:

- Configure the PAN-DB Private Cloud
- Configure the Firewalls to Access the PAN-DB Private Cloud

Configure 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 instructions on registering the M-500 appliance, see Register the Firewall.

 © Palo Alto Networks, Inc.

Set up the PAN-DB Private Cloud	

Step 3 Perform Initial Configuration of the M-500 Appliance.
  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:
	• 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).
	• 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).
	2. When prompted, log in to the appliance. Log in using the default username and password (admin/admin). The appliance will begin to initialize.
	3. Configure an network access settings including the IP address for the MGT interface:
	`set deviceconfig system ip-address <server-IP> netmask <netmask> default-gateway <gateway-IP> dns-setting servers primary <DNS-IP>`
	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.
	4. Configure an network access settings including the IP address for the Eth1 interface:
	`set deviceconfig system eth1 ip-address <server-IP> netmask <netmask> default-gateway <gateway-IP> dns-setting servers primary <DNS-IP>`
	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.
	5. Save your changes to the PAN-DB server.
	commit

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:
	request system system-mode pan-url-db
	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.
	2. Use the following command to verify that the mode is changed:
	`show pan-url-cloud-status`
	`hostname: M-500`
	`ip-address: 1.2.3.4`
	`netmask: 255.255.255.0`
	`default-gateway: 1.2.3.1`
	`ipv6-address: unknown`
	`ipv6-link-local-address: fe80:00/64`
	`ipv6-default-gateway:`
	`mac-address: 00:56:90:e7:f6:8e`
	`time: Mon Apr 27 13:43:59 2015`
	`uptime: 10 days, 1:51:28`
	`family: m`
	`model: M-500`
	`serial: 0073010000xxx`
	`sw-version: 7.0.0`
	`app-version: 492-2638`
	`app-release-date: 2015/03/19 20:05:33`
	`av-version: 0`
	`av-release-date: unknown`
	`wf-private-version: 0`
	`wf-private-release-date: unknown`
	`logdb-version: 7.0.9`
	`platform-family: m`
	`pan-url-db: 20150417-220`
	`system-mode: Pan-URL-DB`
	`operational-mode: normal`
	3. Use the following command to check the version of the cloud database on the appliance:
	`show pan-url-cloud-status`
	`Cloud status: Up`
	`URL database version: 20150417-220`

 © Palo Alto Networks, Inc.

Set up the PAN-DB Private Cloud	

Step 5 Install content and database updates.
  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:
	• If the PAN-DB server has direct Internet access use the following commands:
	a. To check whether a new version is published use:
	`request pan-url-db upgrade check`
	b. To check the version that is currently installed on your server use:
	`request pan-url-db upgrade info`
	c. To download and install the latest version:
	– `request pan-url-db upgrade download latest`
	– `request pan-url-db upgrade install <version latest | file>`
	d. To schedule the M-500 appliance to automatically check for updates:
	`set deviceconfig system update-schedule pan-url-db recurring weekly action download-and-install day-of-week <day of week> at <hr:min>`
	• If the PAN-DB server is offline, access the Palo Alto Networks Customer Support web site 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:
	• `scp import pan-url-db remote-port <port-number> from username@host:path`
	• `request pan-url-db upgrade install file <filename>`

Set up the PAN-DB Private Cloud		

Step 6	Set up administrative access to the PAN-DB private cloud.
	 
	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.
	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:
		a. `configure`
		b. `set mgt-config users <username> permissions role-based <superreader | superuser> yes`
		c. `set mgt-config users <username> password`
		d. `Enter password:xxxxx`
		e. `Confirm password:xxxxx`
		f. `commit`
		• To set up an administrative user with RADIUS authentication:
		a. Create RADIUS server profile.
		`set shared server-profile radius <server_profile_name> server <server_name> ip-address <ip_address> port <port_no> secret <shared_password>`
		b. Create authentication-profile.
		`set shared authentication-profile <auth_profile_name> user-domain <domain_name_for_authentication> allow-list <all> method radius server-profile <server_profile_name>`
		c. Attach the authentication-profile to the user.
		`set mgt-config users <username> authentication-profile <auth_profile_name>`
		d. Commit the changes.
		`commit`
		• To view the list of users:.
		`show mgt-config users`
		`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.	

 © Palo Alto Networks, Inc.

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.

- 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, select **Device > Setup >Content-ID**, edit the URL Filtering section and enter the **PAN-DB Server** IP address(es) or FQDN(s). The list must be comma separated.
- 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
  ```
  
  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.

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
```

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.

- Use Case: Control Web Access
- Use Case: Use URL Categories for Policy Matching



These use cases rely on User-ID to implement policies based on users and groups and a Decryption to identify and control websites that are encrypted using SSL/TLS.

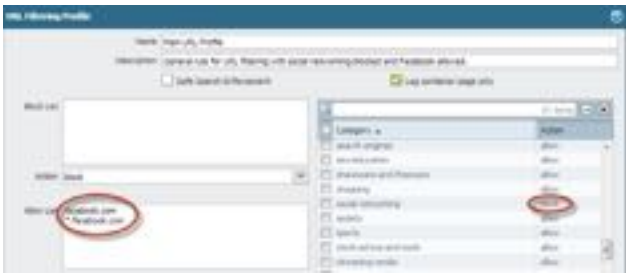
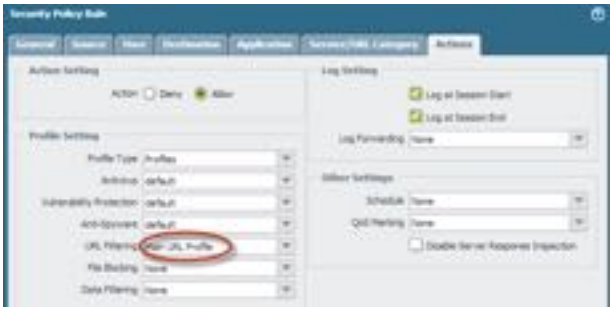
Use Case: Control Web Access

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. 2. If a valid license is not installed, see Enable PAN-DB URL Filtering.

 © Palo Alto Networks, Inc.

Control Web Access (Continued)		
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: **show user group-mapping statistics** 2. To check User Mapping from the CLI, enter the following command: **show user ip-user-mapping-mp all** 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. 2. Click the **Clone** icon. A new profile should appear named **default-1**. 3. Select the new profile and rename it.
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**. 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.  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. 2. On the **Actions** tab, select the URL profile you just created from the **URL Filtering** drop-down.  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. This rule must precede other rules because: • It is a specific rule. More specific rules must precede other rules. • Allow rule 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 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. 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.

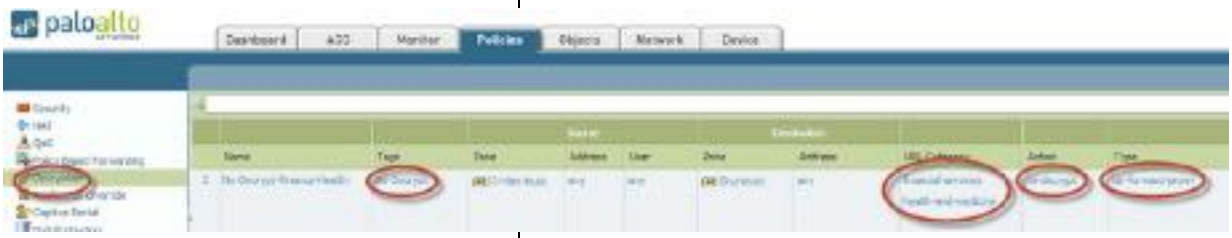
 © Palo Alto Networks, Inc.

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.

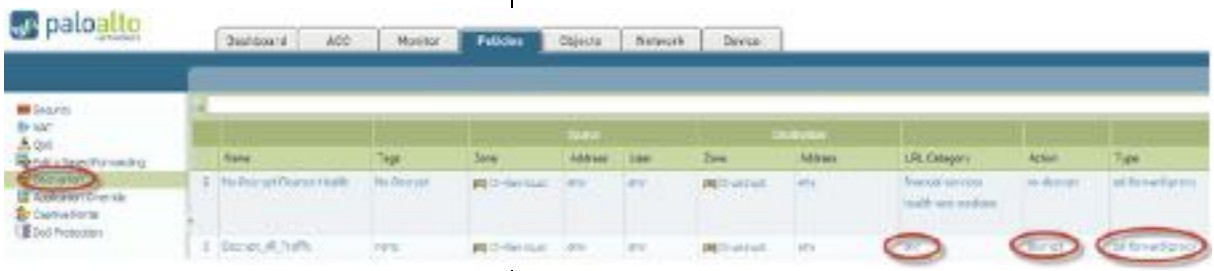
Use Case: Use URL Categories for Policy Matching

You can also use URL categories as match criteria in the following policy types: Captive Portal, Decryption, Security, and QoS. In this use case, Decryption policy rules match on URL categories to control which web categories to decrypt or not decrypt. The first rule is a no-decrypt rule instructing the firewall not to decrypt outbound user traffic to financial-services or health-and-medicine sites and the second rule instructs the firewall to decrypt all other traffic.

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**. 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. (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. 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 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. 2. Enter the following commands to enable Dynamic URL Filtering: a. **configure** b. **set deviceconfig setting url dynamic-url yes** 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

Use the following workflow to troubleshoot PAN-DB activation issues.

Troubleshoot PAN-DB Activation Issues	
Step 1	Access the PAN-OS CLI.
Step 2	Verify whether PAN-DB has been activated by running the following command: `show system setting url-database` If the response is `paloaltonetworks`, PAN-DB is the active vendor.
Step 3	Verify that the firewall has a valid PAN-DB license by running the following command: `request license info` 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 installing the license, download a new PAN-DB seed database by running the following command: `request url-filtering download paloaltonetworks region <region>`

Step 5	Check the download status by running the following command:
	`request url-filtering download status vendor paloaltonetworks`
	• 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.
	• 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:
	`admin@PA-200> **set system setting url-database paloaltonetworks**`

3.	If the problems persists, contact Palo Alto Networks Customer Support.

 © Palo Alto Networks, Inc.

PAN-DB Cloud Connectivity Issues

To check connectivity between the firewall and the PAN-DB cloud:

show url-cloud status

If the cloud is accessible, the expected response is similar to the following:

```
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 not accessible, the expected response is similar to the following:

```
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
```

Use the following checklist to identify and resolve connectivity issues:

- Does the PAN-DB URL Filtering license field shows as invalid? Obtain and install a valid PAN-DB license.
- Does the URL database status show as out of date? Download a new seed database by running the following command:

  ```
  request url-filtering download paloaltonetworks region <region>
  ```
- Does the URL protocol version show as not compatible? Upgrade PAN-OS to the latest version.
- Can you ping the PAN-DB cloud server from the firewall? Run the following command to check:

  ```
  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:

  ping source 10.1.1.5 host s0000.urlcloud.paloaltonetworks.com
- Is the firewall in an HA configuration? Verify that the HA state of the firewalls is in the active, active-primary, or active-secondary state. Access to the PAN-DB cloud will be blocked if the firewall is in a different state. Run the following command on each firewall in the pair to see the state:

  ```
  show high-availability state
  ```

If you still have problems with connectivity between the firewall and the PAN-DB cloud, contact Palo Alto Networks support.

URLs Classified as Not-Resolved

Use the following workflow to troubleshoot why 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: `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.
Step 2	If the cloud connection status shows `connected`, check the current utilization of the firewall. If firewall utilization 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: `show system resources` You can also view system resources on the System Resources widget on the **Dashboard** in the web interface.
Step 3	If the problem persist, contact Palo Alto Networks support.

Incorrect Categorization

Sometimes you may come across a URL that you believe is categorized incorrectly. Use the following workflow to determine the URL categorization for a site and request a category change, if appropriate.

Troubleshoot Incorrect Categorization Issues	

Step 1	Verify the category in the dataplane by running the following command:
	`show running url <URL>`
	For example, to view the category for the Palo Alto Networks website, run the following command:
	`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.

Step 2	Verify if the category in the management plane by running the command:
	`test url-info-host <URL>`
	For example:
	`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:
	`clear url-cache url <URL>`
	The next time the firewall 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: `test url-info-cloud <URL>`

 © Palo Alto Networks, Inc.

Troubleshoot Incorrect Categorization Issues	

Step 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:
	`clear url-cache url <URL>`
	Run the following command to delete a URL from the management plane cache:
	`delete url-database url <URL>`
	The next time the firewall 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.
Step 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: `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 PAN-DB 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 must re-download an initial seed database (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:

  `request url-filtering download paloaltonetworks region <region_name>`



Re-downloading the seed database causes the URL cache in the management plane and dataplane to 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-based 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.

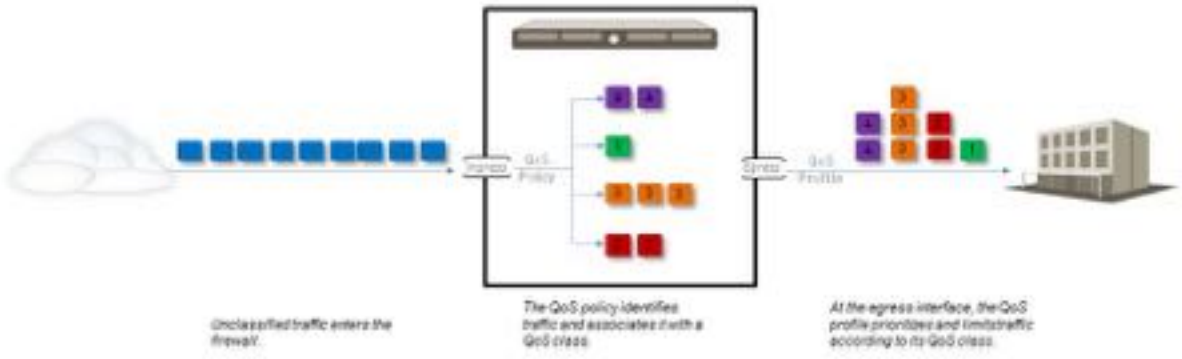
Service quality measurements subject to a QoS implementation are bandwidth (maximum rate of transfer), throughput (actual rate of transfer), latency (delay), and jitter (variance in latency). The capability to shape and control these service quality measurements makes QoS of particular importance to 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. A QoS policy rule allows you to define traffic you want to receive QoS treatment and assign that traffic a QoS class. The matching traffic is then shaped based on the QoS profile class settings as it exits the physical interface.

Each of the QoS configuration components influence each other and the QoS configuration options can be used to create a full and granular QoS implementation or can be used sparingly with minimal administrator action.

 © Palo Alto Networks, Inc.

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 Policy
- QoS Profile
- QoS Classes
- QoS Priority Queuing
- QoS Bandwidth Management
- QoS Egress Interface
- QoS for 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 for which you want to manage and/or guarantee bandwidth.

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.

 © Palo Alto Networks, Inc.

QoS Profile

Use a QoS profile rule to define values of up to eight QoS Classes contained within that single profile rule.

With a QoS profile rule, you can define QoS Priority Queuing and QoS Bandwidth Management for QoS classes. Each QoS profile rule allows you to configure individual bandwidth and priority settings for up eight QoS classes, as well as the total bandwidth alloted for the eight classes combined. Attach the QoS profile rule (or multiple QoS profile rules) to a physical interface to apply the defined priority and bandwidth settings to the traffic exiting that interface.

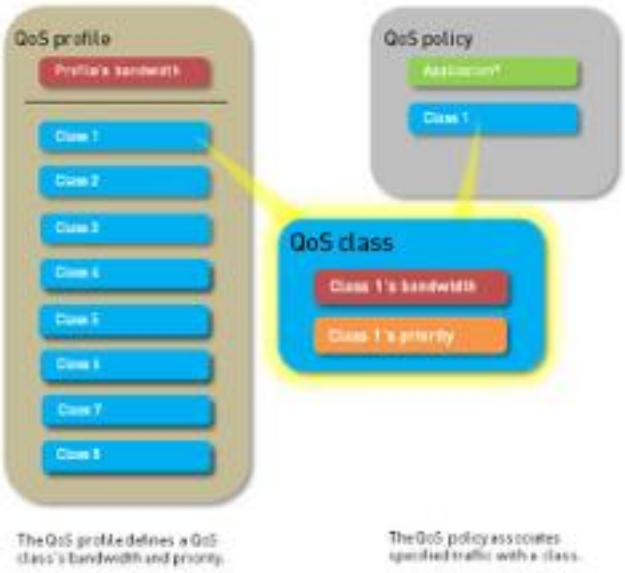
A default QoS profile rule is available on the firewall. The default profile rule and the classes defined in the profile do not have predefined maximum or guaranteed bandwidth limits.

To define priority and bandwidth settings for QoS classes, Add a QoS profile rule.

QoS Classes

A QoS class determines the priority and bandwidth for traffic matching a QoS Policy rule. You can use a QoS Profile rule to define QoS classes. 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 QoS Bandwidth Management, the fundamental mechanisms of a QoS configuration, are configured within the QoS class definition (see Step 4). For each QoS class, you can set a priority (real-time, high, medium, and low) and the maximum and guaranteed bandwidth for matching traffic. QoS priority queuing and bandwidth management determine the order of traffic and how traffic is handled upon entering or leaving a network.



QoS Priority Queuing

One of four priorities can be enforced for a QoS class: real-time, high, medium, and low. Traffic matching a QoS policy rule is assigned the QoS class associated with that rule, and the firewall treats the matching traffic based on the QoS class priority. Packets in the outgoing traffic flow are queued based on their priority until the network is ready to process the packets. Priority queuing allows you to ensure that important traffic, applications, and users take precedence. Real-time priority is typically used for applications that are particularly sensitive to latency, such as voice and video applications.

QoS Bandwidth Management

QoS bandwidth management allows you to control traffic flows on a network so that traffic does not exceed network capacity (resulting in network congestion) and also allows you to allocate bandwidth for certain types of traffic and for applications and users. With QoS, you can enforce bandwidth for traffic on a narrow or a broad scale. A QoS profile rule allows you to set bandwidth limits for individual QoS classes and the total combined bandwidth for all eight QoS classes. As part of the steps to Configure QoS, you can attach the QoS profile rule to a physical interface to enforce bandwidth settings on the traffic exiting that interface—the individual QoS class settings are enforced for traffic matching that QoS class (QoS classes are assigned to traffic matching QoS Policy rules) and the overall bandwidth limit for the profile can be applied to all clear text traffic, specific clear text traffic originating from source interfaces and source subnets, all tunneled traffic, and individual tunnel interfaces. You can add multiple profile rules to a single QoS interface to apply varying bandwidth settings to the traffic exiting that interface.

The following fields support QoS bandwidth settings:

- **Egress Guaranteed**—The amount of bandwidth guaranteed for matching traffic. When the egress guaranteed bandwidth is exceeded, the firewall passes traffic on a best-effort basis. Bandwidth that is guaranteed but is unused continues to remain available for all traffic. Depending on your QoS configuration, you can guarantee bandwidth for a single QoS class, for all or some clear text traffic, and for all or some tunneled traffic.

  Example:

  Class 1 traffic has 5 Gbps of egress guaranteed bandwidth, which means that 5 Gbps is available but is not reserved for class 1 traffic. If Class 1 traffic does not use or only partially uses the guaranteed bandwidth, the remaining bandwidth can be used by other classes of traffic. However, during high traffic periods, 5 Gbps of bandwidth is absolutely available for class 1 traffic. During these periods of congestion, any Class 1 traffic that exceeds 5 Gbps is best effort.
- **Egress Max**—The overall bandwidth allocation for matching traffic. The firewall drops traffic that exceeds the egress max limit that you set. Depending on your QoS configuration, you can set a maximum bandwidth limit for a QoS class, for all or some clear text traffic, for all or some tunneled traffic, and for all traffic exiting the QoS interface.



The cumulative guaranteed bandwidth for the QoS profile rules attached to the interface must not exceed the total bandwidth allocated to the interface.

To define bandwidth settings for QoS classes, Add a QoS profile rule. To then apply those bandwidth settings to clear text and tunneled traffic, and to set the overall bandwidth limit for a QoS interface, Enable QoS on a physical interface.

 © Palo Alto Networks, Inc.

QoS Egress Interface

Enabling a QoS profile rule 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 the egress interface for a traffic flow. 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.



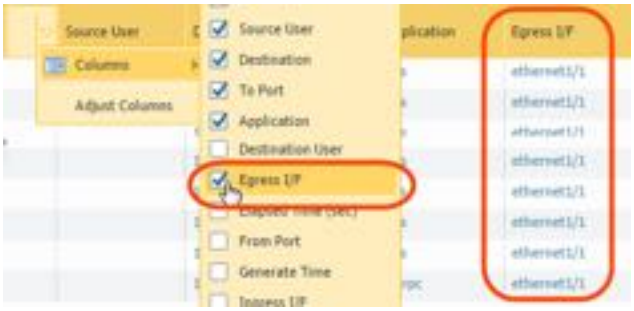
See Step 3 to learn how to Identify the egress interface for applications that you want to receive QoS treatment.

QoS for Clear Text and Tunneled Traffic

At the minimum, enabling a QoS interfaces requires you to select a default QoS profile rule that defines bandwidth and priority settings for clear text traffic egressing the interface. However, when setting up or modifying a QoS interface, you can apply granular QoS settings to outgoing clear text traffic and tunneled traffic. QoS preferential treatment and bandwidth limiting can be enforced for tunneled traffic, 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.

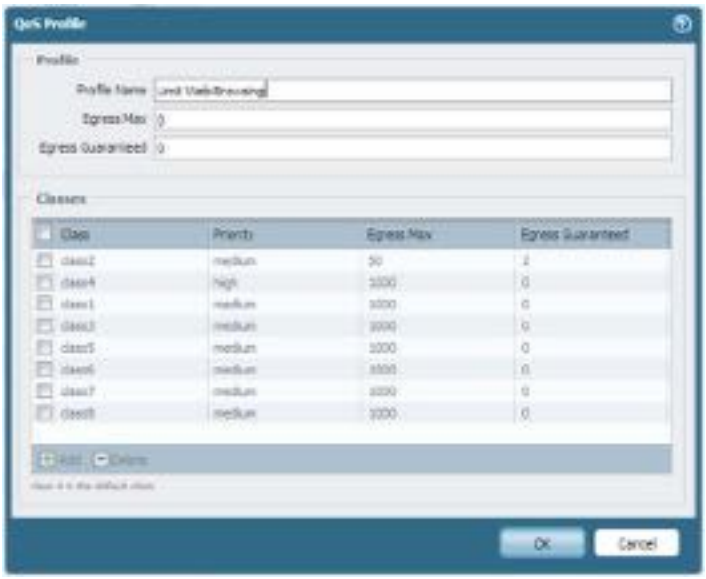
Configure QoS

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 you want to manage with QoS. 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. Click any application name to display detailed application information.
Step 2	Identify the egress interface for applications that you want to receive QoS treatment. 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 Traffic logs. To filter and only show logs for a specific application: • If an entry is displayed for the application, click the underlined link in the Application column then click the Submit icon. • If an entry is not displayed for the application, click the Add Log icon and search for the application. 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: • 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 listed in the Destination section: 

 © Palo Alto Networks, Inc.

Configure QoS (Continued)		
Step 3	Add a QoS policy rule. A QoS policy rule defines the traffic to receive QoS treatment. The firewall assigns a QoS class of service to the traffic matched to the policy rule.	1. Select **Policies > QoS** and **Add** a new policy rule. 2. On the **General** tab, give the QoS Policy Rule a descriptive **Name**. 3. Specify traffic to receive QoS treatment based on **Source**, **Destination**, **Application**, **Service/URL Category**, and **DSCP/ToS** values (the **DSCP/ToS** settings allow you to Enforce QoS Based on DSCP Classification). For example, select the **Application**, click **Add**, and select web-browsing to apply QoS to web browsing traffic. 4. (Optional) Continue to define additional parameters. For example, select **Source** and **Add** a source user to provide QoS for a specific user's web traffic. 5. Select **Other Settings** and assign a **QoS Class** to traffic matching the policy rule. For example, assign Class 2 to the user1's web traffic. 6. Click **OK.**

Configure QoS (Continued)	
Step 4 Add a QoS profile rule. A QoS profile rule allows you to define the eight classes of service that traffic can receive, including priority, and enables QoS Bandwidth Management. You can edit any existing QoS profile, including the default, by clicking the QoS profile name.	1. Select **Network > Network Profiles > QoS Profile** and **Add** a new profile. 2. Enter a descriptive **Profile Name**. 3. Set the overall bandwidth limits for the QoS profile rule: • Enter an **Egress Max** value to set the overall bandwidth allocation for the QoS profile rule. • Enter an **Egress Guaranteed** value to set the guaranteed bandwidth for the QoS Profile. Any traffic that exceeds the Egress Guaranteed value is best effort and not guaranteed. Bandwidth that is guaranteed but is unused continues to remain available for all traffic. 4. In the Classes section, specify how to treat up to eight individual QoS classes: a. **Add** a class to the QoS Profile. b. Select the **Priority** for the class: real-time, high, medium, and low. c. Enter the **Egress Max** and **Egress Guaranteed** bandwidth for traffic assigned to each QoS class. 5. Click **OK**. In the following example, the QoS profile rule Limit Web Browsing limits Class 2 traffic to a maximum bandwidth of 50Mbps and a guaranteed bandwidth of 2Mbps. 

 © Palo Alto Networks, Inc.

Configure QoS (Continued)

Step 5 Enable QoS on a physical interface.

Part of this step includes the option to select clear text and tunneled traffic for unique QoS treatment.

Check if the platform you're using supports enabling QoS on a subinterface by reviewing a summary of the Product Specifications.

1. Select **Network > QoS** and **Add** a QoS interface.
2. Select **Physical Interface** and choose the **Interface Name** of the interface on which to enable QoS.

   In the example, Ethernet 1/1 is the egress interface for web-browsing traffic (see Step 2).
3. Set the **Egress Max** bandwidth for all traffic exiting this interface.

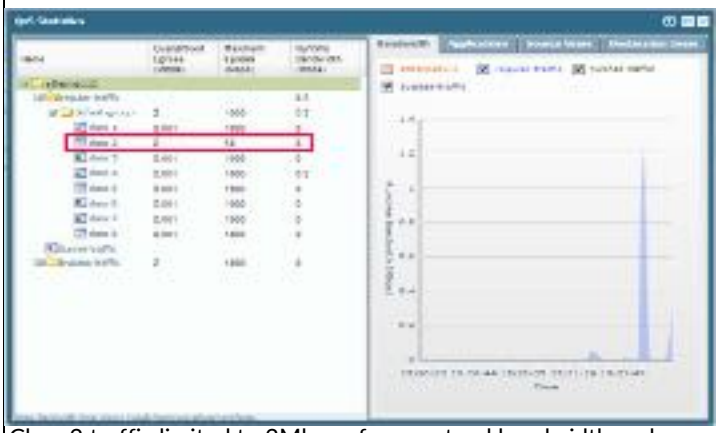
   It is a best practice to always define the Egress Max value for a QoS interface. Ensure that the cumulative guaranteed bandwidth for the QoS profile rules attached to the interface does not exceed the total bandwidth allocated to the interface.
4. Select **Turn on QoS feature on this interface**.
5. In the Default Profile section, select a QoS profile rule to apply to all **Clear Text** traffic exiting the physical interface.
6. (Optional) Select a default QoS profile rule to apply to all tunneled traffic exiting the interface.

For example, enable QoS on ethernet 1/1 and apply the bandwidth and priority settings you defined for the QoS profile rule Limit Web Browsing (Step 4) to be used as the default settings for clear text egress traffic.



7. (Optional) Continue to define more granular settings to provide QoS for Clear Text and Tunneled Traffic. Settings configured on the **Clear Text Traffic** tab and the **Tunneled Traffic** tab automatically override the default profile settings for clear text and tunneled traffic on the Physical Interface tab.
   - Select **Clear Text Traffic** and:
     - Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.
     - Click **Add** and apply a QoS profile rule to enforce clear text traffic based on source interface and source subnet.
   - Select **Tunneled Traffic** and:
     - Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.
     - Click **Add** and attach a QoS profile rule to a single tunnel interface.
8. Click **OK**.

Step 6 **Commit** the configuration.

Configure QoS (Continued)	
Step 7 Verify a 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. For example, see the statistics for ethernet 1/1 with QoS enabled:  Class 2 traffic limited to 2Mbps of guaranteed bandwidth and a maximum bandwidth of 50Mbps. Continue to click the tabs to display further information regarding applications, source users, destination users, security rules and QoS rules. Bandwidth limits shown on the **QoS Statistics** window include a hardware adjustment factor.

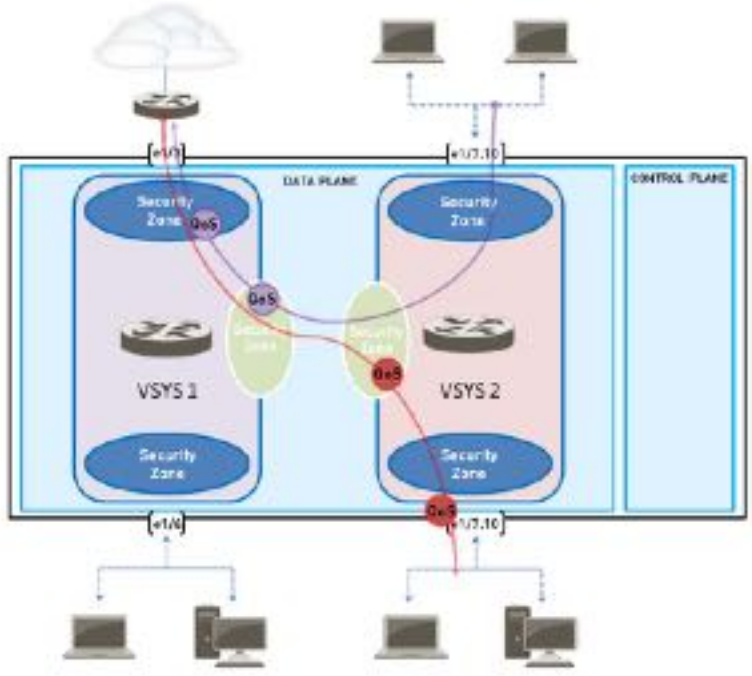
 © Palo Alto Networks, Inc.

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.

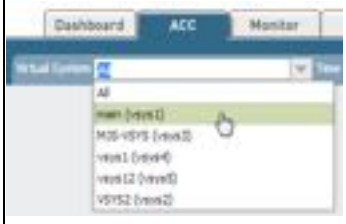
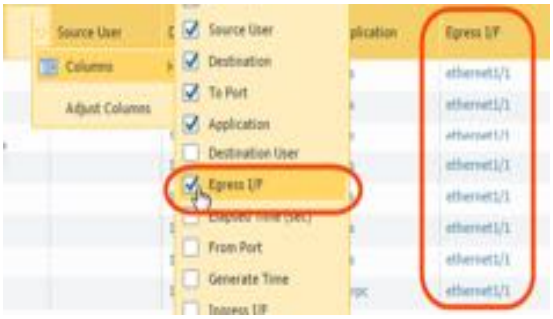
Configuring QoS for a virtual system is similar to configuring QoS on a physical firewall, with the exception that configuring QoS for a virtual system requires specifying the source and destination of traffic. Because a virtual system exists without set physical boundaries and because traffic in a virtual 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.



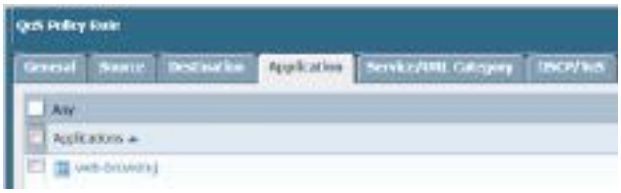
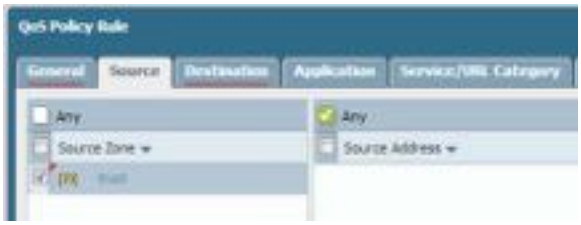
Refer to Virtual Systems for information on Virtual Systems and how to configure them.

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**. • To view configured zones, select **Network > Zones**. • 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. To view information for a specific virtual system, select the virtual system from the **Virtual System** drop-down:  Click any application name to display detailed application information.
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*.

 © Palo Alto Networks, Inc.

Configure QoS in a Virtual System Environment	
Step 4 Create a QoS Profile. 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. 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 of the **QoS Profile**, 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.

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.

© Palo Alto Networks, Inc.

Configure QoS in a Virtual System Environment	
Step 6 Enable the QoS Profile on a physical interface. 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 this example, ethernet 1/1 is the egress interface for web-browsing traffic on vsys 1 (see Step 2).  b. Select **Turn on QoS feature on this interface**. 3. On the **Physical Interface** tab, select the default QoS profile to apply to all **Clear Text** traffic. (Optional) Use the **Tunnel Interface** field to apply a default QoS profile to all tunneled traffic. 4. (Optional) On the **Clear Text Traffic** tab, configure additional QoS settings for clear text 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 additional 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 changes. 7. **Commit** the changes.
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. • 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 request (for example) high priority or best effort delivery for traffic. Session-Based DSCP Classification allows you to both honor DSCP values for incoming traffic and to mark a session with a DSCP value as session traffic exits the firewall. This enables all inbound and outbound traffic for a session can receive continuous QoS treatment as it flows through your network. For example, inbound return traffic from an external server can now be treated with the same QoS priority that the firewall initially enforced for the outbound flow based on the DSCP value the firewall detected at the beginning of the session. 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).

Different types of DSCP markings indicate different levels of service:

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.

- **Expedited Forwarding (EF)**: Can be used to request low loss, low latency and guaranteed bandwidth for traffic. Packets with EF codepoint values 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.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.

 © Palo Alto Networks, Inc.

Apply QoS Based on DSCP/ToS Marking		
Step 1	Perform the preliminary steps to Configure QoS.	
Step 2	Define the traffic to receive QoS treatment based on DSCP value.	1. Select **Policies > QoS** and **Add** or modify an existing QoS rule and populate required fields. 2. **Select DSCP/ToS and select Codepoints.** 3. **Add** a DSCP/ToS codepoints for which you want to enforce QoS. 4. 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. 5. 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 rule matches to traffic marked with any EF codepoint value. 6. Select **Other Settings** and assign a **QoS Class** to traffic matched to the QoS rule. In this example, assign Class 1 to sessions where a DSCP marking of AF11 is detected for the first packet in the session. 7. Click **OK** to save the QoS rule.
Step 3	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. 2. **Add** or modify a profile class. For example, because Step 2 showed steps to classify AF11 traffic as Class 1 traffic, you could add or modify a **class1** entry. 3. Select a **Priority** for the class of traffic, such as **high**. 4. Click **OK** to save the QoS Profile.
Step 4	Enable QoS on an interface.	Select **Network > QoS** and **Add** or modify an existing interface and **Turn on QoS feature on this interface**. 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).

Apply QoS Based on DSCP/ToS Marking		
Step 5	Enable DSCP Marking. 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. 2. Select **Actions** and in the **QoS Marking** drop-down, choose **Follow-Client-to-Server-Flow**. 3. Click **OK** to save your changes. 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 6	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:

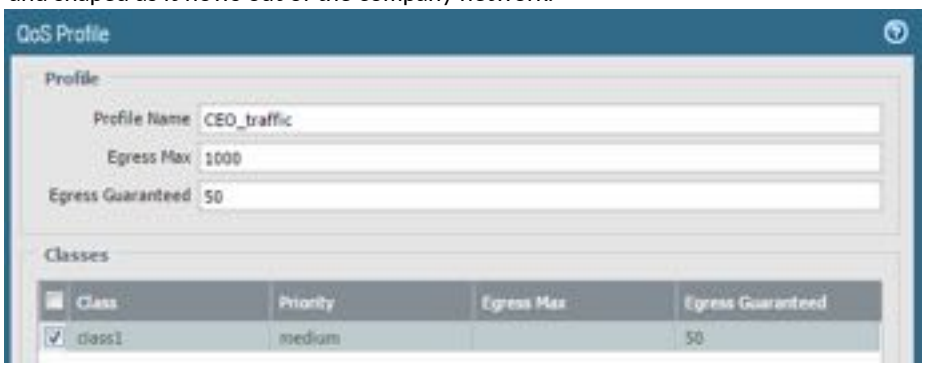
- Use Case: QoS for a Single User
- Use Case: QoS for Voice and Video Applications

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:



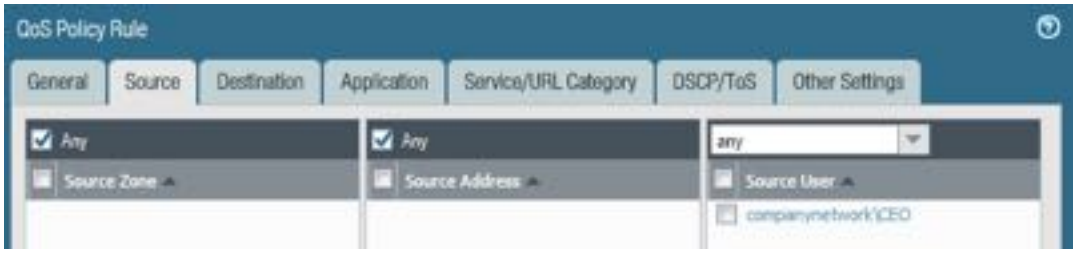
The admin assigns a guaranteed bandwidth (**Egress Guaranteed**) of 50 Mbps to ensure that the CEO will have that amount that bandwidth guaranteed to her at all times (more than she would need to use), regardless of network congestion.

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.

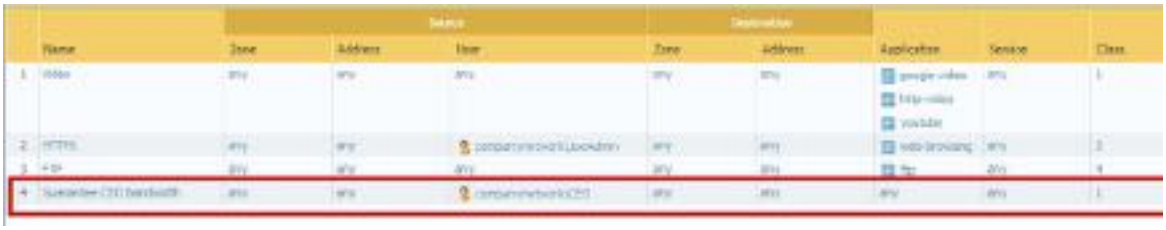
It is a best practice to populate the **Egress Max** field for a QoS profile, even if the max bandwidth of the profile matches the max bandwidth of the interface. The QoS profile's max bandwidth should never exceed the max bandwidth of the interface you are planning to enable QoS on.

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:

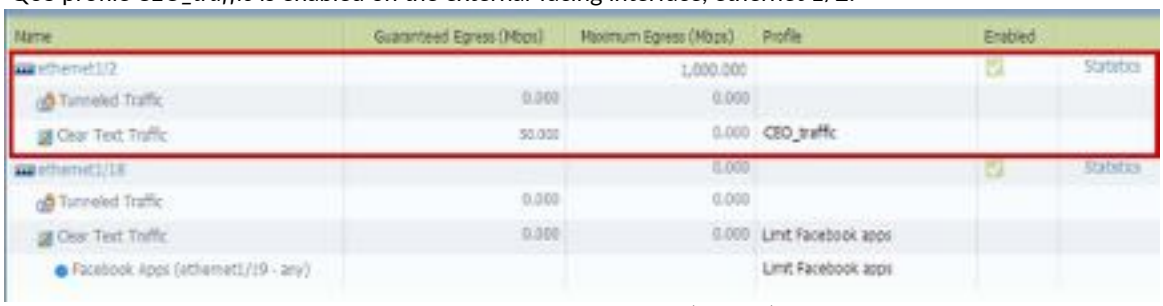


Because the admin wants to ensure that all traffic originating from the CEO is guaranteed by the QoS profile and associated QoS policy he created, he selects the *CEO_traffic* to apply to **Clear Text** traffic flowing from ethernet 1/2.

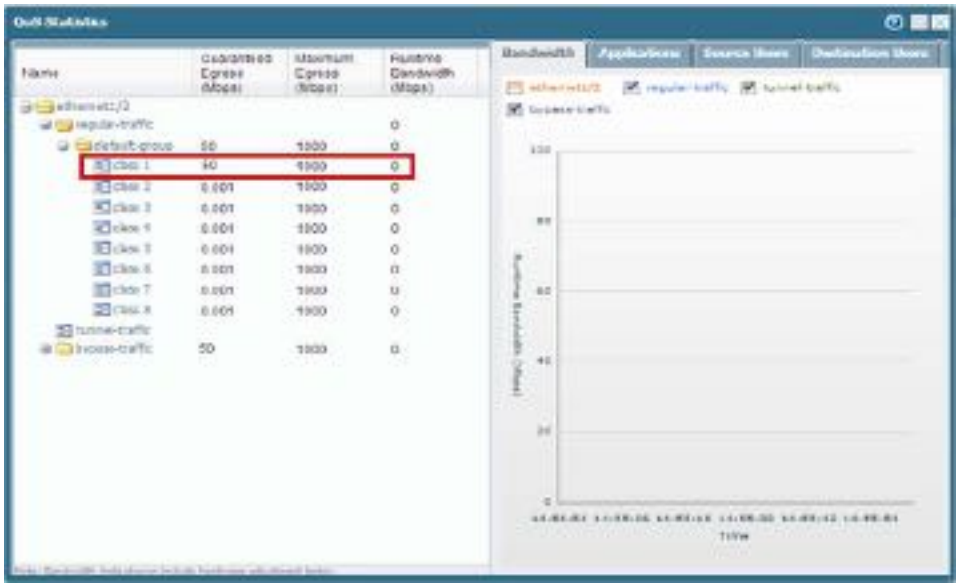
 © Palo Alto Networks, Inc.

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) and then enable QoS on the network's internal-facing interface on the **Network > QoS** page (instead of the external-facing interface).

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.

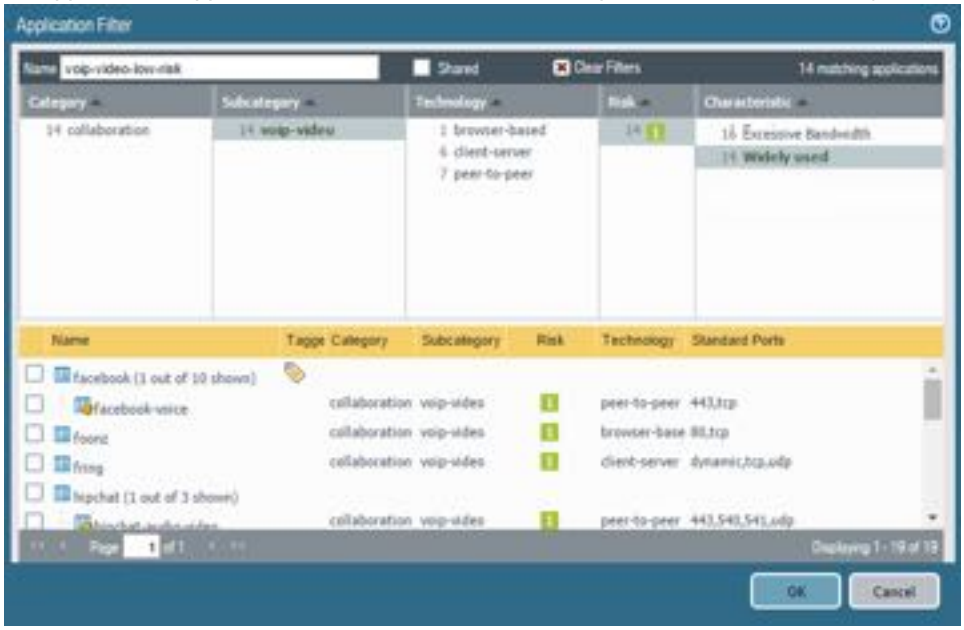
Ensure Quality for Voice and Video Applications
Step 1 The admin creates a QoS profile, defining Class 2 so that Class 2 traffic receives real-time priority and on an interface with a maximum bandwidth of 1000 Mbps, is guaranteed a bandwidth of 250 Mbps at all times, including peak periods of network 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 firewall web interface, the admin selects **Network > Network Profiles > Qos Profile** page, clicks **Add**, enters the **Profile Name** *ensure voip-video traffic* and defines Class 2 traffic. 

 © Palo Alto Networks, Inc.

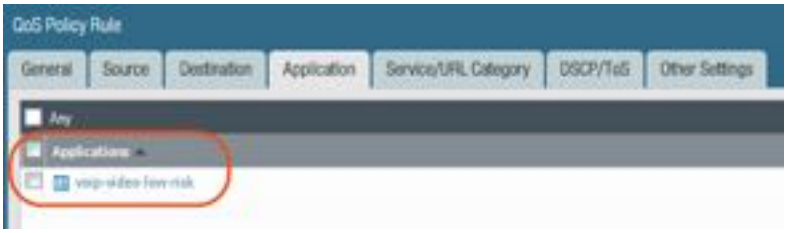
Ensure Quality for Voice and Video Applications (Continued)

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.

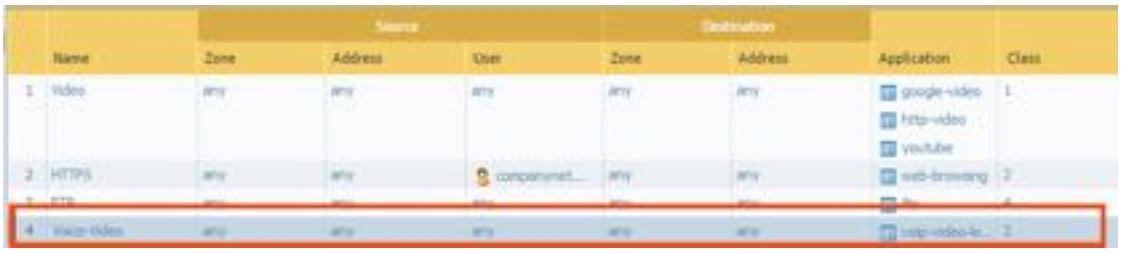
The application filter is a dynamic tool that, when used to filter applications in the QoS policy, allows QoS to be applied to all applications that meet the criteria of *voip-video*, *low risk*, and *widely used* at any given time.



The admin names the **Application Filter** *voip-video-low-risk* and includes it in the QoS policy:



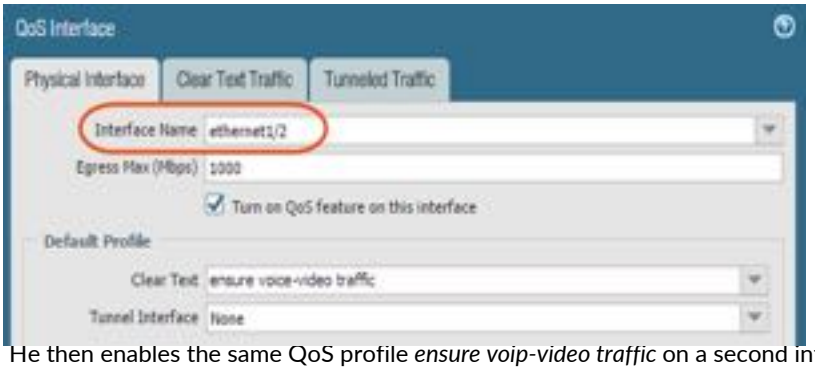
The admin names the QoS policy *Voice-Video* and selects Other Settings to assign all traffic matched to the policy Class 2. He is going to use the *Voice-Video* QoS policy for both incoming and outgoing QoS traffic, so he sets **Source** and **Destination** information to **Any**:



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).

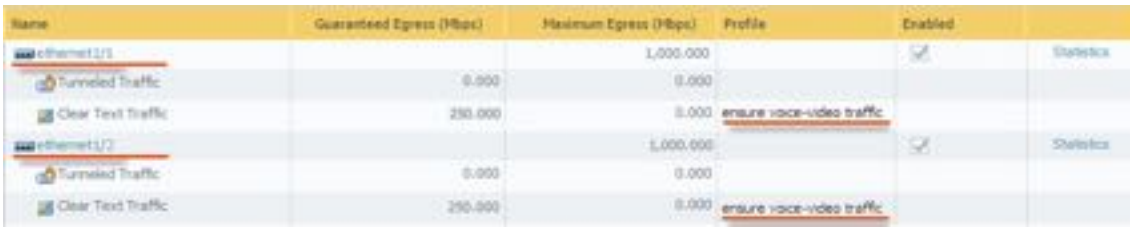
The admin begins by enabling the QoS profile he created, *ensure voice-video traffic* (Class 2 in this profile is associated with policy, *Voice-Video*) on the external-facing interface, in this case, ethernet 1/2.



He then enables the same QoS profile *ensure voip-video traffic* on a second interface, the internal-facing interface (in this case, ethernet 1/1).



Step 4 The admin selects **Network > QoS** to confirm 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.

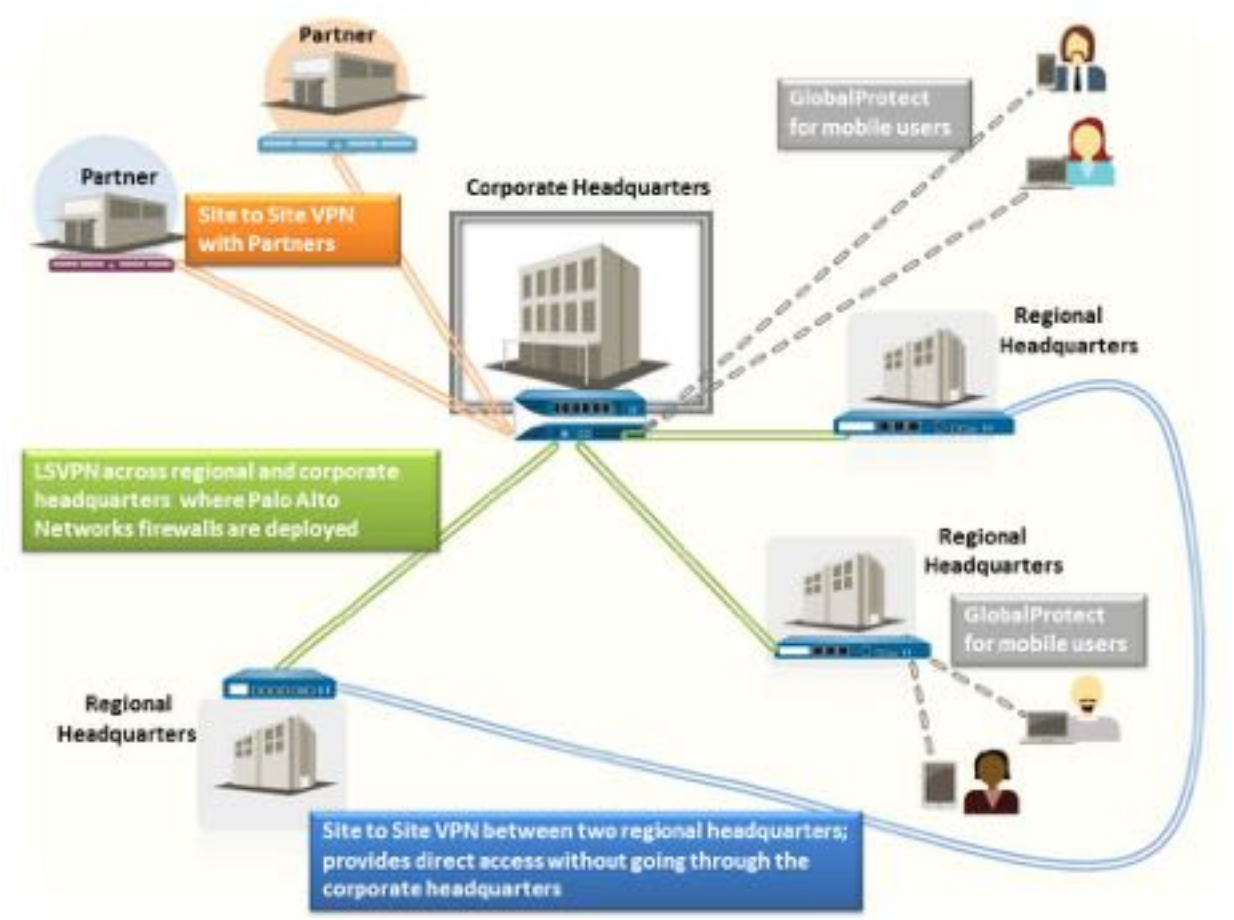
- VPN Deployments
- Site-to-Site VPN Overview
- Site-to-Site VPN Concepts
- Set Up Site-to-Site VPN
- Site-to-Site VPN Quick Configs

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 1,024 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.

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.

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.

 © Palo Alto Networks, Inc.

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.

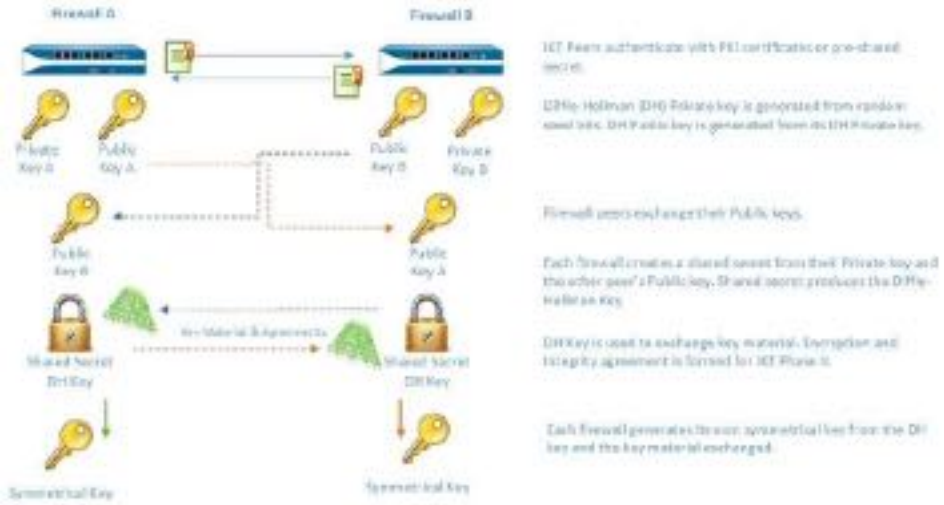
If the destination IP is unreachable, you either configure the firewall to wait for the tunnel to recover or configure automatic failover to another tunnel. In either case, the firewall generates a system log that alerts you to a tunnel failure and renegotiates the IPSec keys to accelerate recovery.

See Set Up Tunnel Monitoring for configuration details.

© Palo Alto Networks, Inc.

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:



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, aes-256-cbc, or des

 © Palo Alto Networks, Inc.

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:

- 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.

Algorithms Supported for IPSEC Authentication and Encryption

ESP	AH
Diffie Hellman (DH) exchange options supported	
• Group 1—768 bits • Group 2—1024 bits (the default) • Group 5—1536 bits • Group 14—2048 bits. • Group 19— 256-bit elliptic curve group • Group 20—384-bit elliptic curve group • 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
• des	Data Encryption Standard (DES) with a security strength of 56 bits

ESP	AH
Authentication algorithms supported	
• md5	• md5
• 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.
- Supports Hash and URL certificate exchange to reduce fragmentation.

 © Palo Alto Networks, Inc.

- 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. With **Strict Cookie Validation** 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.

- 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.

 © Palo Alto Networks, Inc.

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.

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. 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. (The **Version** selection also determines which options are available on the **Advanced Options** tab.)

 © Palo Alto Networks, Inc.

Set Up an IKE Gateway (Continued)		
Step 2	Establish the local endpoint of the tunnel (gateway).	1. For **Address Type**, click **IPv4** or **IPv6**. 2. Select the physical, outgoing **Interface** on the firewall where the local gateway resides. 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. 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 the next step. If you choose Certificate, skip to Configure certificate-based authentication.
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**. Use a maximum of 255 ASCII or non-ASCII characters. Generate a key that is difficult to crack with dictionary attacks; use a pre-shared key generator, if necessary. 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. 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. 4. Proceed to Step 7 and continue from there.

Set Up an IKE Gateway (Continued)	
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. • If you want to **Import** a certificate, Import a Certificate for IKEv2 Gateway Authentication and then return to this task. • If you want to **Generate** a new certificate, generate a certificate on the firewall and then return to this task. 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. 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. 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. 5. Select one type of **Peer ID Check**: • **Exact**—Check this to ensure that the local setting and peer IKE ID payload match exactly. • **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. 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. 7. Choose a **Certificate Profile** from the drop-down. A certificate profile contains information about how to authenticate the peer gateway. 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 (Continued)		
Step 7	Configure advanced options for the gateway.	1. Select the **Advanced Options** tab. 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. 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. 4. If you chose **IKEv1 only mode** earlier, on the IKEv1 tab: • 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. 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. • 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. • (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. • 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. 5. If you chose **IKEv2 only mode** or **IKEv2 preferred mode** in Step 1, on the IKEv2 tab: • 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. • 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. • **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.
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	
Step 1	Select **Device > Certificates**, and if your platform supports multiple virtual systems, for **Location**, select the appropriate virtual system.
Step 2	On the **Device Certificates** tab, select the certificate to **Export** to the server. The status of the certificate should be valid, not expired. The firewall will not stop you from exporting an invalid certificate.
Step 3	For **File Format**, select **Binary Encoded Certificate (DER)**.
Step 4	Leave **Export private key** clear. Exporting the private key is unnecessary for Hash and URL.
Step 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**.

 © Palo Alto Networks, Inc.

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**. 2. In the Import Certificate window, enter a **Certificate Name** for the certificate you are importing. 3. Select **Shared** if this certificate is to be shared among multiple virtual systems. 4. For **Certificate File**, **Browse** to the certificate file. Click on the file name and click **Open**, which populates the **Certificate File** field. 5. For **File Format**, select one of the following: • **Base64 Encoded Certificate (PEM)**—Contains the certificate, but not the key. It is cleartext. • **Encrypted Private Key and Certificate (PKCS12)**—Contains both the certificate and the key. 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: • 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. • Enter a **Passphrase** and **Confirm Passphrase**. 7. Click **OK**.
Step 2	Configure certificate-based authentication.	

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.

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. 2. For the **Key Lifetime**, select a unit (**Seconds**, **Minutes**, **Hours**, or **Days**) and enter a value. The minimum is three minutes. 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). 2. Click **OK**.
Step 2	Save the configuration	Click **OK** and **Commit**.

Configure IKEv2 Traffic Selectors

In IKEv2, you can configure 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. Use the following workflow to configure traffic selectors.

Configure Traffic Selectors for IKEv2	
Step 1	Select **Network > IPSec Tunnels > Proxy IDs**.
Step 2	Select the **IPv4** or **IPv6** tab.
Step 3	Click **Add** and enter the **Name** in the **Proxy ID** field.
Step 4	In the **Local** field, enter the **Source IP Address**.
Step 5	In the **Remote** field, enter the **Destination IP Address**.
Step 6	In the **Protocol** field, select the transport protocol (**TCP** or **UDP**) from the drop-down.
Step 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 are ready for use.

 © Palo Alto Networks, Inc.

▲ 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**. 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. 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: • DH Group—**group20**, **group19**, **group14**, **group5**, **group2**, and **group1**. • Authentication—**sha512**, **sha384**, **sha256**, **sha1**, **md5**. • Encryption—**aes-256-cbc**, **aes-192-cbc**, **aes-128-cbc**, **3des**, **des**.  DES is available to provide backward compatibility with legacy devices that do not support stronger encryption, but as a best practice always use a stronger encryption algorithm, such as 3DES or AES if the peer can support it.
Step 3	Specify the duration for which the key is valid and the re-authentication interval. 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. 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.
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 Configure advanced options for the 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**. 2. Enter a **Name** for the new profile. 3. Select the **IPSec Protocol**—ESP or AH—that you want to apply to secure the data as it traverses across the tunnel. 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. 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: • 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**, **des**. DES is available to provide backward compatibility with legacy devices that do not support stronger encryption, but as a best practice always use a stronger encryption algorithm, such as 3DES or AES if the peer can support it. • 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. 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**. 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. 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. 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 Set up key exchange.

 © 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** and then **Add** a new tunnel configuration.
Step 2	On the General tab, enter a **Name** for the new tunnel.
Step 3	Select the **Tunnel interface** that will be used to set up the IPSec tunnel. To create a new tunnel interface: 1. Select **Tunnel Interface > New Tunnel Interface**. (You can also 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, select 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, select **New Zone**. 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, and **Add** the IP address and network mask, for example 10.31.32.1/32. 6. To save the interface configuration, click **OK**.

Set Up an IPSec Tunnel (Continued)	
Step 4 (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 assign an IPv6 **Address** to the tunnel interface, **Add** the 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.
Step 5 Set up key exchange.	Configure one of the following types of key exchange: **Set up Auto Key exchange** 1. Select the IKE Gateway. To set up an IKE gateway, see Set Up an IKE Gateway. 2. (Optional) Select the default IPSec Crypto Profile. To create a new IPSec Profile, see Define IPSec Crypto Profiles. **Set up Manual Key exchange** 1. 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. 2. 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. 3. Select the protocol to be used—**AH** or **ESP**. 4. For AH, select the **Authentication** method from the drop-down and enter a **Key** and then **Confirm Key**. 5. 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. 6. Specify the **SPI** for the remote peer. 7. Enter the **Remote Address**, the IP address of the remote peer.

 © Palo Alto Networks, Inc.

Set Up an IPSec Tunnel (Continued)		
Step 6	Protect against a replay attack. 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 7	(Optional) 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. If there are multiple sessions inside the tunnel (each with a different TOS value), copying the TOS header can cause the IPSec packets to arrive out of order.
Step 8	Enable Tunnel Monitoring. You must 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: 1. Specify a **Destination IP** address on the other side of the tunnel to determine if the tunnel is working properly. 2. Select a **Profile** to determine the action on tunnel failure. To create a new profile, see Define a Tunnel Monitoring Profile.
Step 9	Create a Proxy ID to identify the VPN peers. This step is required only if the VPN peer uses policy-based VP).	1. Select **Network > IPSec Tunnels** and click **Add**. 2. Select the **Proxy IDs** tab. 3. Select the **IPv4** or **IPv6** tab. 4. Click **Add** and enter the **Proxy ID** name. 5. Enter the **Local** IP address or subnet for the VPN gateway. 6. Enter the **Remote** address for the VPN gateway. 7. Select the **Protocol** from the drop-down: • **Number**—Specify the protocol number (used for interoperability with third-party devices). • **Any**—Allows TCP and/or UDP traffic. • **TCP**—Specify the Local Port and Remote Port numbers. • **UDP**—Specify the Local Port and Remote Port numbers. 8. Click **OK**.
Step 10	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. • **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. • **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. 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. The threshold specifies the number of heartbeats to wait before taking the specified action. The range is 2-100 and the default is 5. 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	
Step 1	Select **Network > IPSec Tunnels**.
Step 2	View the **Tunnel Status**. • Green indicates a valid IPSec SA tunnel. • Red indicates that IPSec SA is not available or has expired.
Step 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.
Step 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.

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 IPSec Tunnel

▲ Refresh and Restart Behaviors

▲ Refresh or Restart an IKE Gateway or IPSec Tunnel

Enable or Disable an IKE Gateway or IPSec Tunnel

Enable or Disable an IKE Gateway or IPSec Tunnel	
• Enable or disable an IKE gateway.	1. Select **Network > Network Profiles > IKE Gateways** and select the gateway you want to enable or disable. 2. At the bottom of the screen, click **Enable** or **Disable**.

Enable or Disable an IKE Gateway or IPSec Tunnel (Continued)	
• Enable or disable an IPSec tunnel.	1. Select **Network > IPSec Tunnels** and select the tunnel you want to enable or disable. 2. At the bottom of the screen, click **Enable** or **Disable**.

Refresh and Restart Behaviors

The refresh and restart behaviors for an IKE gateway and IPSec tunnel are as follows:

Phase	Refresh	Restart
IKE Gateway (IKE Phase 1)	Updates the onscreen statistics for the selected IKE gateway. Equivalent to issuing a second `show` command in the CLI (after an initial `show` command).	Restarts the selected IKE gateway. **IKEv2**: Also restarts any associated child IPSec security associations (SAs). **IKEv1**: Does not restart the associated IPSec SAs. A restart is disruptive to all existing sessions. Equivalent to issuing a `clear, test, show` command sequence in the CLI.
IPSec Tunnel (IKE Phase 2)	Updates the onscreen statistics for the selected IPSec tunnel. Equivalent to issuing a second `show` command in the CLI (after an initial `show` command).	Restarts the IPSec tunnel. A restart is disruptive to all existing sessions. Equivalent to issuing a `clear, test, show` command sequence in the CLI.

Refresh or Restart an IKE Gateway or IPSec Tunnel

Restart 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. 2. In the row for that tunnel, under the Status column, click **IKE Info**. 3. At the bottom of the IKE Info screen, click the action you want: • **Refresh**—Updates the statistics on the screen. • **Restart**—Clears the SAs, so traffic is dropped until the IKE negotiation starts over and the tunnel is recreated.

 © Palo Alto Networks, Inc.

Refresh or Restart an IKE Gateway or IPSec Tunnel	
• Refresh or restart an IPSec tunnel. 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. 2. In the row for that tunnel, under the Status column, click **Tunnel Info**. 3. At the bottom of the Tunnel Info screen, click the action you want: • **Refresh**—Updates the onscreen statistics. • **Restart**—Clears the SAs, so traffic is dropped until the IKE negotiation starts over and the tunnel is recreated.

Test VPN Connectivity

Test VPN Connectivity	
Step 1	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>`
Step 2	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.
Step 3	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>`
Step 4	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.

Step 5	To view the VPN traffic flow information, use the following command:
	`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.

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.`
or
`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.
	• 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...`
or
`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`
or
`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:
	• pfs is either enabled or disabled on both VPN peers
	• 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 Create a Proxy ID to identify the VPN peers..

 © Palo Alto Networks, Inc.

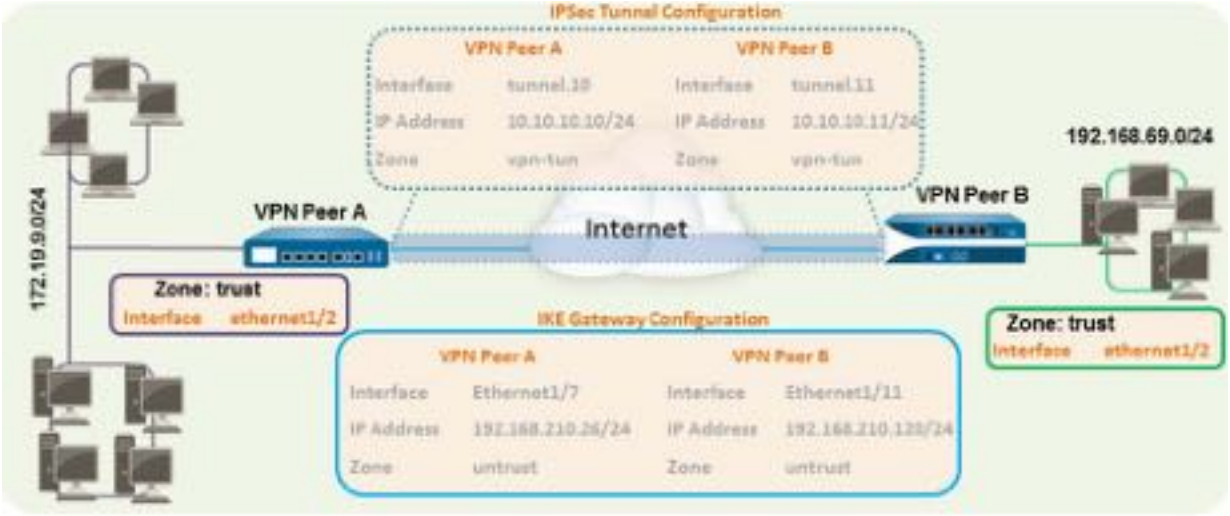
Site-to-Site VPN Quick Configs

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.



Quick Config: Site-to-Site VPN with Static Routing	
Step 1 Configure a Layer 3 interface. 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. 2. Select **Layer3** from the **Interface Type** drop-down. 3. On the **Config** tab, select the **Security Zone** to which the interface belongs: • 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. • 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**. 4. Select the **Virtual Router** to use. 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. 6. To save the interface configuration, click **OK**. In this example, the configuration for VPN Peer A is: • **Interface**—ethernet1/7 • **Security Zone**—untrust • **Virtual Router**—default • **IPv4**—192.168.210.26/24 The configuration for VPN Peer B is: • **Interface**—ethernet1/11 • **Security Zone**—untrust • **Virtual Router**—default • **IPv4**—192.168.210.120/24

 © Palo Alto Networks, Inc.

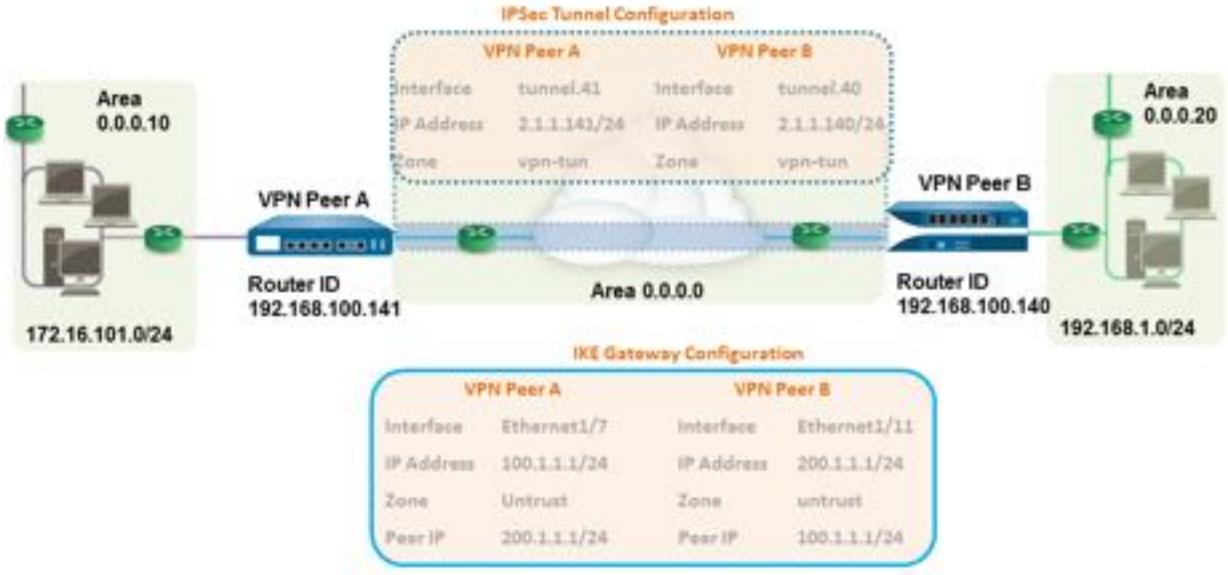
Quick Config: Site-to-Site VPN with Static Routing (Continued)		
Step 2	Create a tunnel interface and attach it to a virtual router and security zone.	1. Select **Network > Interfaces > Tunnel** and click **Add**. 2. In the **Interface Name** field, specify a numeric suffix, such as **.1**. 3. On the **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
Step 4	Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2). 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. 2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.

Quick Config: Site-to-Site VPN with Static Routing (Continued)	
Step 5 Set up the IKE Gateway.	1. Select **Network > Network Profiles > IKE Gateway.** 2. Click **Add** and configure the options in the **General** tab. In this example, the configuration for VPN Peer A is: • **Interface**—ethernet1/7 • **Local IP address**—192.168.210.26/24 • **Peer IP type/address**—static/192.168.210.120 • **Preshared keys**—enter a value • **Local identification**—None; this means that the local IP address will be used as the local identification value. The configuration for VPN Peer B is: • **Interface**—ethernet1/11 • **Local IP address**—192.168.210.120/24 • **Peer IP type/address**—static/192.168.210.26 • **Preshared keys**—enter same value as on Peer A • **Local identification**—None 3. Select **Advanced Phase 1 Options** and select the IKE Crypto profile you created earlier to use for IKE phase 1.
Step 6 Set up the IPSec Tunnel.	1. Select **Network > IPSec Tunnels**. 2. Click **Add** and configure the options in the **General** tab. In this example, the configuration for VPN Peer A is: • **Tunnel Interface**—tunnel.11 • **Type**—Auto Key • **IKE Gateway**—Select the IKE Gateway defined above. • **IPSec Crypto Profile**—Select the IPSec Crypto profile defined in Step 4. The configuration for VPN Peer B is: • **Tunnel Interface**—tunnel.12 • **Type**—Auto Key • **IKE Gateway**—Select the IKE Gateway defined above. • **IPSec Crypto Profile**—Select the IPSec Crypto defined in Step 4. 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. 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**. 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. 2. Select **Layer3** from the **Interface Type** drop-down. 3. On the **Config** tab, select the **Security Zone** to which the interface belongs: • 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. • 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**. 4. Select the **Virtual Router** to use. 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. 6. To save the interface configuration, click **OK**. In this example, the configuration for VPN Peer A is: • **Interface**—ethernet1/7 • **Security Zone**—untrust • **Virtual Router**—default • **IPv4**—100.1.1.1/24 The configuration for VPN Peer B is: • **Interface**—ethernet1/11 • **Security Zone**—untrust • **Virtual Router**—default • **IPv4**—200.1.1.1/24

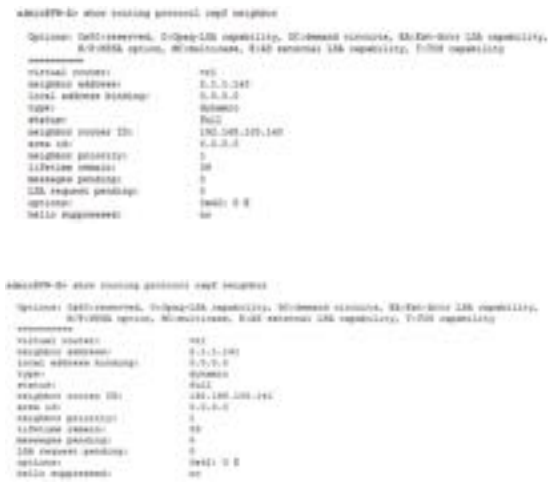
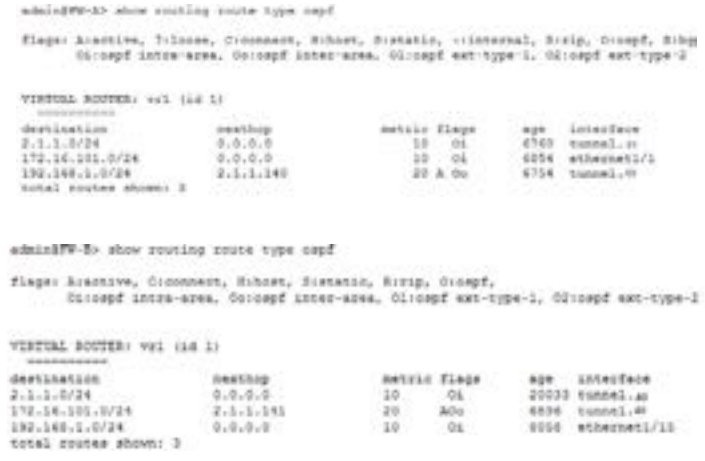
 © Palo Alto Networks, Inc.

Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF (Continued)		
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, **.11**. 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. 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. 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. 6. To save the interface configuration, click **OK**. In this example, the configuration for VPN Peer A is: • **Interface**—tunnel.41 • **Security Zone**—vpn_tun • **Virtual Router**—default • **IPv4**—2.1.1.141/24 The configuration for VPN Peer B is: • **Interface**—tunnel.40 • **Security Zone**—vpn_tun • **Virtual Router**—default • **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). 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. 2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.

Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF (Continued)		
Step 4	Set up the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall. For more information on the OSPF options that are available on the firewall, see Configure OSPF. 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. 2. Select **OSPF** (for IPv4) or **OSPFv3** (for IPv6) and select **Enable**. 3. In this example, the OSPF configuration for VPN Peer A is: – **Router ID**: 192.168.100.141 – **Area ID**: 0.0.0.0 that is assigned to the tunnel.1 interface with Link type: p2p – **Area ID**: 0.0.0.10 that is assigned to the interface Ethernet1/1 and Link Type: Broadcast The OSPF configuration for VPN Peer B is: – **Router ID**: 192.168.100.140 – **Area ID**: 0.0.0.0 that is assigned to the tunnel.1 interface with Link type: p2p – **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. 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.** 2. Click **Add** and configure the options in the **General** tab**.** In this example, the configuration for VPN Peer A is: • **Interface**—ethernet1/7 • **Local IP address**—100.1.1.1/24 • **Peer IP address**—200.1.1.1/24 • **Preshared keys**—enter a value The configuration for VPN Peer B is: • **Interface**—ethernet1/11 • **Local IP address**—200.1.1.1/24 • **Peer IP address**—100.1.1.1/24 • **Preshared keys**—enter same value as on Peer A 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 (Continued)		
Step 6	Set up the IPSec Tunnel.	1. Select **Network > IPSec Tunnels**. 2. Click **Add** and configure the options in the **General** tab. In this example, the configuration for VPN Peer A is: • **Tunnel Interface**—tunnel.41 • **Type**—Auto Key • **IKE Gateway**—Select the IKE Gateway defined above. • **IPSec Crypto Profile**—Select the IKE Gateway defined above. The configuration for VPN Peer B is: • **Tunnel Interface**—tunnel.40 • **Type**—Auto Key • **IKE Gateway**—Select the IKE Gateway defined above. • **IPSec Crypto Profile**—Select the IKE Gateway defined above. 3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity. 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**. 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.

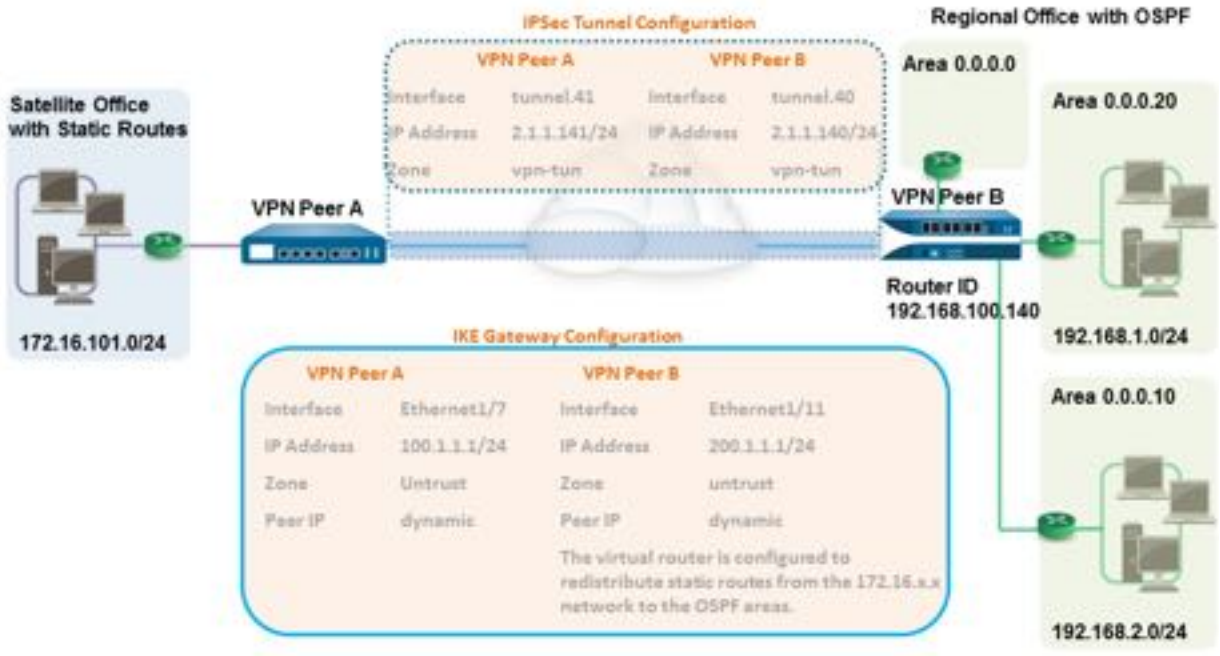
Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF (Continued)		
Step 8	Verify OSPF adjacencies and routes from the CLI.	Verify that both the firewalls can see each other as neighbors with full status. Also confirm that the IP address of the VPN peer's tunnel interface and the OSPF Router ID. Use the following CLI commands on each VPN peer. • **show routing protocol ospf neighbor**  • **show routing route type ospf** 
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. 2. Select **Layer3** from the **Interface Type** drop-down. 3. On the **Config** tab, select the **Security Zone** to which the interface belongs: • 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. • 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**. 4. Select the **Virtual Router** to use. 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. 6. To save the interface configuration, click **OK**. In this example, the configuration for VPN Peer A is: • **Interface**—ethernet1/7 • **Security Zone**—untrust • **Virtual Router**—default • **IPv4**—100.1.1.1/24 The configuration for VPN Peer B is: • **Interface**—ethernet1/11 • **Security Zone**—untrust • **Virtual Router**—default • **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). 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. 2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.

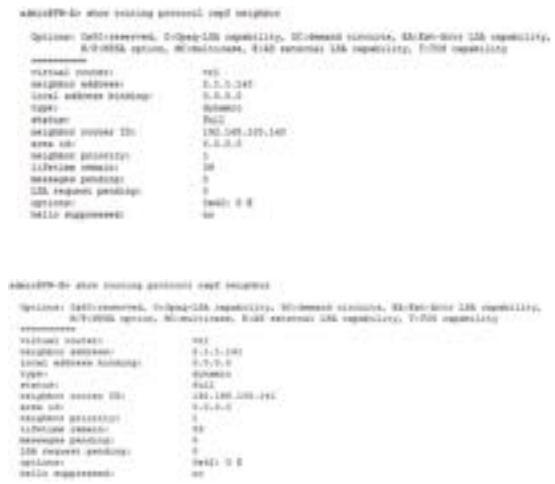
 © Palo Alto Networks, Inc.

Quick Config: Site-to-Site VPN with Static and Dynamic Routing (Continued)	
Step 3 Set up the IKE Gateway. 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.** 2. Click **Add** and configure the options in the **General** tab. In this example, the configuration for VPN Peer A is: • **Interface**—ethernet1/7 • **Local IP address**—100.1.1.1/24 • **Peer IP type**—dynamic • **Preshared keys**—enter a value • **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer A. • **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer B The configuration for VPN Peer B is: • **Interface**—ethernet1/11 • **Local IP address**—200.1.1.1/24 • **Peer IP address**—dynamic • **Preshared keys**—enter same value as on Peer A • **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer B • **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer A 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 (Continued)		
Step 4	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, say, **.41**. 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. 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. This IP address will be used to route traffic to the tunnel and to monitor the status of the tunnel. 6. To save the interface configuration, click **OK**. In this example, the configuration for VPN Peer A is: • **Interface**—tunnel.41 • **Security Zone**—vpn_tun • **Virtual Router**—default • **IPv4**—2.1.1.141/24 The configuration for VPN Peer B is: • **Interface**—tunnel.42 • **Security Zone**—vpn_tun • **Virtual Router**—default • **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. 2. Select **Static Routes**, and **Add** tunnel.41 as the **Interface** for routing traffic with a **Destination** in the 192.168.x.x network.
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. 2. Select **Static Routes** and **Add** the tunnel IP address as the next hop for traffic in the 172.168.x.x. network. Assign the desired route metric; using a lower the value makes the a higher priority for route selection in the forwarding table. 3. Select **OSPF** (for IPv4) or **OSPFv3** (for IPv6) and select **Enable**. 4. In this example, the OSPF configuration for VPN Peer B is: • Router ID: 192.168.100.140 • Area ID: 0.0.0.0 is assigned to the interface Ethernet 1/12 Link type: Broadcast • Area ID: 0.0.0.10 that is assigned to the interface Ethernet1/1 and Link Type: Broadcast • Area ID: 0.0.0.20 is assigned to the interface Ethernet1/15 and Link Type: Broadcast

 © Palo Alto Networks, Inc.

Quick Config: Site-to-Site VPN with Static and Dynamic Routing (Continued)		
Step 7	Create a redistribution profile to inject the static routes into the OSPF autonomous system.	1. Create a redistribution profile on VPN Peer B. a. Select **Network > Virtual Routers**, and select the router you used above. b. Select **Redistribution Profiles,** and click **Add**. 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. d. Set **Source Type** as **static**, and click **OK**. The static route defined in Step 6-2 will be used for the redistribution. 2. Inject the static routes in to the OSPF system. a. Select **OSPF> Export Rules** (for IPv4) or **OSPFv3> Export Rules** (for IPv6). b. Click **Add**, and select the redistribution profile that you just created. 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.** 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. e. Click **OK** to save the changes.
Step 8	Set up the IPSec Tunnel.	1. Select **Network > IPSec Tunnels**. 2. Click **Add** and configure the options in the **General** tab. In this example, the configuration for VPN Peer A is: • **Tunnel Interface**—tunnel.41 • **Type**—Auto Key • **IKE Gateway**—Select the IKE Gateway defined above. • **IPSec Crypto Profile**—Select the IKE Gateway defined above. The configuration for VPN Peer B is: • **Tunnel Interface**—tunnel.40 • **Type**—Auto Key • **IKE Gateway**—Select the IKE Gateway defined above. • **IPSec Crypto Profile**—Select the IKE Gateway defined above. 3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity. 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**. 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 (Continued)	
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. • **`show routing protocol ospf neighbor`**  • **show routing route** The following is an example of the output on each VPN peer. 
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 *satellites*. This solution uses certificates for firewall 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:

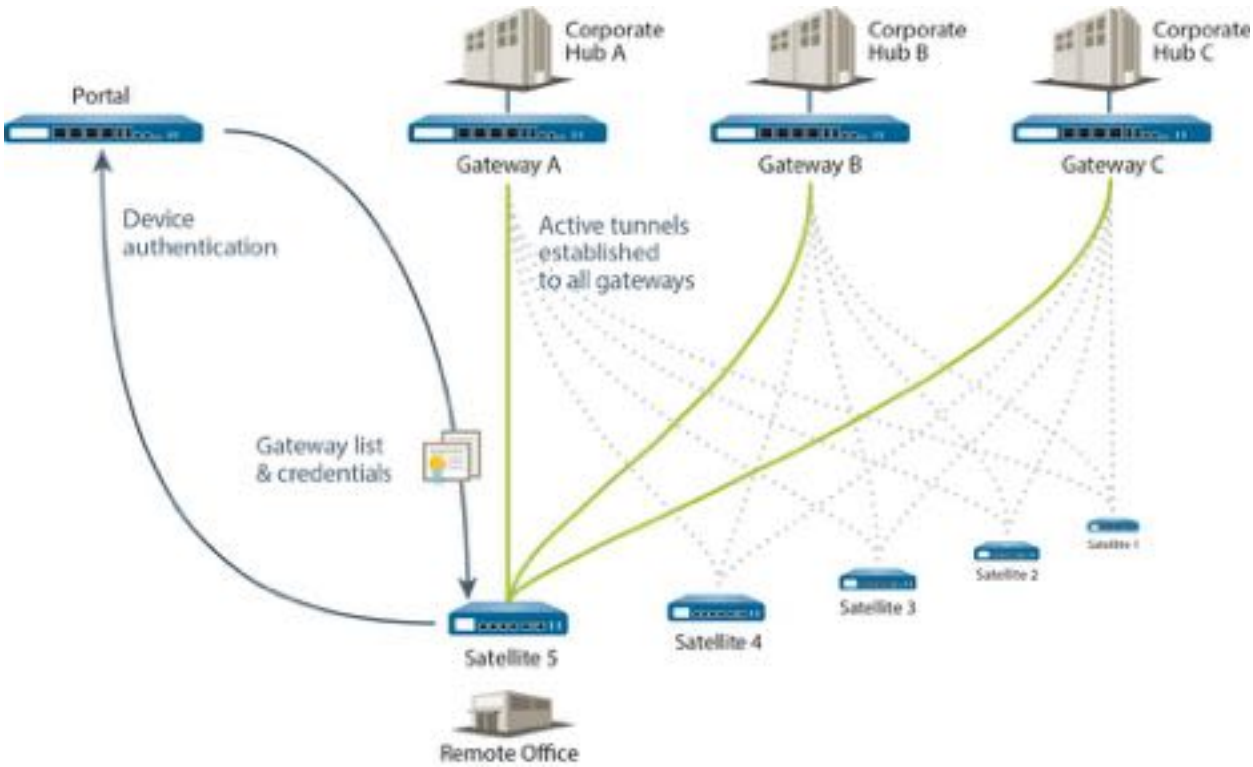
- LSVPN Overview
- Create Interfaces and Zones for the LSVPN
- Enable SSL Between GlobalProtect LSVPN Components
- Configure the Portal to Authenticate Satellites
- Configure GlobalProtect Gateways for LSVPN
- Configure the GlobalProtect Portal for LSVPN
- Prepare the Satellite to Join the LSVPN
- Verify 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. The portal and each gateway and satellite all require a Layer 3 interface to enable traffic to be routed between sites. If the gateway and portal are on the same firewall, you can use a single interface for both components. IPv6 addresses are not supported with LSVPN.	1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for GlobalProtect LSVPN. 2. Select **Layer3** from the **Interface Type** drop-down. 3. On the **Config** tab, select the **Security Zone** to which the interface belongs: • 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. • 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**. 4. Select the **Virtual Router** to use. 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. 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.  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.  Make sure to enable User-ID in the zone where the VPN tunnels terminate.	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. • (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**. 4. Select the **Virtual Router**. 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. 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, a policy rule enables traffic between the *lsvpn-tun* zone and the *L3-Trust* zone.
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
- Deploy Client Certificates to the GlobalProtect Satellites Using SCEP

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. You can also configure the GlobalProtect portal to act as a Simple Certificate Enrollment Protocol (SCEP) client to issue client certificates to GlobalProtect 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 satellites. You don't need to create SSL/TLS service profiles for the satellites because the portal will issue a server certificate for each satellite during the first connection as part of the satellite 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. 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:

© 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: 1. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Generate**. 2. Enter a **Certificate Name**, such as `LSVPN_CA`. 3. Do not select a value in the **Signed By** field (this is what indicates that it is self-signed). 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. For the portal and each gateway, you must assign an SSL/TLS service profile that references a unique self-signed server certificate. 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. 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 for each gateway you will deploy: a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Generate**. b. Enter a **Certificate Name**. c. Enter the FQDN (recommended) or IP address of the interface where you plan to configure the gateway in the **Common Name** field. d. In the **Signed By** field, select the `LSVPN_CA` certificate you just created. 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**. f. **Generate** the certificate. 2. Configure an SSL/TLS Service Profile for the portal and each gateway: a. Select **Device > Certificate Management > SSL/TLS Service Profile** and click **Add**. b. Enter a **Name** to identify the profile and select the server **Certificate** you just created for the portal or gateway. c. Define the range of TLS versions (**Min Version** to **Max Version**) allowed for communicating with satellites 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.  **Best Practices:** •Export the self-signed server certificates issued by the root CA from the portal and import them onto the gateways. •Be sure to issue a unique server certificate for each gateway. •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**. 2. Select **Encrypted Private Key and Certificate (PKCS12)** from the **File Format** drop-down. 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. 4. On the gateway, select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**. 5. Enter a **Certificate Name**. 6. Enter the path and name to the **Certificate File** you just downloaded from the portal, or **Browse** to find the file. 7. Select **Encrypted Private Key and Certificate (PKCS12)** as the **File Format**. 8. Enter the path and name to the PKCS12 file in the **Key File** field or **Browse** to find it. 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. 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. a. Select **Device > Certificate Management > Certificates > Device Certificates**. b. Select the root CA certificate used to issue certificates for the LSVPN components and click **Export**. 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.) 2. On the firewalls hosting the gateways and satellites, import the root CA certificate. a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**. b. Enter a **Certificate Name** that identifies the certificate as your client CA certificate. c. **Browse** to the **Certificate File** you downloaded from the CA. d. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**. e. Select the certificate you just imported on the **Device Certificates** tab to open it. f. Select **Trusted Root CA** and then click **OK**. g. **Commit** the changes.

Deploy SSL Server Certificates to the GlobalProtect Components (Continued)		
Step 5	Create a certificate profile. 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**. 2. Make sure **Username Field** is set to **None**. 3. In the **CA Certificates** field, click **Add**, select the Trusted Root CA certificate you imported in the previous step. 4. (Optional, but recommended) Enable use of CRL and/or OCSP to enable certificate status verification. 5. Click **OK** to save the profile.
Step 6	Save the configuration.	Click **Commit**.

Deploy Client Certificates to the GlobalProtect Satellites Using SCEP

As an alternative method for deploying client certificates to satellites, you can configure your GlobalProtect portal to act as a Simple Certificate Enrollment Protocol (SCEP) client to a SCEP server in your enterprise PKI. SCEP operation is dynamic in that the enterprise PKI generates a certificate when the portal requests it and sends the certificate to the portal.

When the satellite device requests a connection to the portal or gateway, it also includes its serial number with the connection request. The portal submits a CSR to the SCEP server using the settings in the SCEP profile and automatically includes the serial number of the device in the subject of the client certificate. After receiving the client certificate from the enterprise PKI, the portal transparently deploys the client certificate to the satellite device. The satellite device then presents the client certificate to the portal or gateway for authentication.

Deploy Server Certificates to the GlobalProtect Components Using SCEP		
Step 1	Create a SCEP profile.	1. Select **Device > Certificate Management > SCEP** and then **Add** a new profile. 2. Enter a **Name** to identify the SCEP profile. 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.
Step 2	(Optional) To make the SCEP-based certificate generation more secure, configure a SCEP challenge-response mechanism between the PKI and portal for each certificate request. After you configure this mechanism, its operation is invisible, and no further input from you is necessary. To comply with the U.S. Federal Information Processing Standard (FIPS), use a **Dynamic** SCEP challenge and specify a **Server URL** that uses HTTPS (see Step 7).	Select one of the following options: • **None**—(Default) The SCEP server does not challenge the portal before it issues a certificate. • **Fixed**—Obtain the enrollment challenge password from the SCEP server (for example, `http://10.200.101.1/CertSrv/mscep_admin/`) in the PKI infrastructure and then copy or enter the password into the **Password** field. • **Dynamic**—Enter the SCEP **Server URL** where the portal-client submits these credentials (for example, `http://10.200.101.1/CertSrv/mscep_admin/`), and a username and OTP of your choice. The username and password can be the credentials of the PKI administrator.

 © Palo Alto Networks, Inc.

Deploy Server Certificates to the GlobalProtect Components Using SCEP (Continued)		
Step 3	Specify the settings for the connection between the SCEP server and the portal to enable the portal to request and receive client certificates. To identify the satellite, the portal automatically includes the device serial number in the CSR request to the SCEP server. Because the SCEP profile requires a value in the **Subject** field, you can leave the default `$USERNAME` token even though the value is not used in client certificates for LSVPN.	1. Configure the **Server URL** that the portal uses to reach the SCEP server in the PKI (for example, `http://10.200.101.1/certsrv/mscep/`). 2. Enter a string (up to 255 characters in length) in the **CA-IDENT Name** field to identify the SCEP server. 3. Select the **Subject Alternative Name Type**: • **RFC 822 Name**—Enter the email name in a certificate's subject or Subject Alternative Name extension. • **DNS Name**—Enter the DNS name used to evaluate certificates. • **Uniform Resource Identifier**—Enter the name of the resource from which the client will obtain the certificate. • **None**—Do not specify attributes for the certificate.
Step 4	(Optional) Configure cryptographic settings for the certificate.	• Select the key length (**Number of Bits**) for the certificate. If the firewall is in FIPS-CC mode and the key generation algorithm is RSA. The RSA keys must be 2048 bits or larger. • Select the **Digest for CSR** which indicates the digest algorithm for the certificate signing request (CSR): SHA1, SHA256, SHA384, or SHA512.
Step 5	(Optional) Configure the permitted uses of the certificate, either for signing or encryption.	• To use this certificate for signing, select the **Use as digital signature** check box. This enables the endpoint use the private key in the certificate to validate a digital signature. • To use this certificate for encryption, select the **Use for key encipherment** check box. This enables the client use the private key in the certificate to encrypt data exchanged over the HTTPS connection established with the certificates issued by the SCEP server.
Step 6	(Optional) To ensure that the portal is connecting to the correct SCEP server, enter the **CA Certificate Fingerprint**. Obtain this fingerprint from the SCEP server interface in the Thumbprint field.	1. Enter the URL for the SCEP server's administrative UI (for example, `http://<hostname or IP>/CertSrv/mscep_admin/`). 2. Copy the thumbprint and enter it in the **CA Certificate Fingerprint** field.
Step 7	Enable mutual SSL authentication between the SCEP server and the GlobalProtect portal. This is required to comply with the U.S. Federal Information Processing Standard (FIPS). FIPS-CC operation is indicated on the firewall login page and in its status bar.	Select the SCEP server's root **CA Certificate**. Optionally, you can enable mutual SSL authentication between the SCEP server and the GlobalProtect portal by selecting a **Client Certificate**.
Step 8	Save and commit the configuration.	1. Click **OK** to save the settings and close the SCEP configuration. 2. **Commit** the configuration. The portal attempts to request a CA certificate using the settings in the SCEP profile and saves it to the firewall hosting the portal. If successful, the CA certificate is shown in **Device > Certificate Management > Certificates**.

Deploy Server Certificates to the GlobalProtect Components Using SCEP (Continued)		
Step 9	(Optional) If after saving the SCEP profile, the portal fails to obtain the certificate, you can manually generate a certificate signing request (CSR) from the portal.	1. Select **Device > Certificate Management > Certificates > Device Certificates** and then click **Generate**. 2. Enter a **Certificate Name**. This name cannot contain spaces. 3. Select the **SCEP Profile** to use to submit a CSR to your enterprise PKI. 4. Click **OK** to submit the request and generate the certificate.

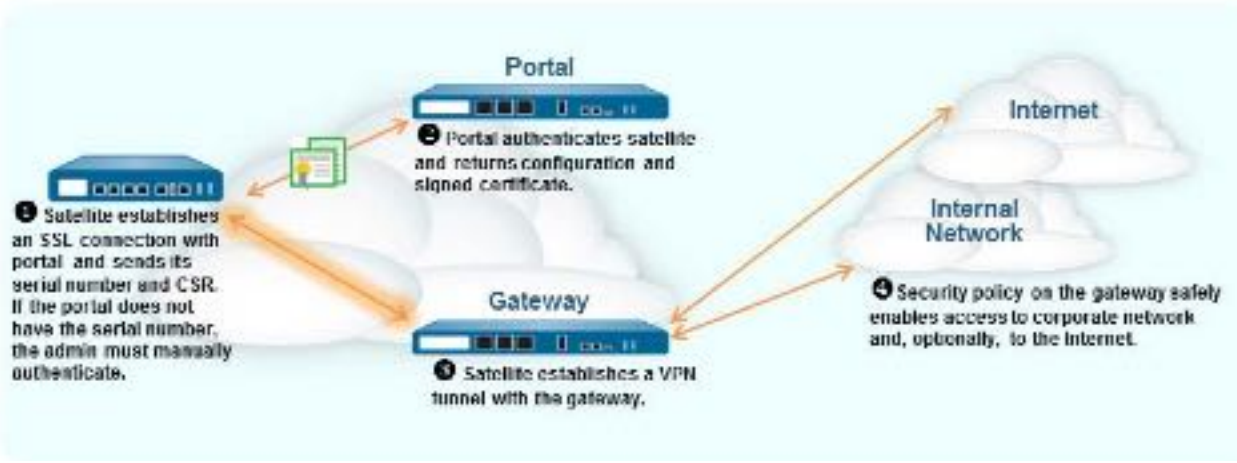
 © Palo Alto Networks, Inc.

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 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 satellites 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. The server profile defines how the firewall connects to an external authentication service to validate the authentication credentials that the satellite administrator enters. If you use local authentication, skip this step and instead add a local user for the satellite administrator: see Configure the user account.	Configure a server profile for the authentication service type: • Configure a RADIUS Server Profile. • Configure a TACACS+ Server Profile. • 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. • Configure a Kerberos Server Profile.
Step 2	Configure an authentication profile. The authentication profile defines which server profile to use to authenticate satellites.	1. Select **Device > Authentication Profile** and click **Add**. 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 the previous step. If you added a local user instead, set the **Type** to **Local Database**. 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 satellites 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**. 2. In the **General** screen, enter a **Name** for the gateway. The gateway name should have no spaces and, as a best practice, should include the location or other descriptive information to help users and administrators identify the gateway. 3. (Optional) Select the virtual system to which this gateway belongs from the **Location** field.
Step 2	Specify the network information that enables satellite devices to connect to the gateway. If you haven't created the network interface for the gateway, see Create Interfaces and Zones for the LSVPN for instructions.	1. Select the **Interface** that satellites will use for ingress access to the gateway. 2. Select the **IP Address** for gateway access. 3. Click **OK** to save changes.

Configure the Gateway for LSVPN (Continued)		
Step 3	Specify how the gateway authenticates satellites attempting to establish tunnels. If you haven't yet created an SSL/TLS Service profile for the gateway, see Deploy Server Certificates to the GlobalProtect LSVPN Components. If you haven't set up the authentication profiles or certificate profiles, see Configure the Portal to Authenticate Satellites for instructions. If you have not yet set up the certificate profile, see Enable SSL Between GlobalProtect LSVPN Components for instructions.	On the GlobalProtect Gateway Configuration dialog, select Authentication and then configure any of the following: • To secure communication between the gateway and the satellites, select the **SSL/TLS Service Profile** for the gateway. • To specify the authentication profile to use to authenticate satellites, **Add** a Client Authentication. Then, enter a **Name** to identify the configuration, select **OS**: **Satellite** to apply the configuration to all satellites, and specify the **Authentication Profile** to use to authenticate the satellite. You can also select a **Certificate Profile** for the gateway to use to authenticate satellite devices attempting to establish tunnels.
Step 4	Configure the tunnel parameters and enable tunneling.	1. On the GlobalProtect Gateway Configuration dialog, select **Satellite > Tunnel Settings**. 2. Select the **Tunnel Configuration** check box to enable tunneling. 3. Select the **Tunnel Interface** you defined to terminate VPN tunnels established by the GlobalProtect satellites when you performed the task to Create Interfaces and Zones for the LSVPN. 4. (Optional) If you want to preserve the Type of Service (ToS) information in the encapsulated packets, select **Copy TOS**. If there are multiple sessions inside the tunnel (each with a different TOS value), copying the TOS header can cause the IPSec packets to arrive out of order.
Step 5	(Optional) Enable tunnel monitoring. Tunnel monitoring enables satellites 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. 2. Specify the **Destination IP** address the satellites 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. 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. 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 satellites during establishment of the IPSec tunnel.  You can also configure the satellite 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 Configuration dialog, select **Satellite > Network Settings**. 2. (Optional) If clients local to the satellite need to resolve FQDNs on the corporate network, configure the gateway to push DNS settings to the satellites in one of the following ways: • If the gateway has an interface that is configured as a DHCP client, you can set the **Inheritance Source** to that interface and assign the same settings received by the DHCP client to GlobalProtect satellites. You can also inherit the DNS suffix from the same source. • Manually define the **Primary DNS**, **Secondary DNS**, and **DNS Suffix** settings to push to the satellites. 3. To specify the **IP Pool** of addresses to assign the tunnel interface on the satellites when the VPN is established, click **Add** and then specify the IP address range(s) to use. 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: • 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. • 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. • 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. 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 > Route Filter**. 2. Select the **Accept published routes** check box. 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 Configuration dialog. 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 satellites.
- Configure the Portal to Authenticate Satellites by defining the authentication profile that the portal will use to authenticate satellites 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**. 2. On the **General** tab, enter a **Name** for the portal. The portal name should not contain any spaces. 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. If you haven't yet created the network interface for the portal, see Create Interfaces and Zones for the LSVPN for instructions.	1. Select the **Interface** that satellites will use for ingress access to the portal. 2. Select the **IP Address** for satellite access to the portal.

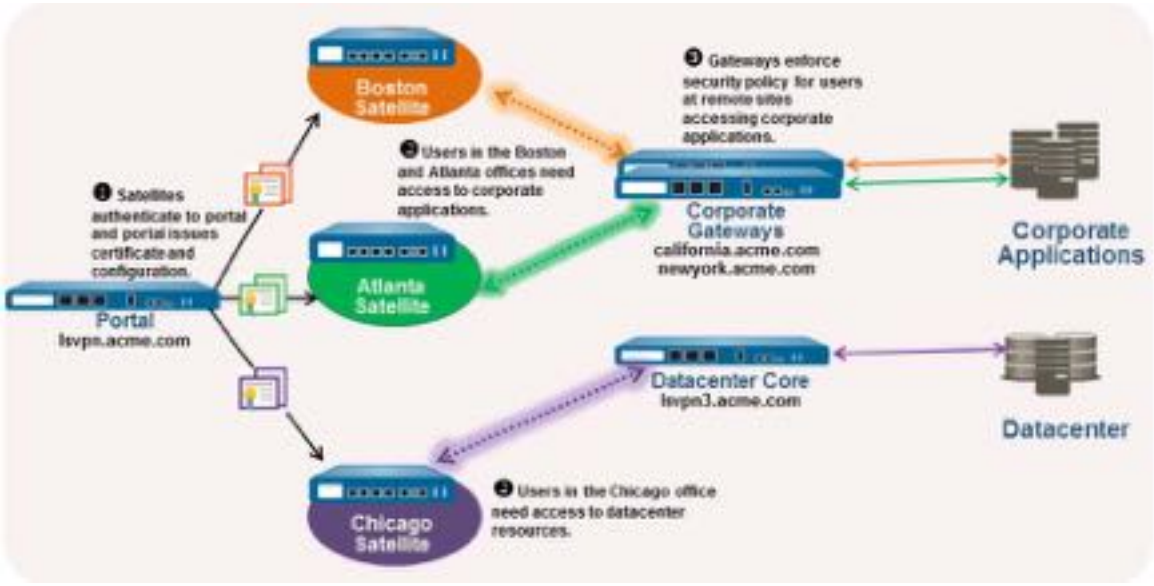
 © Palo Alto Networks, Inc.

Configure the Portal for LSVPN (Continued)		
Step 3	Specify an SSL/TLS Service profile to use to enable the satellite to establish an SSL/TLS connection to the portal. 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. On the GlobalProtect Portal Configuration dialog, select **Authentication**. 2. Select the **SSL/TLS Service Profile**.
Step 4	Specify an authentication profile and optional certificate profile for authenticating satellites. If the portal can't validate the serial numbers of connecting satellites, it will fall back to the authentication profile. Therefore, before you can save the portal configuration (by clicking **OK**), you must Configure an authentication profile.	**Add** a Client Authentication, and then enter a **Name** to identify the configuration, select **OS**: **Satellite** to apply the configuration to all satellites, and specify the **Authentication Profile** to use to authenticate satellite devices. You can also specify a **Certificate Profile** for the portal to use to authenticate satellite devices.
Step 5	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 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 data center.



Use the following procedure to create one or more satellite configurations.

Create a GlobalProtect Satellite Configuration	
Step 1 Add a satellite configuration. The satellite configuration specifies the GlobalProtect LSVPN configuration settings to deploy to the connecting satellites. You must define at least one satellite configuration.	1. Select **Network > GlobalProtect > Portals** and select the portal configuration for which you want to add a satellite configuration and then select the **Satellite** tab. 2. In the Satellite section, click **Add** 3. Enter a **Name** for the configuration. If you plan to create multiple configurations, make sure the name you define for each is descriptive enough to allow you to distinguish them. 4. To change how often a satellite should check the portal for configuration updates specify a value in the **Configuration Refresh Interval (hours)** field (range is 1-48; default is 24).
Step 2 Specify the satellites to which to deploy this configuration. 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 5 for instructions on ordering the list of satellite configurations.	Specify the match criteria for the satellite configuration as follows: • To restrict this configuration to satellites 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. • 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). Before you can restrict the configuration to specific groups, you must Map Users to Groups.

 © Palo Alto Networks, Inc.

Create a GlobalProtect Satellite Configuration (Continued)	
Step 3 Specify the gateways that satellites with this configuration can establish VPN tunnels with.  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**. 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. 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. 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 4 Save the satellite configuration.	1. Click **OK** to save the satellite configuration. 2. If you want to add another satellite configuration, repeat the previous steps.
Step 5 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**. • To move a satellite configuration down on the list of configurations, select the configuration and click **Move Down**.

Create a GlobalProtect Satellite Configuration (Continued)		
Step 6	Specify the certificates required to enable satellites to participate in the LSVPN.	1. 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. 2. Select the method of **Client Certificate** distribution: • **To store the client certificates on the portal**—select **Local** and select the Root CA certificate that the portal will use to issue client certificates to satellites upon successfully authenticating them from the **Issuing Certificate** drop-down.  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. • **To enable the portal to act as a SCEP client to dynamically request and issue client certificates**—select **SCEP** and then select the **SCEP** profile used to generate CSRs to your SCEP server.  If the you have not yet set up the portal to act as a SCEP client, you can add a **New** SCEP profile now. See Deploy Client Certificates to the GlobalProtect Satellites Using SCEP for details.
Step 7	Save the portal configuration.	1. Click **OK** to save the settings and close the GlobalProtect Portal Configuration dialog. 2. **Commit** your changes.

 © Palo Alto Networks, Inc.

Prepare the Satellite to Join the LSVPN

To participate in the LSVPN, the satellites require a minimal amount of configuration. Because the required configuration is minimal, you can pre-configure the satellites before shipping them to your branch offices for installation.

Prepare the Satellite 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. 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**. 2. In the **Interface Name** field, specify a numeric suffix, such as **.2**. 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*). 4. In the **Virtual Router** drop-down, select **default**. 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. 6. To save the interface configuration, click **OK**.
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. The root CA certificate is required to enable the satellite 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: a. Select **Device > Certificate Management > Certificates > Device Certificates**. b. Select the CA certificate, and click **Export**. 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.) 2. Import the root CA certificate you just exported onto each satellite as follows. a. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**. b. Enter a **Certificate Name** that identifies the certificate as your client CA certificate. c. **Browse** to the **Certificate File** you downloaded from the CA. d. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**. e. Select the certificate you just imported on the **Device Certificates** tab to open it. f. Select **Trusted Root CA** and then click **OK**.

Prepare the Satellite to Join the GlobalProtect LSVPN (Continued)		
Step 4	Configure the IPSec tunnel configuration.	1. Select **Network > IPSec Tunnels** and click **Add**. 2. On the **General** tab, enter a descriptive **Name** for the IPSec configuration. 3. Select the **Tunnel Interface** you created for the satellite. 4. Select **GlobalProtect Satellite** as the **Type**. 5. Enter the IP address or FQDN of the portal as the **Portal Address**. 6. Select the Layer 3 **Interface** you configured for the satellite. 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. 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 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**. 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: • Default routes • Routes within a virtual router other than the virtual router associated with the tunnel interface • Routes using the tunnel interface • Routes using the physical interface associated with the tunnel interface 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. 2. Click **Commit**.
Step 7	If required, provide the credentials to allow the satellite to authenticate to the portal. 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. 2. Click the **enter credentials** link in the **Portal Status** field and username and password required to authenticate the satellite to the portal. 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**.

 © Palo Alto Networks, Inc.

Verify the LSVPN Configuration

After configuring the portal, gateways, and satellites, 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.
Step 2	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.
Step 3	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.

LSVPN Quick Configs

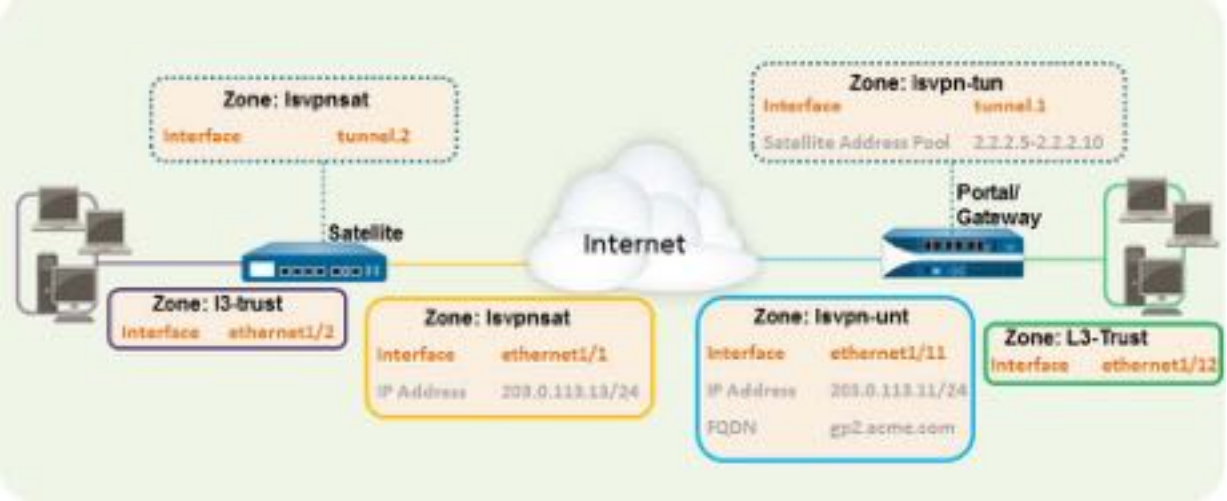
The following sections provide step-by-step instructions for configuring some common GlobalProtect LSVPN deployments:

- Basic LSVPN Configuration with Static Routing
- Advanced LSVPN Configuration with Dynamic Routing
- Advanced LSVPN Configuration with iBGP

 © 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. Satellites 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: • **Interface**—ethernet1/11 • **Security Zone**—lsvpn-tun • **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. 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: • **Interface**—tunnel.1 • **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).	See Create a Security Policy Rule.

Quick Config: Basic LSVPN with Static Routing (Continued)		
Step 4	Assign an SSL/TLS Service profile to the portal/gateway. The profile must reference a self-signed server certificate. 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 satellites. 2. Create SSL/TLS service profiles for the GlobalProtect portal and gateways. 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 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 one type of server profile on the portal: • Configure a RADIUS Server Profile. • Configure a TACACS+ Server Profile. • 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. • Configure a Kerberos Server Profile. 2. Configure 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: • **Interface**—ethernet1/11 • **IP Address**—203.0.113.11/24 • **SSL/TLS Server Profile**—lsvpnserver • **Certificate Profile**—lsvpn-profile • **Tunnel Interface**—tunnel.1 • **Primary DNS**/**Secondary DNS**—4.2.2.1/4.2.2.2 • **IP Pool**—2.2.2.111-2.2.2.120 • **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: • **Interface**—ethernet1/11 • **IP Address**—203.0.113.11/24 • **SSL/TLS Server Profile**—lsvpnserver • **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** 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: • **Gateway**—203.0.113.11 • **Issuing Certificate**—lsvpn-CA • **Trusted Root CA**—lsvpn-CA
Step 10 Prepare the Satellite to Join the LSVPN.	The satellite configuration in this example requires the following settings: Interface Configuration • Layer 3 interface—ethernet1/1, 203.0.113.13/24 • Tunnel interface—tunnel.2 • Zone—lsvpnsat Root CA Certificate from Portal • lsvpn-CA IPSec Tunnel Configuration • **Tunnel Interface**—tunnel.2 • **Portal Address**—203.0.113.11 • **Interface**—ethernet1/1 • **Local IP Address**—203.0.113.13/24 • **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.	Complete the following steps on each gateway and each satellite: 1. Select **Network > Interfaces > Tunnel** and select the tunnel configuration you created for the LSVPN to open the Tunnel Interface dialog. If you have not yet created the tunnel interface, see Step 2 in Quick Config: Basic LSVPN with Static Routing. 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. 3. Click **OK** to save the configuration.
Step 2	Configure the dynamic routing protocol on the gateway.	To configure OSPF on the gateway: 1. Select **Network > Virtual Routers** and select the virtual router associated with your VPN interfaces. 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. 3. If you are creating a new area, enter an **Area ID** on the **Type** tab. 4. On the **Interface** tab, click **Add** and select the tunnel **Interface** you created for the LSVPN. 5. Select **p2mp** as the **Link Type**. 6. Click **Add** in the Neighbors section and enter the IP address of the tunnel interface of each satellite, for example 2.2.2.111. 7. Click **OK** twice to save the virtual router configuration and then **Commit** the changes on the gateway. 8. Repeat this step each time you add a new satellite to the LSVPN.
Step 3	Configure the dynamic routing protocol on the satellite.	To configure OSPF on the satellite: 1. Select **Network > Virtual Routers** and select the virtual router associated with your VPN interfaces. 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. 3. If you are creating a new area, enter an **Area ID** on the **Type** tab. 4. On the **Interface** tab, click **Add** and select the tunnel **Interface** you created for the LSVPN. 5. Select **p2mp** as the **Link Type**. 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. 7. Click **OK** twice to save the virtual router configuration and then **Commit** the changes on the gateway. 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.

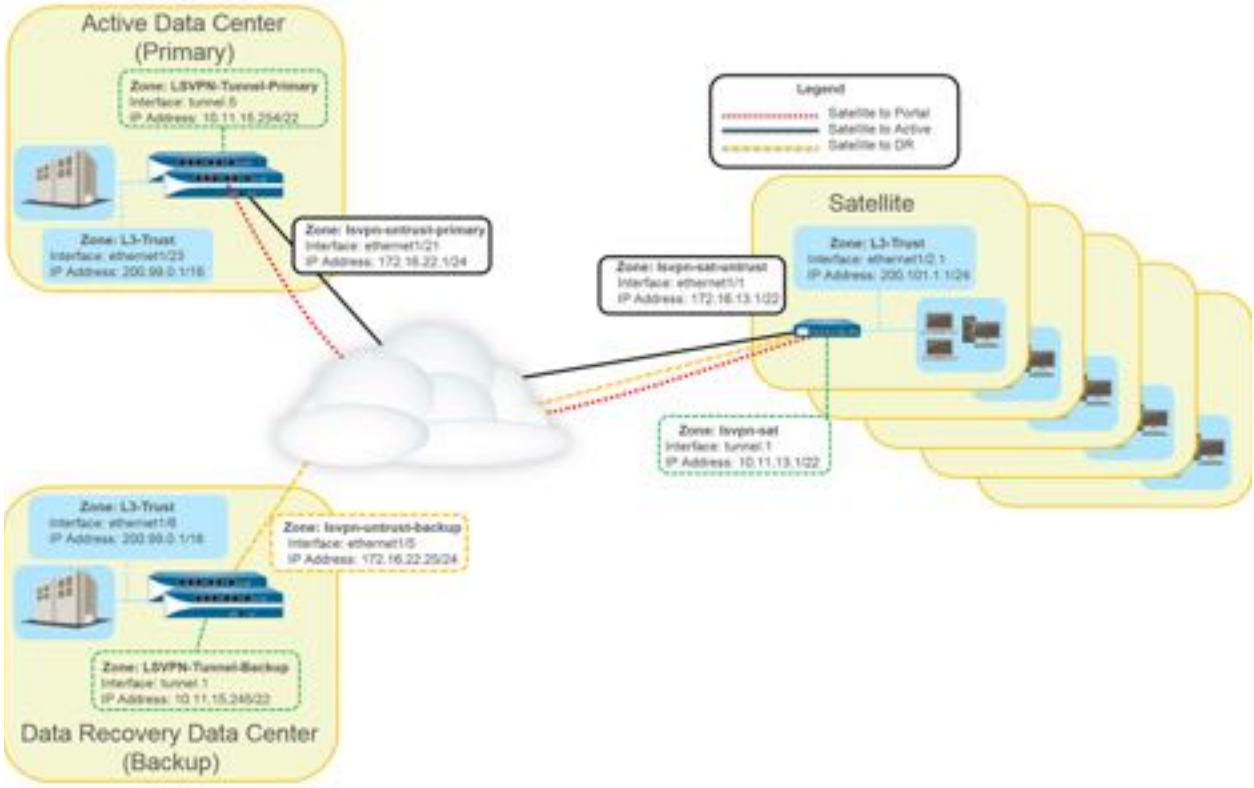
Advanced LSVPN Configuration with iBGP

This use case illustrates how GlobalProtect LSVPN securely connects distributed office locations with primary and disaster recovery data centers that house critical applications for users and how internal border gateway protocol (iBGP) eases deployment and upkeep. Using this method, you can extend up to 500 satellite offices connecting to a single gateway.

BGP is a highly scalable, dynamic routing protocol that is ideal for hub-and-spoke deployments such as LSVPN. As a dynamic routing protocol, it eliminates much of the overhead associated with access routes (static routes) by making it relatively easy to deploy additional satellite firewalls. Due to its route filtering capabilities and features such as multiple tunable timers, route dampening, and route refresh, BGP scales to a much higher number of routing prefixes with greater stability than other routing protocols like RIP and OSPF. In the case of iBGP, a peer group, which includes all the satellites and gateways in the LSVPN deployment, establishes adjacencies over the tunnel endpoints. The protocol then implicitly takes control of route advertisements, updates, and convergence.

In this example configuration, an active/passive HA pair of PA-5050 firewalls is deployed in the primary (active) data center and acts as the portal and primary gateway. The disaster recovery data center also has two PA-5050s in an active/passive HA pair acting as the backup LSVPN gateway. The portal and gateways serve 500 PA-200s deployed as LSVPN satellites in branch offices.

Both data center sites advertise routes but with different metrics. As a result, the satellites prefer and install the active data center's routes. However, the backup routes also exist in the local routing information base (RIB). If the active data center fails, the routes advertised by that data center are removed and replaced with routes from the disaster recovery data center's routes. The failover time depends on selection of iBGP times and routing convergence associated with iBGP.



The following workflow shows the steps for configuring this deployment:

Configure LSVPN with iBGP	

Step 1 Create Interfaces and Zones for the LSVPN.	Portal and Primary gateway:
	• **Zone**: LSVPN-Untrust-Primary
	  • Interface: ethernet1/21
	  • IPv4: 172.16.22.1/24
	• **Zone**: L3-Trust
	  • **Interface**: ethernet1/23
	  • **IPv4**: 200.99.0.1/16
	Backup gateway:
	• **Zone**: LSVPN-Untrust-Primary
	  • **Interface**: ethernet1/5
	  • **IPv4**: 172.16.22.25/24
	• **Zone**: L3-Trust
	  • **Interface**: ethernet1/6
	  • **IPv4**: 200.99.0.1/16
	Satellite:
	• **Zone**: LSVPN-Sat-Untrust
	  • **Interface**: ethernet1/1
	  • **IPv4**: 172.16.13.1/22
	• **Zone**: L3-Trust
	  • **Interface**: ethernet1/2.1
	  • **IPv4**: 200.101.1.1/24
	  Configure the zones, interfaces, and IP addresses on each satellite. The interface and local IP address will be different for each satellite. This interface is used for the VPN connection to the portal and gateway.

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.	Primary gateway: • **Interface**: tunnel.5 • **IPv4**: 10.11.15.254/22 • **Zone**: LSVPN-Tunnel-Primary Backup gateway: • **Interface**: tunnel.1 • **IPv4**: 10.11.15.245/22 • **Zone**: LSVPN-Tunnel-Backup

 © Palo Alto Networks, Inc.

Configure LSVPN with iBGP		
Step 3	Enable SSL Between GlobalProtect LSVPN Components. The gateway uses the self-signed root certificate authority (CA) to issue certificates for the satellites in a GlobalProtect LSVPN. Because one firewall houses the portal and primary gateway, a single certificate is used for authenticating to the satellites. The same CA is used to generate a certificate for the backup gateway. The CA generates certificates that pushed to the satellites from the portal and then used by the satellites to authenticate to the gateways.	You must also generate a certificate from the same CA for the backup gateway, allowing it to authenticate with the satellites. 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 is called CA-cert. 2. Create SSL/TLS service profiles for the GlobalProtect portal and gateways. Because the GlobalProtect portal and primary gateway are the same firewall interface, you can use the same server certificate for both components. • **Root CA Certificate**: CA-Cert • **Certificate Name**: LSVPN-Scale 3. Deploy the self-signed server certificates to the gateways. 4. Import the root CA certificate used to issue server certificates for the LSVPN components. 5. Create a certificate profile. 6. Repeat steps 2 through 5 on backup gateway with the following settings: • **Root CA Certificate**: CA-cert • **Certificate Name**: LSVPN-back-GW-cert
Step 4	Configure GlobalProtect Gateways for LSVPN.	1. Select **Network > GlobalProtect > Gateways** and click **Add**. 2. On the **General** tab, name the primary gateway `LSVPN-Scale`. 3. Under **Network Settings**, select `ethernet1/21` as the primary gateway interface and enter `172.16.22.1/24` as the IP address. 4. On the **Authentication** tab, select the LSVPN-Scale certificate created in Step 3. 5. Select **Satellite > Tunnel Settings** and select **Tunnel Configuration**. Set the **Tunnel Interface** to tunnel.5. All satellites in this use case connect to a single gateway, so a single satellite configuration is needed. Satellites are matched based on their serial numbers, so no satellites will need to authenticate as a user. 6. On **Satellite > Network Settings**, define the pool of IP address to assign to the tunnel interface on the satellite once the VPN connection is established. Because this use case uses dynamic routing, the Access Routes setting remains blank. 7. Repeat steps 1 through 5 on the backup gateway with the following settings: • **Name**: LSVPN-backup • **Gateway interface**: ethernet1/5 • **Gateway IP**: 172.16.22.25/24 • **Server cert**: LSVPN-backup-GW-cert • **Tunnel interface**: tunnel.1

Configure LSVPN with iBGP	
Step 5 Configure iBGP on the primary and backup gateways and add a redistribution profile to allow the satellites to inject local routes back to the gateways. Each satellite office manages its own network and firewall, so the redistribution profile called ToAllSat is configured to redistribute local routes back to the GlobalProtect gateway.	1. Select **Network > Virtual Routers** and **Add** a virtual router. 2. On **Router Settings**, add the **Name** and **Interface** for the virtual router. 3. On **Redistribution Profile** and select **Add**. a. Name the redistribution profile `ToAllSat` and set the **Priority** to 1. b. Set Redistribute to **Redist**. c. **Add** `ethernet1/23` from the Interface drop-down. d. Click **OK**. 4. Select **BGP** on the Virtual Router to configure BGP. a. On **BGP > General**, select **Enable**. b. Enter the gateway IP address as the **Router ID** (`172.16.22.1`) and `1000` as the **AS Number**. c. In the Options section, select **Install Route**. d. On **BGP > Peer Group**, click **Add** a peer group with all the satellites that will connect to the gateway. e. On **BGP > Redist Rules**, **Add** the `ToAllSat` redistribution profile you created previously. 5. Click **OK**. 6. Repeat steps 1 through 5 on the backup gateway using `ethernet1/6` for the redistribution profile.
Step 6 Prepare the Satellite to Join the LSVPN. The configuration shown is a sample of a single satellite. Repeat this configuration each time you add a new satellite to the LSVPN deployment.	1. Configure a tunnel interface as the tunnel endpoint for the VPN connection to the gateways. 2. Set the IPSec tunnel type to GlobalProtect Satellite and enter the IP address of the GlobalProtect Portal. 3. Select **Network > Virtual Routers** and **Add** a virtual router. 4. On **Router Settings**, add the **Name** and **Interface** for the virtual router. 5. Select **Virtual Router > Redistribution Profile** and **Add** a profile with the following settings. a. Name the redistribution profile `ToLSVPNGW` and set the **Priority** to 1. b. **Add** an **Interface** `ethernet1/2.1`. c. Click **OK**. 6. Select **BGP > General**, **Enable** BGP and configure the protocol as follows: a. Enter the gateway IP address as the **Router ID** (`172.16.22.1`) and `1000` as the **AS Number**. b. In the Options section, select **Install Route**. c. On **BGP > Peer Group**, **Add** a peer group containing all the satellites that will connect to the gateway. d. On **BGP > Redist Rules**, **Add** the `ToLSVPNGW` redistribution profile you created previously. 7. Click **OK**.

 © Palo Alto Networks, Inc.

Configure LSVPN with iBGP		
Step 7	Configure the GlobalProtect Portal for LSVPN. Both data centers advertise their routes but with different routing priorities to ensure that the active data center is the preferred gateway.	1. Select **Network > GlobalProtect > Portals** and click **Add**. 2. On **General**, enter `LSVPN-Portal` as the portal name. 3. On **Network Settings**, select `ethernet1/21` as the **Interface** and select `172.16.22.1/24` as the **IP Address**. 4. On the **Authentication** tab, select the previously created primary gateway SSL/TLS Profile `LSVPN-Scale` from the **SSL/TLS Service Profile** drop-down menu. 5. On the **Satellite** tab, **Add** a satellite and **Name** it `sat-config-1`. 6. Set the **Configuration Refresh Interval** to `12`. 7. On **GlobalProtect Satellite > Devices**, add the serial number and hostname of each satellite device in the LSVPN. 8. On **GlobalProtect Satellite > Gateways**, add the name and IP address of each gateway. Set the routing priority of the primary gateway to 1 and the backup gateway to 10 to ensure that the active data center is the preferred gateway.
Step 8	Verify the LSVPN Configuration.	
Step 9	(Optional) Add a new site to the LSVPN deployment.	1. Select **Network > GlobalProtect > Portals > GlobalProtect Portal> Satellite Configuration > GlobalProtect Satellite > Devices** to add the serial number of the new satellite to the GlobalProtect portal. 2. Configure the IPSec tunnel on the satellite with the GlobalProtect Portal IP address. 3. Select **Network > Virtual Router > BGP > Peer Group** to add the satellite to the BGP Peer Group configuration on each gateway. 4. Select **Network > Virtual Router > BGP > Peer Group** to add the gateways to the BGP Peer Group configuration on the new satellite.

 © 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
- Configure an Aggregate Interface Group
- Use Interface Management Profiles to Restrict Access
- Virtual Routers
- Static Routes
- RIP
- OSPF
- BGP
- Session Settings and Timeouts
- DHCP
- DNS
- NAT
- NPTv6
- ECMP
- LLDP
- BFD

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 supported deployments.

▲ Virtual Wire Deployments

▲ Layer 2 Deployments

▲ Layer 3 Deployments

▲ Tap Mode Deployments

Virtual Wire Deployments

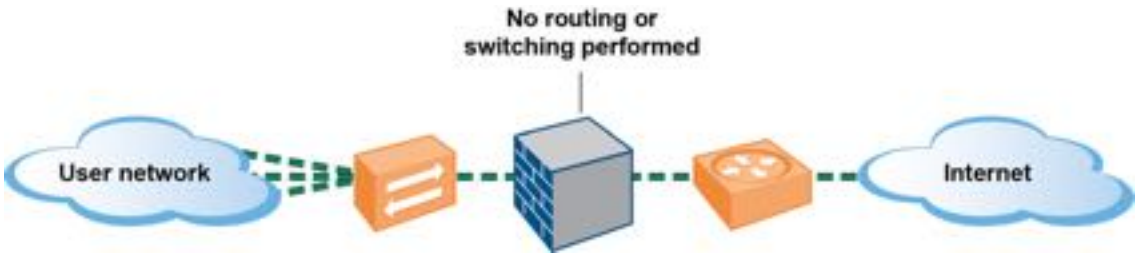
In a virtual wire deployment, the firewall is installed transparently on a network segment by binding two ports together and should be used only when no switching or routing is needed.

A virtual wire deployment:

- 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:

- **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
• Configure two Ethernet interfaces as type virtual wire, and assign these interfaces to a virtual wire.
• 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.
• 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. 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)

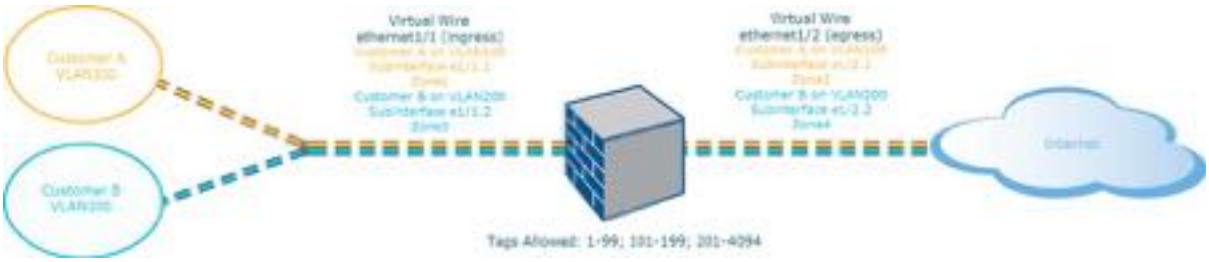


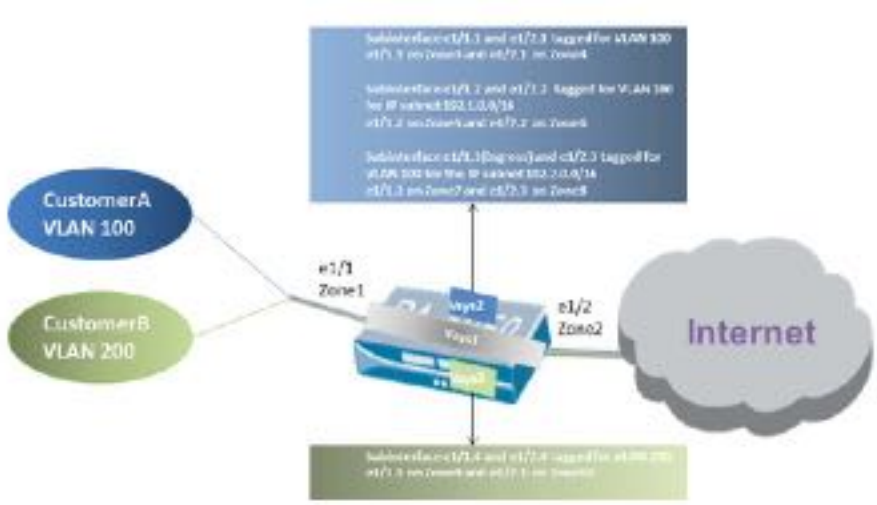
Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags only) depicts CustomerA and CustomerB connected to the firewall through one physical interface, ethernet1/1, configured as a Virtual Wire; it is the ingress interface. A second physical interface, ethernet1/2, is also part of the Virtual Wire; 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) e1/2.1 (egress)	Zone3 Zone4	100 100	None
	2	e1/1.2 (ingress) e1/2.2 (egress)	Zone5 Zone6	100 100	IP subnet 192.1.0.0/16
	2	e1/1.3 (ingress) e1/2.3 (egress)	Zone7 Zone8	100 100	IP subnet 192.2.0.0/16
B	3	e1/1.4 (ingress) e1/2.4 (egress)	Zone9 Zone10	200 200	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.

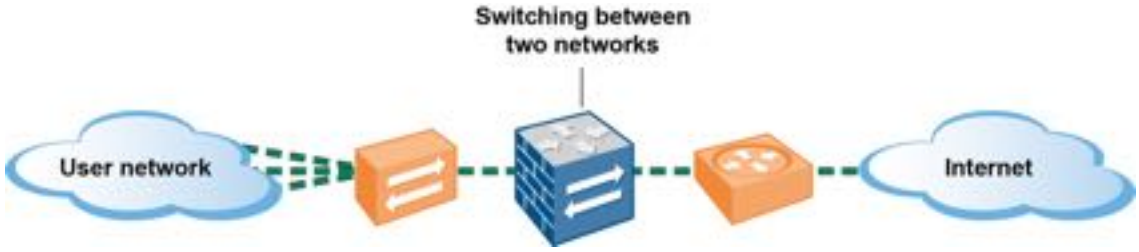
For return-path traffic, the firewall compares the destination IP address as defined in the IP classifier on the customer-facing subinterface and selects the appropriate virtual wire to route traffic through the accurate subinterface.

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. You must assign a group of interfaces to a VLAN object in order for the firewall to switch between them. The firewall performs 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



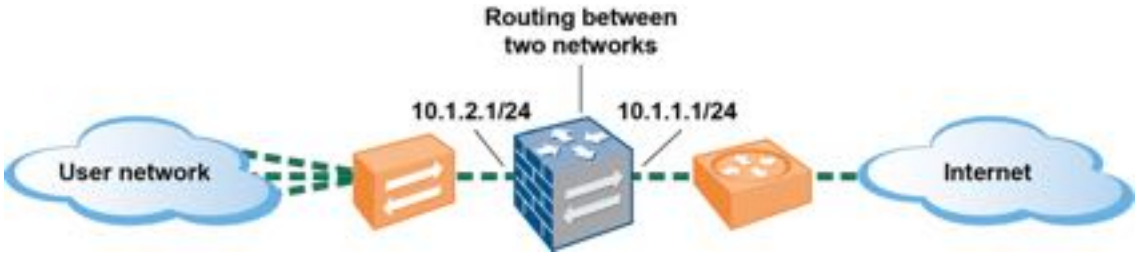
In a Layer 2 deployment, the firewall rewrites the inbound Port VLAN ID (PVID) number in a Cisco per-VLAN spanning tree (PVST+) or Rapid PVST+ bridge protocol data unit (BPDU) to the proper outbound VLAN ID number and forwards it out. The firewall rewrites such BPDUs on Layer 2 Ethernet and Aggregated Ethernet (AE) interfaces only.

A Cisco switch must have the loopguard disabled for the PVST+ or Rapid PVST+ BPDU rewrite to function properly on the firewall.

Layer 3 Deployments

In a Layer 3 deployment, the firewall routes traffic between multiple ports. This deployment requires that you assign an IP address to each interface and configure Virtual Routers to route the traffic. Choose this option when routing is required.

Figure: Layer 3 Deployment



The following Layer 3 interface deployments are also supported:

- Point-to-Point Protocol over Ethernet Support
- DHCP Client

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.

 © Palo Alto Networks, Inc.

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.

Configure an Aggregate Interface Group

An aggregate interface group uses IEEE 802.1AX link aggregation to combine multiple Ethernet interfaces into a single virtual interface that connects the firewall to another network device or another firewall. An aggregate group increases the bandwidth between peers by load balancing traffic across the combined interfaces. It also provides redundancy; when one interface fails, the remaining interfaces continue supporting traffic.

By default, interface failure detection is automatic only at the physical layer between directly connected peers. However, if you enable Link Aggregation Control Protocol (LACP), failure detection is automatic at the physical and data link layers regardless of whether the peers are directly connected. LACP also enables automatic failover to standby interfaces if you configured hot spares. All Palo Alto Networks firewalls except the PA-200 and VM-Series platforms support aggregate groups. You can add up to eight aggregate groups per firewall and each group can have up to eight interfaces.

Before configuring an aggregate group, you must configure its interfaces. All the interfaces in an aggregate group must be the same with respect to bandwidth and interface type. The options are:

- **Bandwidth**—1Gbps or 10Gbps
- **Interface type**—HA3, virtual wire, Layer 2, or Layer 3. You can aggregate the HA3 (packet forwarding) interfaces in an active/active high availability (HA) deployment but only for PA-500, PA-3000 Series, PA-4000 Series, and PA-5000 Series firewalls.



This procedure describes configuration steps only for the Palo Alto Networks firewall. You must also configure the aggregate group on the peer device. Refer to the documentation of that device for instructions.

Configure an Aggregate Interface Group	
Step 1 Configure the general interface group parameters.	1. Select **Network > Interfaces > Ethernet** and **Add Aggregate Group**. 2. In the field adjacent to the read-only **Interface Name**, enter a number (1–8) to identify the aggregate group. 3. For the **Interface Type**, select **HA**, **Virtual Wire**, **Layer2**, or **Layer3**. 4. Configure the remaining parameters for the **Interface Type** you selected.

 © Palo Alto Networks, Inc.

Configure an Aggregate Interface Group (Continued)	
Step 2 Configure the LACP settings. Perform this step only if you want to enable LACP for the aggregate group. You cannot enable LACP for virtual wire interfaces.	1. Select the **LACP** tab and **Enable LACP**. 2. Set the **Mode** for LACP status queries to **Passive** (the firewall just responds—the default) or **Active** (the firewall queries peer devices). As a best practice, set one LACP peer to active and the other to passive. LACP cannot function if both peers are passive. The firewall cannot detect the mode of its peer device. 3. Set the **Transmission Rate** for LACP query and response exchanges to **Slow** (every 30 seconds—the default) or **Fast** (every second). Base your selection on how much LACP processing your network supports and how quickly LACP peers must detect and resolve interface failures. 4. Select **Fast Failover** if you want to enable failover to a standby interface in less than one second. By default, the option is disabled and the firewall uses the IEEE 802.1ax standard for failover processing, which takes at least three seconds. As a best practice, use **Fast Failover** in deployments where you might lose critical data during the standard failover interval. 5. Enter the **Max Ports** (number of interfaces) that are active (1–8) in the aggregate group. If the number of interfaces you assign to the group exceeds the **Max Ports**, the remaining interfaces will be in standby mode. The firewall uses the **LACP Port Priority** of each interface you assign (Step 3) to determine which interfaces are initially active and to determine the order in which standby interfaces become active upon failover. If the LACP peers have non-matching port priority values, the values of the peer with the lower **System Priority** number (default is 32,768; range is 1–65,535) will override the other peer. 6. (Optional) For active/passive firewalls only, select **Enable in HA Passive State** if you want to enable LACP pre-negotiation for the passive firewall. LACP pre-negotiation enables quicker failover to the passive firewall (for details, see LACP and LLDP Pre-Negotiation for Active/Passive HA). If you select this option, you cannot select **Same System MAC Address for Active-Passive HA**; pre-negotiation requires unique interface MAC addresses on each HA firewall. 7. (Optional) For active/passive firewalls only, select **Same System MAC Address for Active-Passive HA** and specify a single **MAC Address** for both HA firewalls. This option minimizes failover latency if the LACP peers are virtualized (appearing to the network as a single device). By default, the option is disabled: each firewall in an HA pair has a unique MAC address. If the LACP peers are not virtualized, use unique MAC addresses to minimize failover latency.

Configure an Aggregate Interface Group (Continued)		
Step 3	Assign interfaces to the aggregate group.	Perform the following steps for each interface (1–8) that will be a member of the aggregate group. 1. Select **Network > Interfaces > Ethernet** and click the interface name to edit it. 2. Set the **Interface Type** to **Aggregate Ethernet**. 3. Select the **Aggregate Group** you just defined. 4. Select the **Link Speed**, **Link Duplex**, and **Link State**. As a best practice, set the same link speed and duplex values for every interface in the group. For non-matching values, the firewall defaults to the higher speed and full duplex. 5. (Optional) Enter an **LACP Port Priority** (default is 32,768; range is 1–65,535) if you enabled LACP for the aggregate group. If the number of interfaces you assign exceeds the **Max Ports** value of the group, the port priorities determine which interfaces are active or standby. The interfaces with the lower numeric values (higher priorities) will be active. 6. Click **OK**.
Step 4	If the firewalls have an active/active configuration and you are aggregating HA3 interfaces, enable packet forwarding for the aggregate group.	1. Select **Device > High Availability > Active/Active Config** and edit the Packet Forwarding section. 2. Select the aggregate group you configured for the **HA3 Interface** and click **OK**.
Step 5	Commit your changes and verify the aggregate group status.	1. Click **Commit**. 2. Select **Network > Interfaces > Ethernet**. 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. 4. If you configured LACP, verify that the Features column displays the LACP enabled icon for the aggregate group.

 © Palo Alto Networks, Inc.

Use Interface Management Profiles to Restrict Access

An Interface Management profile protects the firewall from unauthorized access by defining the protocols, services, and IP addresses that a firewall interface permits for management traffic. For example, you might want to prevent users from accessing the firewall web interface over the ethernet1/1 interface but allow that interface to receive SNMP queries from your network monitoring system. In this case, you would enable SNMP and disable HTTP/HTTPS in an Interface Management profile and assign the profile to ethernet1/1.

You can assign an Interface Management profile to Layer 3 Ethernet interfaces (including subinterfaces) and to logical interfaces (aggregate group, VLAN, loopback, and tunnel interfaces). If you do not assign an Interface Management profile to an interface, it denies access for all IP addresses, protocols, and services by default.



The management (MGT) interface does not require an Interface Management profile. You restrict protocols, services, and IP addresses for the MGT interface when you Perform Initial Configuration of the firewall. In case the MGT interface goes down, allowing management access over another interface enables you to continue managing the firewall. However, as a best practice, use additional methods besides Interface Management profiles to prevent unauthorized access over that interface. These methods include role-based access control and access restrictions based on VLANs, virtual routers, or virtual systems.

Configure and Assign an Interface Management Profile	
Step 1 Configure the Interface Management profile.	1. Select **Network > Network Profiles > Interface Mgmt** and click **Add**. 2. Select the protocols that the interface permits for management traffic: **Ping**, **Telnet**, **SSH**, **HTTP**, **HTTP OCSP**, **HTTPS**, or **SNMP**. 3. Select the services that the interface permits for management traffic: • **Response Pages**—Use to enable response pages for: – **Captive Portal**—To serve Captive Portal response pages, the firewall leaves ports open on Layer 3 interfaces: port 6080 for NT LAN Manager (NTLM), 6081 for Captive Portal in transparent mode, and 6082 for Captive Portal in redirect mode. For details, see Configure Captive Portal. – **URL Admin Override**—For details, see Configure URL Admin Override. • **User-ID**—Use to Configure Firewalls to Redistribute User Mapping Information. • **User-ID Syslog Listener-SSL** or **User-ID Syslog Listener-UDP**—Use to Configure User-ID to Receive User Mappings from a Syslog Sender over SSL or UDP. 4. (Optional) **Add** the Permitted IP Addresses that can access the interface. If you don't add entries to the list, the interface has no IP address restrictions. 5. Click **OK**.

Configure and Assign an Interface Management Profile (Continued)		
Step 2	Assign the Interface Management profile to an interface.	1. Select **Network > Interfaces**, select the type of interface (**Ethernet**, **VLAN**, **Loopback**, or **Tunnel**), and select the interface. 2. Select **Advanced > Other info** and select the Interface **Management Profile** you just added. 3. Click **OK** and **Commit**.

 © 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 virtual router on the firewall supports the following Layer 3:

- 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 • Administrative distances for static, OSPF internal, OSPF external, IBGP, EBGP and RIP
Step 2	Create the virtual router and name it.	1. Select **Network > Virtual Routers**. 2. Click **Add** and enter a name for the virtual router. 3. Select interfaces to apply to the virtual router. 4. Click **OK**.
Step 3	Select interfaces to apply to the virtual router.	1. Click **Add** in the Interfaces box. 2. Select an already defined interface from the drop-down. Repeat this step for all interfaces that you want to add to the virtual router.

Define a Virtual Router General Configuration (Continued)		
Step 4	Set Administrative Distances for static and dynamic routing.	Set Administrative Distances as required. • **Static**—Range is 10-240; default is 10. • **OSPF Internal**—Range is 10-240; default is 30. • **OSPF External**—Range is 10-240; default is 110. • **IBGP**—Range is 10-240; default is 200. • **EBGP**—Range is 10-240; default is 20. • **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 firewall can 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. 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). 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). 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. 2. Select the **Interface Type**. Although your choice here depends on your network topology, this example shows the steps for **Layer3**. 3. In the **Virtual Router** drop-down, select **default**. 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**. 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. 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**. 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.  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. 2. Select **Layer3** from the **Interface Type** drop-down. 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**. 4. Select the same Virtual Router you used in the previous step, default in this example. 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. 6. To enable you to ping the interface, select the management profile that you created in Step 2-6. 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. 2. Select **Layer3** from the **Interface Type** drop-down. In this example, we are configuring Ethernet1/13 as the DMZ interface. 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**. 4. Select the Virtual Router you used in Step 2, default in this example. 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. 6. To enable you to ping the interface, select the management profile that you created in Step 2-6. 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**.

 © 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. 2. Select **Enable** to enable the RIP protocol. 3. Select **Reject Default Route** if you do not want to learn any default routes through RIP. This is the recommended default setting. 4. De-select **Reject Default Route** if you want to permit redistribution of default routes through RIP.
Step 3	Configure interfaces for the RIP protocol.	1. On the **Interfaces** tab, select an interface from the drop-down in the Interface configuration section. 2. Select an already defined interface. 3. Select **Enable**. 4. Select **Advertise** to advertise a default route to RIP peers with the specified metric value. 5. (Optional) Select a profile from the **Auth Profile** drop-down. 6. Select normal, passive or send-only from the **Mode** drop-down. 7. Click **OK**.
Step 4	Configure RIP timers.	1. On the **Timers** tab, enter a value for **Interval Seconds (sec)**. This setting defines the length of the following RIP timer intervals in seconds (range is 1-60; default is 1). 2. Specify the **Update Intervals** to define the number of intervals between route update announcements (range is 1-3600; default is 30). 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). 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).

Configure RIP (Continued)	
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 MD5 authentication. MD5 authentication is recommended; it is more secure than a simple password. Simple Password RIP authentication 1. Select **Auth Profiles** and click **Add**. 2. Enter a name for the authentication profile to authenticate RIP messages. 3. Select **Simple Password** as the **Password Type**. 4. Enter a simple password and then confirm. MD5 RIP authentication 1. Select **Auth Profiles** and click **Add**. 2. Enter a name for the authentication profile to authenticate RIP messages. 3. Select **MD5** as the **Password Type**. 4. Click **Add**. 5. Enter one or more password entries, including: • Key-ID (range is 0-255) • Key 6. (Optional) Select **Preferred** status. 7. Click **OK** to specify the key to be used to authenticate outgoing message. 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

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 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.

OSPF 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.

 © Palo Alto Networks, Inc.

OSPF Area Type	Description
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.

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	Enable OSPF.	1. Select the **OSPF** tab. 2. Select **Enable** to enable the OSPF protocol. 3. (Optional) Enter the **Router ID**. 4. Select **Reject Default Route** if you do not want to learn any default routes through OSPF. This is the recommended default setting. De-select **Reject Default Route** if you want to permit redistribution of default routes through OSPF.

Configure OSPF (Continued)		
Step 3	Configure Areas - Type for the OSPF protocol.	1. On the **Areas** tab, click **Add**. 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. 3. On the **Type** tab, select one of the following from the area **Type** drop-down: • **Normal**—There are no restrictions; the area can carry all types of routes. • **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: – **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. – **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. • **NSSA** (Not-So-Stubby Area)—The firewall can leave the area only 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: – **Type**—Select either **Ext 1** or **Ext 2** route type to advertise the default LSA. – **Ext Ranges**—Click **Add** in the section to enter ranges of external routes that you want to enable or suppress advertising for. 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. • **Auth Profile**—Select a previously-defined authentication profile. • **Timing**—It is recommended that you keep the default timing settings. • **Neighbors**—For p2pmp interfaces, enter the neighbor IP address for all neighbors that are reachable through this interface. 5. Select **normal**, **passive** or **send-only** as the **Mode**. 6. Click **OK**.
Step 4	Configure Areas - Range for the OSPF protocol	1. On the **Range** tab, click **Add** to aggregate LSA destination addresses in the area into subnets. 2. **Advertise** or **Suppress** advertising LSAs that match the subnet, and click **OK**. Repeat to add additional ranges.

 © Palo Alto Networks, Inc.

Configure OSPF (Continued)		
Step 5	Configure Areas - Interfaces for the OSPF protocol	1. On the **Interface** tab, click **Add** and enter the following information for each interface to be included in the area: • **Interface**—Select an interface from the drop-down. • **Enable**—Selecting this option causes the OSPF interface settings to take effect. • **Passive**—Select 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. • **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. • **Metric**—Enter an OSPF metric for this interface (range is 0-65535; default is 10). • **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. • **Auth Profile**—Select a previously-defined authentication profile. • **Timing**—Modify the timing settings if desired (not recommended). For details on these settings, refer to the online help. • If **p2mp** is selected for **Link Type** interfaces, enter the neighbor IP addresses for all neighbors that are reachable through this interface. 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: • **Name**—Enter a name for the virtual link. • **Neighbor ID**—Enter the router ID of the router (neighbor) on the other side of the virtual link. • **Transit Area**—Enter the area ID of the transit area that physically contains the virtual link. • **Enable**—Select to enable the virtual link. • **Timing**—It is recommended that you keep the default timing settings. • **Auth Profile**—Select a previously-defined authentication profile. 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. MD5 authentication is recommended; it is more secure than a simple password. Simple Password OSPF authentication 1. On the **Auth Profiles** tab, click **Add**. 2. Enter a name for the authentication profile to authenticate OSPF messages. 3. Select **Simple Password** as the **Password Type**. 4. Enter a simple password and then confirm. MD5 OSPF authentication 1. On the **Auth Profiles** tab, click **Add**. 2. Enter a name for the authentication profile to authenticate OSPF messages. 3. Select **MD5** as the **Password Type**. 4. Click **Add**. 5. Enter one or more password entries, including: • Key-ID (range is 0-255) • Key • Select the **Preferred** option to specify that the key be used to authenticate outgoing messages. 6. Click **OK**. 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 **RFC 1583 Compatibility** to ensure compatibility with RFC 1583. 2. Configure a value for the **SPF Calculation Delay** (sec) timer. 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. 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

OSPFv3 supports both IPv4 and IPv6. You must use OSPFv3 if you are using IPv6.

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. 2. Select **Enable** to enable the OSPF protocol. 3. Select **Reject Default Route** if you do not want to learn any default routes through OSPF. This is the recommended default setting. 4. Clear **Reject Default Route** if you want to permit redistribution of default routes through OSPF.
Step 3	Configure general OSPFv3 configuration settings.	1. Select the **OSPFv3** tab. 2. Select **Enable** to enable the OSPF protocol. 3. Select **Reject Default Route** if you do not want to learn any default routes through OSPFv3 This is the recommended default setting. Clear **Reject Default Route** if you want to permit redistribution of default routes through OSPFv3.

Configure OSPFv3 (Continued)	
Step 4 Configure Auth Profile for the OSPFv3 protocol. 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) (which is recommended) or IPv6 Authentication Header (AH). **ESP OSPFv3 authentication** 1. On the **Auth Profiles** tab, click **Add**. 2. Enter a name for the authentication profile to authenticate OSPFv3 messages. 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. 4. Select **ESP** for **Protocol**. 5. Select a **Crypto Algorithm** from the drop-down. You can enter none or one of the following algorithms: SHA1, SHA256, SHA384, SHA512 or MD5. 6. If a **Crypto Algorithm** other than none was selected, enter a value for **Key** and then confirm. **AH OSPFv3 authentication** 1. On the **Auth Profiles** tab, click **Add**. 2. Enter a name for the authentication profile to authenticate OSPFv3 messages. 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. 4. Select **AH** for **Protocol**. 5. Select a **Crypto Algorithm** from the drop-down. You must enter one of the following algorithms: SHA1, SHA256, SHA384, SHA512 or MD5. 6. Enter a value for **Key** and then confirm. 7. Click **OK**. 8. Click **OK** again in the Virtual Router - OSPF Auth Profile dialog.

 © Palo Alto Networks, Inc.

Configure OSPFv3 (Continued)		
Step 5	Configure Areas - Type for the OSPF protocol.	1. On the **Areas** tab, click **Add**. 2. Enter an Area ID. This is the identifier that each neighbor must accept to be part of the same area. 3. On the **General** tab, select one of the following from the area **Type** drop-down: • **Normal**—There are no restrictions; the area can carry all types of routes. • **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: – **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. – **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. • **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: – **Type**—Select either **Ext 1** or **Ext 2** route type to advertise the default LSA. – **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 1. On the **Areas** tab, select an existing area from the table. 2. On the **General** tab, select a previously defined **Authentication Profile** from the **Authentication** drop-down. 3. Click **OK**. To an Interface 1. On the **Areas** tab, select an existing area from the table. 2. Select the **Interface** tab and click **Add**. 3. Select the authentication profile you want to associate with the OSPF interface from the **Auth Profile** drop-down.

Configure OSPFv3 (Continued)	
Step 7 (Optional) Configure Export Rules	1. On the **Export** tab, click **Add**. 2. Select **Allow Redistribute Default Route** to permit redistribution of default routes through OSPFv3. 3. Select the name of a redistribution profile. The value must be an IP subnet or valid redistribution profile name. 4. Select a metric to apply for **New Path Type**. 5. Specify a **New Tag** for the matched route that has a 32-bit value. 6. Assign a metric for the new rule (range is 1 - 65535). 7. Click **OK**.
Step 8 Configure Advanced OSPFv3 options.	1. On the **Advanced** tab, select **Disable Transit Routing for SPF Calculation** if you want the firewall to participate in OSPF topology distribution without being used to forward transit traffic. 2. Configure a value for the **SPF Calculation Delay** (sec) timer. 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. 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. 4. (Optional) 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.

 © Palo Alto Networks, Inc.

- **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 OSPF 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 are selected (they are enabled by default): • **Enable Graceful Restart** • **Enable Helper Mode** • **Enable Strict LSA checking** These 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 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`

If you are using the web interface to view the routing table, use the following workflow:

View the Routing Table
1. Select **Network > Virtual Routers** and in the same row as the virtual router you are interested in, click the **More Runtime Stats** link.

View the Routing Table
2. Select **Routing > Route Table** and examine the **Flags** column of the routing table for routes that were learned by OSPF.

Confirm OSPF Adjacencies

Use the following workflow to confirm that OSPF adjacencies have been established:

View the Neighbor Tab to Confirm OSPF Adjacencies
1. Select **Network > Virtual Routers** and in the same row as the virtual router you are interested in, click the **More Runtime Stats** link.
2. Select **OSPF > Neighbor** and examine the **Status** column to determine if OSPF adjacencies have been established.

Confirm that OSPF Connections are Established

View the System log to confirm that the firewall has established OSPF connections.

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 (are full).

 © Palo Alto Networks, Inc.

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. Authentication helps prevent route leaking and successful DoS attacks.
- 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.
Step 2	Configure standard BGP configuration settings.	1. Select the **BGP** tab. 2. Select **Enable** to enable the BGP protocol. 3. For **Router ID**, assign an IP address to the virtual router. 4. For **AS Number**, enter the number of the AS to which the virtual router belongs, based on the router ID. Range is 1-4294967295.

Configure BGP (Continued)	
Step 3 Configure general BGP configuration settings.	1. Select **BGP > General**. 2. Select **Reject Default Route** to ignore any default routes that are advertised by BGP peers. 3. Select **Install Route** to install BGP routes in the global routing table. 4. Select **Aggregate MED** to enable route aggregation even when routes have different Multi-Exit Discriminator (MED) values. 5. Enter a value for the **Default Local Preference** that specifies a value than can be used to determine preferences among different paths. 6. Select one of the following values for the AS format for interoperability purposes: • 2 Byte (default value) • 4 Byte 7. Enable or disable each of the following values for **Path Selection**: • **Always Compare MED**—Enable this comparison to choose paths from neighbors in different autonomous systems. • **Deterministic MED Comparison**—Enable this comparison to choose between routes that are advertised by IBGP peers (BGP peers in the same autonomous system). 8. Click **Add** to include a new authentication profile and configure the following settings: • **Profile Name**—Enter a name to identify the profile. • **Secret/Confirm Secret**—Enter and confirm a passphrase for BGP peer communications. The Secret is used as a key in MD5 authentication.

 © Palo Alto Networks, Inc.

Configure BGP (Continued)	
Step 4 (Optional) Configure BGP Advanced settings.	1. On the **Advanced** tab, select **Graceful Restart** and configure the following timers: • **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). • **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). • **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). 2. Specify an IPv4 identifier to represent the reflector cluster in the **Reflector Cluster ID** box. 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. 4. Click **Add** and enter the following information for each Dampening Profile that you want to configure, select **Enable**, and click **OK**: • **Profile Name**—Enter a name to identify the profile. • **Cutoff**—Specify a route withdrawal threshold above which a route advertisement is suppressed (range is 0.0-1000.0; default is 1.25). • **Reuse**—Specify a route withdrawal threshold below which a suppressed route is used again (range is 0.0-1000.0; default is 5). • **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). • **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). • **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). 5. Click **OK**.

Configure BGP (Continued)		
Step 5	Configure the BGP peer group.	1. Select the **Peer Group** tab and click **Add**. 2. Enter a **Name** for the peer group and select **Enable**. 3. Select **Aggregated Confed AS Path** to include a path to the configured aggregated confederation AS. 4. Select **Soft Reset with Stored Info** to perform a soft reset of the firewall after updating the peer settings. 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). • **IBGP**—**Export Next Hop:** Specify **Original** or **Use self** • **EBGP Confed**—**Export Next Hop:** Specify **Original** or **Use self** • **EBGP Confed**—**Export Next Hop:** Specify **Original** or **Use self** • **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. 6. Click **OK** to save.
Step 6	Configure Import and Export rules. 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 **Enable**. 2. Click **Add** and select the **Peer Group** to which the routes will be imported from. 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. 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. 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. 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 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. 2. Select **Enable**. 3. Click **Add** and in the **Used By** section enter the peer group(s) that will use the conditional advertisement policy. 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. 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 summarize routes in the BGP configuration. 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. 2. In the **Prefix** field, enter the network prefix that will be the primary prefix for the aggregated prefixes. 3. Select the **Suppress Filters** tab and define the attributes that will cause the matched routes to be suppressed. 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. 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**. 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. 3. Click **Enable** to enable the rule. 4. In the **Metric** field, enter the route metric that will be used for the rule. 5. In the **Set Origin** drop-down, select **incomplete**, **igp**, or **egp**. 6. (Optional) Set MED, local preference, AS path limit and community values.

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 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

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).

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

 © Palo Alto Networks, Inc.

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.

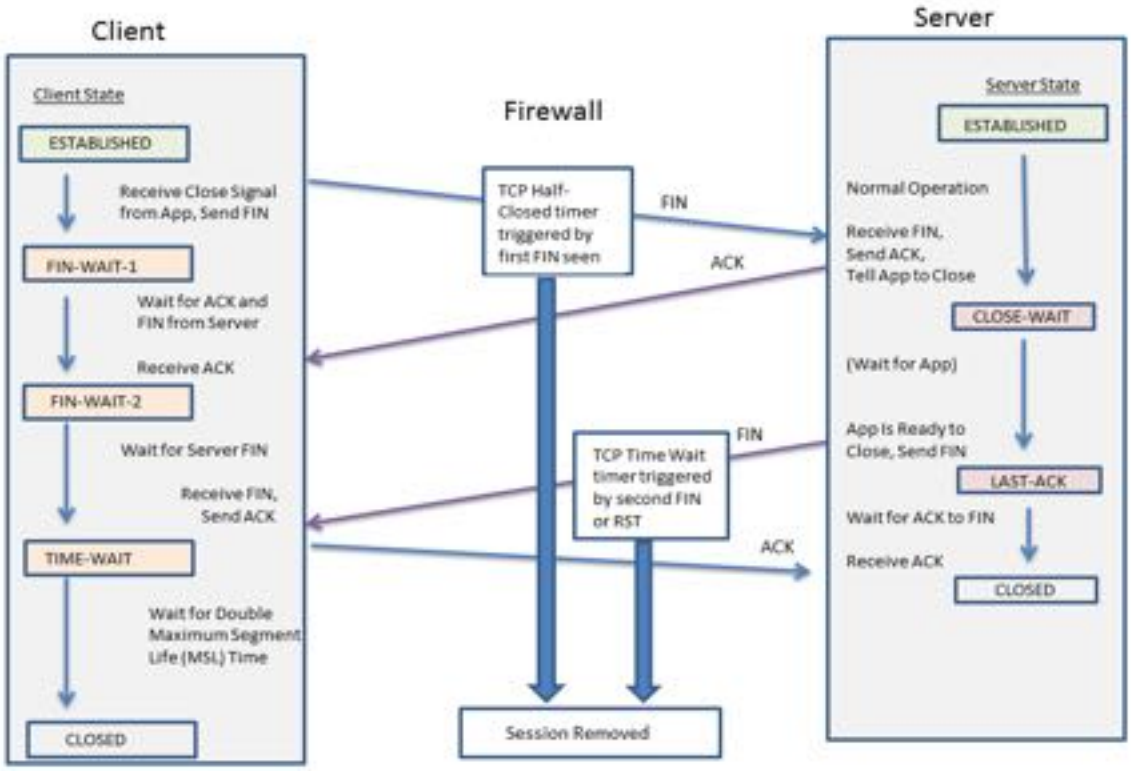
- TCP Half Closed and TCP Time Wait Timers
- Unverified RST Timer
- TCP Split Handshake Drop
- Maximum Segment Size (MSS)

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.

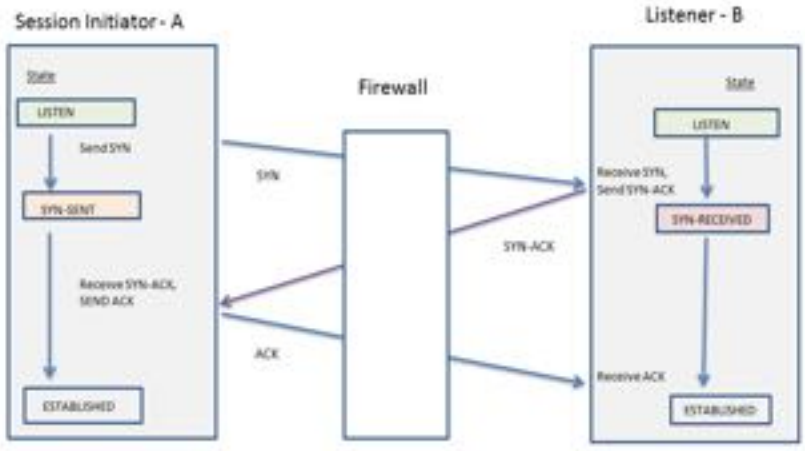
TCP Split Handshake Drop

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.

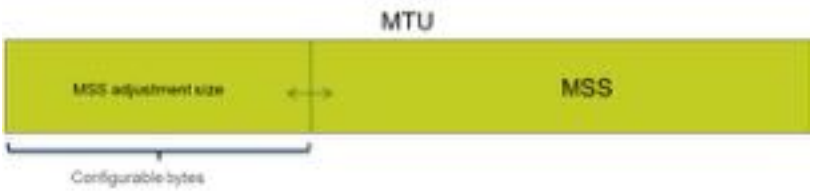
4-Way Split Handshake (Version 1)	4-Way Split Handshake (Version 2)	Simultaneous Open	5-Way Split Handshake
1. A → B SYN 2. A ← B ACK 3. A ← B SYN 4. A → B ACK	1. A → B SYN 2. A ← B SYN 3. A → B SYN-ACK 4. A ← B ACK	1. A → B SYN 2. A ← B SYN 3. A → B SYN-ACK 4. A ← B SYN-ACK	1. A → B SYN 2. A ← B ACK 3. A ← B SYN 4. A → B SYN-ACK 5. A ← B ACK

You can Prevent TCP Split Handshake Session Establishment.

Maximum Segment Size (MSS)

The maximum transmission unit (MTU) is a value indicating the largest number of bytes that can be transmitted in a single TCP packet. The MTU includes the length of headers, so the MTU minus the number of bytes in the headers equals the maximum segment size (MSS), which is the maximum number of data bytes that can be transmitted in a single packet.

A configurable MSS adjustment size (shown below) allows your firewall to pass traffic that has longer headers than the default setting allows. Encapsulation adds length to headers, so you would increase the MSS adjustment size to allow bytes, for example, to accommodate an MPLS header or tunneled traffic that has a VLAN tag.



If the DF (don't fragment) bit is set for a packet, it is especially helpful to have a larger MSS adjustment size and smaller MSS so that longer headers do not result in a packet length that exceeds the allowed MTU. If the DF bit were set and the MTU were exceeded, the larger packets would be dropped.

The firewall supports a configurable MSS adjustment size for IPv4 and IPv6 addresses on the following Layer 3 interface types: Ethernet, subinterfaces, Aggregated Ethernet (AE), VLAN, and loopback. The IPv6 MSS adjustment size applies only if IPv6 is enabled on the interface.

If IPv4 and IPv6 are enabled on an interface and the MSS Adjustment Size differs between the two IP address formats, the proper MSS value corresponding to the IP type is used for TCP traffic.

For IPv4 and IPv6 addresses, the firewall accommodates larger-than-expected TCP header lengths. In the case where a TCP packet has a larger header length than you planned for, the firewall chooses as the MSS adjustment size the larger of the following two values:

- The configured MSS adjustment size

 © Palo Alto Networks, Inc.

- The sum of the length of the TCP header (20) + the length of IP headers in the TCP SYN

This behavior means that the firewall overrides the configured MSS adjustment size if necessary. For example, if you configure an MSS adjustment size of 42, you expect the MSS to equal 1458 (the default MTU size minus the adjustment size [1500 - 42]). However, the TCP packet has 4 extra bytes of IP options in the header, so the MSS adjustment size (20+20+4) equals 44, which is larger than the configured MSS adjustment size of 42. The resulting MSS is 1500-44=1456 bytes, smaller than you expected.

To configure the MSS adjustment size, see Step 8 in Configure Session Settings.

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.

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.	Select **Device > Setup > Session** and edit the Session Timeouts.

 © Palo Alto Networks, Inc.

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 (range is 1-1599999; default is 30). • **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). • **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). • **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). 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.
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. • **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. • **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. • **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. • **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. • **TCP Time Wait**—Maximum length of time after receiving the second FIN or a RST. Default: 15. Range: 1-600. • **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. • 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. • **UDP**—Maximum length of time that a UDP session remains open without a UDP response. Default: 30. Range: 1-1599999. • See also the **Scan** timeout in the section (Optional) Change miscellaneous timeouts.

Change Session Timeouts (Continued)		
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. • 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 tasks if you need to change the default settings.

Configure Session Settings		
Step 1	Change the session settings.	Select **Device > Setup > Session** and edit the Session Settings.
Step 2	Specify whether to apply newly configured Security policy rules to sessions that are in progress.	Select **Rematch all sessions on config policy change** to apply newly configured Security policy rules to sessions that are already in progress. This capability is enabled by default. If you clear this check box, any policy rule changes you make apply only to sessions initiated after you commit the policy change. For example, if a Telnet session started while an associated policy rule 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.
Step 3	Configure IPv6 settings.	• **ICMPv6 Token Bucket Size**—Default: 100 tokens. See the section ICMPv6 Rate Limiting. • **ICMPv6 Error Packet Rate (per sec)**—Default: 100. See the section ICMPv6 Rate Limiting. • **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.
Step 4	Enable jumbo frames and set the MTU.	1. Select **Enable Jumbo Frame** to enable jumbo frame support on Ethernet interfaces. Jumbo frames have a maximum transmission unit (MTU) of 9216 bytes and are available on certain platforms. 2. Set the **Global MTU**, depending on whether or not you enabled jumbo frames: • If you did not enable jumbo frames, the **Global MTU** defaults to 1500 bytes; the range is 576 to 1500 bytes. • If you enabled jumbo frames, the **Global MTU** defaults to 9192 bytes; the range is 9192 to 9216 bytes.  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.

 © Palo Alto Networks, Inc.

Configure Session Settings (Continued)		

Step 5	Tune NAT session settings.	• **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.
		• **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.
		  • A rate of 1 means no oversubscription; each translated IP address and port pair can be used only once at a time.
		  • If the setting is **Platform Default**, user configuration of the rate is disabled and the default oversubscription rate for the platform applies.
		Reducing the oversubscription rate decreases the number of source device translations, but provides higher NAT rule capacities.

Step 6	Tune accelerated aging settings.	Select **Accelerated Aging** to enable faster aging-out of idle sessions. You can also change the threshold (%) and scaling factor:
		  • **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.
		  • **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.
		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.
		3. Click **OK**.

Configure Session Settings (Continued)		
Step 7	Enable buffering of multicast route setup packets.	1. Select **Multicast Route Setup Buffering** to enable the firewall to preserve the first packet in a multicast session when the multicast route or forwarding information base (FIB) entry does not yet exist for the corresponding multicast group. By default, the firewall does not buffer the first multicast packet in a new session; instead, it uses the first packet to set up the multicast route. This is expected behavior for multicast traffic. You only need to enable multicast route setup buffering if your content servers are directly connected to the firewall and your custom application cannot withstand the first packet in the session being dropped. This option is disabled by default. 2. If you enable buffering, you can also tune the **Buffer Size**, which specifies the buffer size per flow. The firewall can buffer a maximum of 5,000 packets.  You can also tune the duration, in seconds, for which a multicast route remains in the routing table on the firewall after the session ends by configuring the multicast settings on the virtual router that handles your virtual router (set the **Multicast Route Age Out Time (sec)** on the **Multicast > Advanced** tab in the virtual router configuration.
Step 8	Tune the Maximum Segment Size (MSS) adjustment size settings for a Layer 3 interface.	1. Select **Network > Interfaces**, select **Ethernet**, **VLAN**, or **Loopback**, and select a Layer 3 interface. 2. Select **Advanced**. 3. Select **Other Info**. 4. Select **Adjust TCP MSS** and enter a value for one or both of the following: • **IPv4 MSS Adjustment Size** (range is 40-300 bytes; default is 40 bytes). • **IPv6 MSS Adjustment Size** (range is 60-300 bytes; default is 60 bytes). 5. Click **OK**.
Step 9	Save the changes.	Click **Commit**.
Step 10	Reboot the firewall after changing the jumbo frame configuration.	1. Select **Device > Setup > Operations**. 2. Click **Reboot Device**.

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 you assigned a security zone for the interface where you want to prevent TCP split handshakes from establishing a session.

 © Palo Alto Networks, Inc.

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). 2. If creating a new profile, enter a **Name** for the profile and an optional **Description**. 3. Select **Packet Based Attack Protection > TCP Drop** and select **Split Handshake**. 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. 2. In the Zone window, from the **Zone Protection Profile** drop-down, select the profile you configured in the previous step. Alternatively, you could start creating a new profile here by clicking **Zone Protection Profile**, in which case you would continue accordingly. 3. Click **OK**. 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
- Firewall as a DHCP Server and Client
- DHCP Messages
- DHCP Addressing
- DHCP Options
- Configure an Interface as a DHCP Server
- Configure an Interface as a DHCP Client
- Configure the Management 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. You configure the firewall 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.

 © Palo Alto Networks, Inc.

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.

Firewall as a DHCP Server and Client

The firewall can function as a DHCP server and as a DHCP client. 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 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.

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.

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: • Request parameters from one server and implicitly decline offers from other servers. • Confirm that a previously allocated address is correct after, for example, a system reboot. • 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

▲ 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.

 © Palo Alto Networks, Inc.

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.

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 <interface> expired-only` command to clear expired leases, making those addresses available in the pool again. You can use the `clear dhcp lease interface <interface> ip <ip_address>` command to release a particular IP address. Use the `clear dhcp lease interface <interface> 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. 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
69	Simple Mail Transfer Protocol (SMTP) server address

 © Palo Alto Networks, Inc.

DHCP Option	DHCP Option Name
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. 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.

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.

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**. 2. Enter an **Interface** name or select one from the drop-down. 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. 4. (Optional) Select **Ping IP when allocating new IP** if you want the server to ping the IP address before it assigns that address to its client. 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. After you set options and return to the DHCP server tab, the **Probe IP** column for the interface indicates if **Ping IP when allocating new IP** was selected.

 © Palo Alto Networks, Inc.

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:
	  • **Unlimited** causes the server to dynamically choose IP addresses from the **IP Pools** and assign them permanently to clients.
	  • **Timeout** determines how long the lease will last. Enter the number of **Days** and **Hours**, and optionally the number of **Minutes**.
	• **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.
	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.
	  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.
	• **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.
	• **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.
	• **Subnet Mask**—Network mask used with the addresses in the **IP Pools**.

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**. • **Primary DNS, Secondary DNS**—IP address of the preferred and alternate Domain Name System (DNS) servers. • **Primary WINS, Secondary WINS**—IP address of the preferred and alternate Windows Internet Naming Service (WINS) servers. • **Primary NIS, Secondary NIS**—IP address of the preferred and alternate Network Information Service (NIS) servers. • **Primary NTP, Secondary NTP**—IP address of the available Network Time Protocol servers. • **POP3 Server**—IP address of a Post Office Protocol (POP3) server. • **SMTP Server**—IP address of a Simple Mail Transfer Protocol (SMTP) server. • **DNS Suffix**—Suffix for the client to use locally when an unqualified hostname is entered that it cannot resolve.
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. 2. Enter the **Option Code** you want to configure the server to offer (range is 1-254). (See RFC 2132 for option codes.) 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. 4. **Inherit from DHCP server inheritance source**—Select it 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. 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. 6. If you did not select **Inherit from DHCP server inheritance source**, select an **Option Type**: **IP Address**, **ASCII**, or **Hexadecimal**. Hexadecimal values must start with the 0x prefix. 7. Enter the **Option Value** you want the DHCP server to offer for that **Option Code**. You can enter multiple values on separate lines. 8. Click **OK.**

 © Palo Alto Networks, Inc.

Configure an Interface as a DHCP Server (Continued)		
Step 4	(Optional) Add another vendor-specific or custom DHCP option.	1. Repeat Step 3 to enter another custom DHCP Option. • 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. 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**. • 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. 2. Click **OK.**
Step 5	Identify the stateful pool of IP addresses from which the DHCP server chooses an address and assigns it to a DHCP client. 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). • An IP Pool or a **Reserved Address** is mandatory for dynamic IP address assignment. • An IP Pool is optional for static IP address assignment as long as the static IP addresses that you assign fall into the subnet that the firewall interface services. 2. (Optional) Repeat Step 1 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. See the DHCP Addressing section for an explanation of allocation of a **Reserved Address**.	1. In the **Reserved Address** field, click **Add.** 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. 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. 4. (Optional) Repeat Step 2 and Step 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.



You can also Configure the Management Interface as a DHCP Client.

You can also Configure the Management Interface as a DHCP Client.

Configure an Interface as a DHCP Client		
Step 1	Configure an interface as a DHCP client.	1. Select **Network** > **Interfaces**. 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. 3. Click the **IPv4** tab; for **Type**, select **DHCP Client**. 4. Select **Enable**. 5. (Optional) Select **Automatically create default route pointing to default gateway provided by server**. 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. 6. (Optional) 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. 7. (Optional) Select **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. 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. 2. Select **Network > Interfaces > VLAN** and look in the **IP Address** field to see which interfaces indicate DHCP Client.

 © Palo Alto Networks, Inc.

Configure the Management Interface as a DHCP Client

The management interface on the firewall supports DHCP client for IPv4, which allows the management interface to receive its IPv4 address from a DHCP server. The management interface also supports DHCP Option 12 and Option 61, which allow the firewall to send its hostname and client identifier, respectively, to DHCP servers.

By default, VM-Series firewalls deployed in AWS and Azure™ use the management interface as a DHCP client to obtain its IP address, rather than a static IP address, because cloud deployments require the automation this feature provides. DHCP on the management interface is turned off by default for the VM-Series firewall except for the VM-Series firewall in AWS and Azure. The management interfaces on WildFire and Panorama platforms do not support this DHCP functionality.



- For hardware-based firewall platforms (not VM-Series), configure the management interface with a static IP address when possible.
- If the firewall acquires a management interface address through DHCP, assign a MAC address reservation on the DHCP server that serves that firewall. The reservation ensures that the firewall retains its management IP address after a restart. If the DHCP server is a Palo Alto Networks firewall, see Step 6 of Configure an Interface as a DHCP Server for reserving an address.

If you configure the management interface as a DHCP client, the following restrictions apply:

- You cannot use the management interface in an HA configuration for control link (HA1 or HA1 backup), data link (HA2 or HA2 backup), or packet forwarding (HA3) communication.
- You cannot select **MGT** as the Source Interface when you customize service routes (**Device > Setup > Services > Service Route Configuration > Customize**). However, you can select **Use default** to route the packets via the management interface.
- You cannot use the dynamic IP address of the management interface to connect to a Hardware Security Module (HSM). The IP address on the HSM client firewall must be a static IP address because HSM authenticates the firewall using the IP address, and operations on HSM would stop working if the IP address were to change during runtime.

A prerequisite for this task is that the management interface must be able to reach a DHCP server.

Configure the Management Interface as a DHCP Client		
Step 1	Configure the Management interface as a DHCP client so that it can receive its IP address (IPv4), netmask (IPv4), and default gateway from a DHCP server. Optionally, you can also send the hostname and client identifier of the management interface to the DHCP server if the orchestration system you use accepts this information.	1. Select **Device > Setup > Management** and edit Management Interface Settings. 2. For **IP Type**, select **DHCP Client**. 3. (Optional) Select one or both options for the firewall to send to the DHCP server in DHCP Discover or Request messages: • **Send Hostname**—Sends the **Hostname** (as defined in **Device > Setup > Management**) as part of DHCP Option 12. • **Send Client ID**—Sends the client identifier as part of DHCP Option 61. A client identifier uniquely identifies a DHCP client, and the DHCP Server uses it to index its configuration parameter database. 4. Click **OK**.

Configure the Management Interface as a DHCP Client (Continued)		
Step 2	(Optional) Configure the firewall to accept the host name and domain from the DHCP server.	1. Select **Device > Setup > Management** and edit General Settings. 2. Select one or both options: • **Accept DHCP server provided Hostname**—Allows the firewall to accept the hostname from the DHCP server (if valid). When enabled, the hostname from the DHCP server overwrites any existing **Hostname** specified in **Device > Setup > Management**. Do not select this option if you want to manually configure a hostname. • **Accept DHCP server provided Domain**—Allows the firewall to accept the domain from the DHCP Server. The domain (DNS suffix) from the DHCP Server overwrites any existing **Domain** specified in **Device > Setup > Management**. Do not select this option if you want to manually configure a domain. 3. Click **OK**.
Step 3	Save the configuration.	Click **Commit**.
Step 4	View DHCP client information.	1. Select **Device > Setup > Management** and Management Interface Settings. 2. Click **Show DHCP Client Runtime Info**.
Step 5	(Optional) Renew the DHCP lease with the DHCP server, regardless of the lease term. This option is convenient if you are testing or troubleshooting network issues.	1. Select **Device > Setup > Management** and edit Management Interface Settings. 2. Click **Show DHCP Client Runtime Info**. 3. Click **Renew**.
Step 6	(Optional) Release the following DHCP options that came from the DHCP server: • IP Address • Netmask • Default Gateway • DNS Server (primary and secondary) • NTP Server (primary and secondary) • Domain (DNS Suffix) A release frees the IP address, which drops your network connection and renders the firewall unmanageable if no other interface is configured for management access.	Use the CLI operational command `request dhcp client management-interface release`.

 © Palo Alto Networks, Inc.

Configure an Interface as a DHCP Relay Agent

To enable a firewall interface to transmit DHCP messages between clients and servers, you must configure the firewall as a DHCP relay agent. The interface can forward messages to a maximum of eight external IPv4 DHCP servers and eight external IPv6 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. Before configuring 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.

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. 2. Select either **IPv4** or **IPv6**, indicating the type of DHCP server address you will specify. 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. 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**. 5. (Optional) Repeat Steps 2-4 to enter a maximum of eight DHCP server addresses per IP address family.
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.

- 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 DHCP 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.

DNS

Domain Name System (DNS) is a protocol that translates (resolves) a user-friendly domain name, such as www.paloaltonetworks.com, to an IP address so that users can access computers, websites, services, or other resources on the internet or private networks.

- DNS Overview
- DNS Proxy Object
- DNS Server Profile
- Multi-Tenant DNS Deployments
- Configure a DNS Proxy Object
- Configure a DNS Server Profile
- 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
- Reference: DNS Proxy Rule and FQDN Matching

DNS Overview

DNS performs a crucial role in enabling user access to network resources so that users need not remember IP addresses and individual computers need not store a huge volume of domain names mapped to IP addresses. DNS employs a client/server model; a DNS server resolves a query for a DNS client by looking up the domain in its cache and if necessary sending queries to other servers until it can respond to the client with the corresponding IP address.

The DNS structure of domain names is hierarchical: the top-level domain (TLD) in a domain name can be a generic TLD (gTLD): com, edu, gov, int, mil, net, or org (gov and mil are for the United States only) or a country code (ccTLD), such as au (Australia) or us (United States). ccTLDs are generally reserved for countries and dependent territories; they are sometimes used in an unrelated context.

A fully qualified domain name (FQDN) includes at a minimum a host name, a second-level domain, and a TLD to completely specify the location of the host in the DNS structure. For example, www.paloaltonetworks.com is an FQDN.

Wherever a Palo Alto Networks firewall uses an FQDN in the user interface or CLI, the firewall must resolve that FQDN using DNS. Depending on where the FQDN query originates, the firewall determines which DNS settings to use to resolve the query. The firewall uses DNS in a number of ways:

- You must configure your firewall with at least one DNS server so it can resolve hostnames. Configure primary and secondary DNS servers or a DNS proxy object that specifies such servers, as shown in Use Case 1: Firewall Requires DNS Resolution for Management Purposes.
- Security policies, reporting, and services often require DNS, as shown in Use Case 2: ISP Tenant Uses DNS Proxy to Handle DNS Resolution for Security Policies, Reporting, and Services within its Virtual System.
- The firewall can act as a DNS server for a client, as shown in Use Case 3: Firewall Acts as DNS Proxy Between Client and Server.

- You can configure an Anti-Spyware profile to Use DNS Queries to Identify Infected Hosts on the Network.
- You can Enable Passive DNS Collection for Improved Threat Intelligence.
- You can Enable DNS Proxy and then enable evasion signatures for threat prevention.
- When you Configure an Interface as a DHCP Server, the firewall acts as a DHCP Server and sends DNS information to its DHCP clients so the provisioned DHCP clients can reach their respective DNS servers.

DNS Proxy Object

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 it doesn't find the domain name in its 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). The firewall forwards the query to the appropriate DNS server based on the match results. If no match is found, the firewall uses default DNS servers.

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 doesn't 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.

When you Configure a DNS Proxy Object, you can supply the DNS proxy with static FQDN-to-address mappings. You can also create DNS proxy rules that control to which DNS server the domain name queries (that match the proxy rules) are directed. You can configure a maximum of 256 DNS proxy objects on a firewall.

When the firewall receives an FQDN query (and the domain name is not in the DNS proxy cache), the firewall compares the domain name from the FQDN query to the domain names in DNS Proxy rules of the DNS Proxy object. If you specify multiple domain names in a single DNS Proxy rule, a query that matches any one of the domain names in the rule means the query matches the rule. Reference: DNS Proxy Rule and FQDN Matching describes how the firewall determines whether an FQDN matches a domain name in a DNS proxy rule. A DNS query that matches a rule is sent to the primary DNS server configured for the proxy object to be resolved.

 © 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 of 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 route table of the virtual router where the source interface is assigned. It's possible that the result of the destination IP egress interface differs from the source interface. The packet would egress out 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 firewall is queried.

You Configure a DNS Server Profile for a virtual system only; it's not for a global Shared location.

Multi-Tenant DNS Deployments

The firewall determines how to handle DNS requests based on where the request originated. An environment where an ISP has multiple tenants on a firewall is known as multi-tenancy. 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, the request comes from the management plane to resolve an FQDN for a management event such as a software update service. The firewall uses the service route to get to a DNS server because the DNS request isn't coming in on a specific virtual router.
- **Policy and Report FQDN Resolution for a Virtual System**—For DNS queries from a security policy, a report, or a service, 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, you must configure a DNS Proxy Object. 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 the virtual system, the firewall uses the global DNS settings.
- **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 Object 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.

The following table summarizes the DNS resolution types. 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 Illustrated in Use Case 1	N/A
Security profile, reporting, and server profile resolution—performed by management plane	Binding: Global Same behavior as Use Case 1	Binding: Specific vsys 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 Service Route: Interface and IP address on which the DNS Request was received. 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

Configure a DNS Proxy Object

If your firewall is to act as a DNS proxy, 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. 2. Verify that **Enable** is selected. 3. Enter a **Name** for the object. 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. 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. 6. For Inheritance Source, select a source from which to inherit default DNS server settings. The default is **None**. 7. For **Interface**, click **Add** and specify the interfaces to which the DNS Proxy object applies. • 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. • If you use the DNS Proxy object for a service route, the interface is optional.

 © Palo Alto Networks, Inc.

Configure a DNS Proxy Object (Continued)		
Step 2	(Optional) Specify DNS Proxy rules.	1. On the **DNS Proxy Rules** tab, click **Add** and enter a **Name** for the rule. 2. **Turn on caching of domains resolved by this mapping** if you want the firewall to cache the resolved domains. 3. For **Domain Name**, click **Add** and enter one or more domains, one entry per row, to which the firewall compares FQDN queries. If a query matches one of the domains in the rule, the query is sent to one of the following servers to be resolved (depending on what you configured in the prior step): • The **Primary** or **Secondary** DNS Server directly specified for this proxy object. • The **Primary** or **Secondary** DNS Server specified in the DNS Server profile for this proxy object. Reference: DNS Proxy Rule and FQDN Matching describes how the firewall matches domain names in an FQDN to a DNS proxy rule. If no match is found, default DNS servers resolve the query. 4. Do one of the following, depending on what you set the **Location** to: • If you chose a virtual system, select a **DNS Server profile**. • If you chose **Shared**, enter a **Primary** and optionally a **Secondary** address. 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 sending a query to the DNS server.	1. On the **Static Entries** tab, click **Add** and enter a **Name**. 2. Enter the **FQDN**. 3. For **Address**, click **Add** and enter the IP address to which the FQDN should be mapped. You can enter additional IP addresses for an FQDN; the firewall provides all of the IP addresses in its DNS response and the client chooses which address to use. 4. Click **OK**.
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.  Leave **Size** and **Timeout** settings with default values. Beginning with PAN-OS 7.1.1 and later releases, the DNS proxy automates these settings to maximize efficiency. 2. Select **TCP Queries** to enable DNS queries using TCP. • **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). 3. For **UDP Queries Retries**, enter the following: • **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). • **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.)

Configure a DNS Proxy Object (Continued)		
Step 5	Save the configuration.	Click **OK** and **Commit**.

Configure a DNS Server Profile

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.

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**. 2. Enter a **Name** for the DNS server profile. 3. Select the **Location** (virtual system) to which the profile applies. 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. 5. Specify the IP address of the **Primary DNS** server, or leave as **inherited** if you chose an **Inheritance Source**.  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**. 6. Specify the IP address of the **Secondary DNS** server, or leave as **inherited** if you chose an **Inheritance Source**.

 © Palo Alto Networks, Inc.

Configure a DNS Server Profile		
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. 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). 3. Specify the IPv4 **Source Address** from which packets going to the DNS server are sourced. 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. 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). 6. Specify the IPv6 **Source Address** from which packets going to the DNS server are sourced. 7. Click **OK**.
Step 3	Save the configuration.	Click **OK** and **Commit**.

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.



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. 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 firewalls that do not support multiple virtual systems, there is no **Global** tab; simply edit the Services.) 2. On the **Services** tab, for **DNS**, click **Servers** and enter the **Primary DNS Server** address and **Secondary DNS Server** address. 3. Click **OK** and **Commit**.
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. Select **Device > Setup > Services > Global** and Edit. 2. On the **Services** tab, for **DNS**, select **DNS Proxy Object**. 3. 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 follows: a. Click **Enable** and enter a **Name** for the DNS proxy object. b. For **Location**, select **Shared** for global, firewall-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. c. Enter the **Primary** DNS server IP address. Optionally enter a **Secondary** DNS server IP address. 4. Click **OK** and **Commit**.

Use Case 2: ISP Tenant Uses DNS Proxy to Handle DNS Resolution for Security Policies, Reporting, and Services within its Virtual System

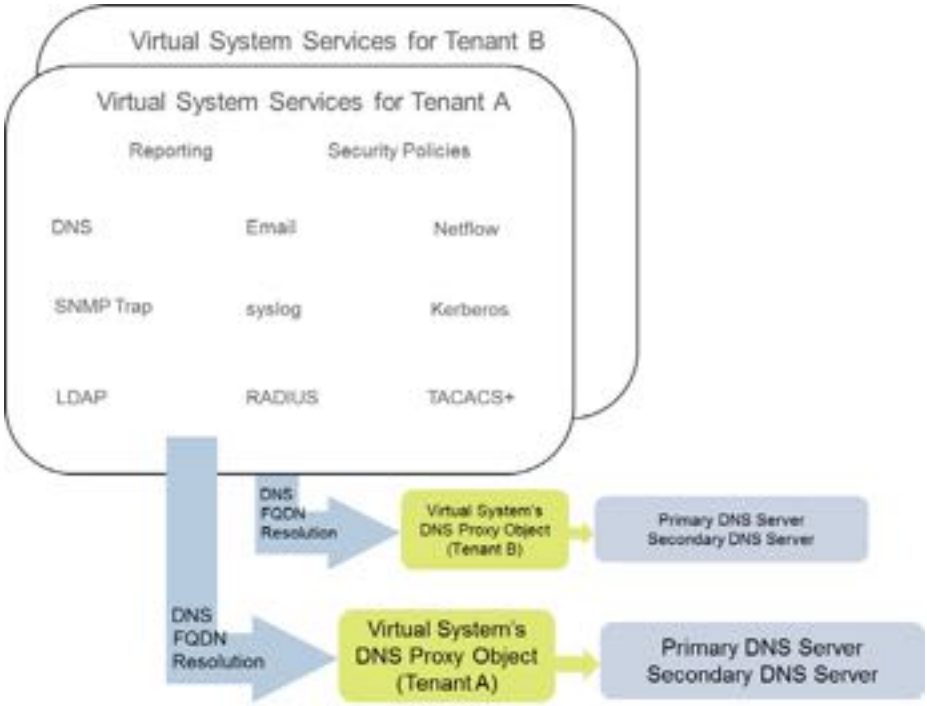
In this use case, multiple tenants (ISP subscribers) are defined on the firewall and each tenant is allocated a separate virtual system (vsys) and virtual router in order to segment its services and administrative domains. The following figure illustrates several virtual systems within a firewall.



 © Palo Alto Networks, Inc.

Each tenant has its own server profiles for Security policy rules, 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.



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**. 2. Enter the **ID** of the virtual system (range is 1-255), and an optional **Name**, in this example, Corp1 Corporation. 3. On the **General** tab, choose a **DNS Proxy** or create a new one. 4. For **Interfaces**, click **Add**. In this example, Ethernet1/20 is dedicated to this tenant. 5. For **Virtual Routers**, click **Add**. A virtual router is assigned to the virtual system in order to separate routing functions. 6. Click **OK** to save the configuration.

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. (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. 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.



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.

 © Palo Alto Networks, Inc.



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 to send the request, 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	
Configure a DNS Proxy and DNS proxy rules.	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. 4. For **Interface**, select the interface that will receive the DNS requests from the tenant's hosts. 5. Choose or create a **Server Profile** to customize DNS servers to resolve DNS requests for this tenant. 6. On the **DNS Proxy Rules** tab, click **Add** and enter a **Name** for the rule. 7. Select **Turn on caching of domains resolved by this mapping**. 8. Click **Add** and enter one or more **Domain Name**(s) for the rule, one entry per row. Reference: DNS Proxy Rule and FQDN Matching describes how the firewall matches FQDNs to domain names in a DNS proxy rule. 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. 10. Click **OK** to save the rule. 11. Click **OK** to save the DNS Proxy.

Reference: DNS Proxy Rule and FQDN Matching

When you configure the firewall with a DNS Proxy Object that uses DNS proxy rules, the firewall compares an FQDN from a DNS query to the domain name of a DNS proxy rule. The firewall comparison works as follows:

FQDN Comparison to DNS Proxy Rule	For Example
The firewall first tokenizes the FQDNs and the domains names in the DNS proxy rules. In a domain name, a string delimited by a period (.) is a token.	`*.boat.fish.com` consists of four tokens: `[*][boat][fish][com]`
The matching process is an exact token match between the FQDN and the domain name in the rule; partial strings are not matched.	Rule: `fishing` FQDN: `fish` — **Not a Match**
An exception to the exact match requirement is the use of the wildcard—an asterisk (*). The * matches one or more tokens. This means a rule consisting of only a wildcard (*) matches any FQDN with one or more tokens.	Rule: `*.boat.com` FQDN: `www.boat.com` — Match FQDN: `www.blue.boat.com` — Match FQDN: `boat.com` — **Not a Match**
	Rule: * FQDN: `boat` — Match FQDN: `boat.com` — Match FQDN: `www.boat.com` — Match
You can use an * in any position: preceding tokens, between tokens, or trailing tokens (but not with other characters within a single token).	Rule: `www.*.com` FQDN: `www.boat.com` — Match FQDN: `www.blue.boat.com` — Match
	Rule: `www.boat.*` FQDN: `www.boat.com` — Match FQDN: `www.boat.fish.com` — Match
	Rule: `www.boat*.com` — **Invalid**
Multiple wildcards (*) can appear in any position of the domain name: preceding tokens, between tokens, or trailing tokens. Each non-consecutive * matches one or more tokens.	Rule: `a.*.d.*.com` FQDN: `a.b.d.e.com` — Match FQDN: `a.b.c.d.e.f.com` — Match FQDN: `a.d.d.e.f.com` — Match (First `*` matches `d`; second `*` matches `e` and `f`) FQDN: `a.d.e.f.com` — **Not a Match** (First `*` matches `d`; subsequent `d` in the rule is not matched)

 © Palo Alto Networks, Inc.

FQDN Comparison to DNS Proxy Rule	For Example
When wildcards are used in consecutive tokens, the first * matches one or more tokens; the second * matches one token. This means a rule consisting of only *.* matches any FQDN with two or more tokens.	Consecutive wildcards preceding tokens: Rule: `*.*.boat.com` FQDN: `www.blue.boat.com` — Match FQDN: `www.blue.sail.boat.com` — Match
	Consecutive wildcards between tokens: Rule: `www.*.*.boat.com` FQDN: `www.blue.sail.boat.com` — Match FQDN: `www.big.blue.sail.boat.com` — Match
	Consecutive wildcards trailing tokens: Rule: `www.boat.*.*` FQDN: `www.boat.fish.com` — Match FQDN: `www.boat.fish.ocean.com` — Match
	Consecutive wildcards only: Rule: *.* FQDN: `boat` — **Not a Match** FQDN: `boat.com` — Match FQDN: `www.boat.com` — Match
Consecutive and non-consecutive wildcards can appear in the same rule.	Rule: `a.*.d.*.*.com` FQDN: `a.b.c.d.e.f.com` — Match (First `*` matches `b` and `c`; second `*` matches `e`; third `*` matches `f`) FQDN: `a.b.c.d.e.com` — **Not a Match** (First `*` matches `b` and `c`; second `*` matches `e`; third `*` not matched)
The Implicit-tail-match behavior provides an additional shorthand: As long as the last token of the rule is not an *, a comparison will match if all tokens in the rule match the FQDN, even when the FQDN has additional trailing tokens that the rule doesn't have.	Rule: `www.boat.fish` FQDN: `www.boat.fish.com` — Match FQDN: `www.boat.fish.ocean.com` — Match FQDN: `www.boat.fish` — Match
This rule ends with *, so the Implicit-tail-match rule doesn't apply. The * behaves as stated; it matches one or more tokens.	Rule: `www.boat.fish.*` FQDN: `www.boat.fish.com` — Match FQDN: `www.boat.fish.ocean.com` — Match FQDN: `www.boat.fish` — **Not a Match** (This FQDN does not have a token to match the * in the rule.)

FQDN Comparison to DNS Proxy Rule	For Example
In the case where an FQDN matches more than one rule, a tie-breaking algorithm selects the most specific (longest) rule; that is, the algorithm favors the rule with more tokens and fewer wildcards (*).	Rule 1: `*.fish.com` — Match Rule 2: `*.com` — Match Rule 3: `boat.fish.com` — Match and Tie-Breaker FQDN: `boat.fish.com` FQDN matches all three rules; the firewall uses Rule 3 because it is the most specific.
	Rule 1: `*.fish.com` — **Not a Match** Rule 2: `*.com` — Match Rule 3: `boat.fish.com` — **Not a Match** FQDN: `fish.com` FQDN does not match Rule 1 because the * does not have a token to match.
	Rule 1: `*.fish.com` — Match and Tie-Breaker Rule 2: `*.com` — Match Rule 3: `boat.fish.com` — **Not a Match** FQDN: `blue.boat.fish.com` FQDN matches Rule 1 and Rule 2 (because the * matches one or more tokens). The firewall uses Rule 1 because it is the most specific.
When working with wildcards (*) and Implicit-tail-match rules, there can be cases when the FQDN matches more than one rule and the tie-breaking algorithm weighs the rules equally. To avoid ambiguity, if rules with an Implicit-tail-match or a wildcard (*) can overlap, replace an Implicit-tail-match rule by specifying the tail token.	Replace this: Rule: `www.boat` with this: Rule: `www.boat.com`

When creating DNS proxy rules, the following best practices will help you avoid ambiguity and unexpected results:

Best Practices for Creating DNS Proxy Rules	For Example
Avoid invoking an Implicit-tail-match by including a top-level domain in the domain name.	`boat.com`
If you use a wildcard (*), use it only as the leftmost token. This practice follows the common understanding of wildcard DNS records and the hierarchical nature of DNS.	`*.boat.com`
Use no more than one * in a rule.	

 © Palo Alto Networks, Inc.

Best Practices for Creating DNS Proxy Rules	For Example
Use the * to establish a base rule associated with a DNS server, and use rules with more tokens to build exceptions to the rule, which you associate with different servers. The tie-breaking algorithm will select the most specific match, based on the number of matched tokens.	Rule: `*.corporation.com` — DNS server A Rule: `www.corporation.com` — DNS server B Rule: `*.internal.corporation.com` — DNS server C Rule: `www.internal.corporation.com` — DNS server D FQDN: `mail.internal.corporation.com` — matches DNS server C FQDN: `mail.corporation.com` — matches DNS server A

NAT

This section describes Network Address Translation (NAT) and how to configure the firewall for NAT. NAT allows you to translate private, non-routable IPv4 addresses to one or more globally-routable IPv4 addresses, thereby conserving an organization's routable IP addresses. NAT allows you to not disclose the real IP addresses of hosts that need access to public addresses and to manage traffic by performing port forwarding. You can use NAT to solve network design challenges, enabling networks with identical IP subnets to communicate with each other. The firewall supports NAT on Layer 3 and virtual wire interfaces.

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 and ports need translation and what the translated addresses and ports 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 Overview
- NAT Address Pools Identified as Address Objects
- Proxy ARP for NAT Address Pools

NAT Policy Overview

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.

 © Palo Alto Networks, Inc.

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. You must configure security policy rules to allow the NAT traffic.

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.



A SIP call sometimes experiences one-way audio when going through the firewall because the call manager sends a SIP message on behalf of the phone to set up the connection. When the message from the call manager reaches the firewall, the SIP ALG must put the IP address of the phone through NAT. If the call manager and the phones are not in the same security zone, the NAT lookup of the IP address of the phone is done using the call manager zone. The NAT policy should take this into consideration.

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.

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
<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

▲ 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 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:

- **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.

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 oversubscription setting (8, 4, 2, or 1) of the firewall 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 firewall. 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

 © Palo Alto Networks, Inc.

Platform	Default Oversubscription Rate
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
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.

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.



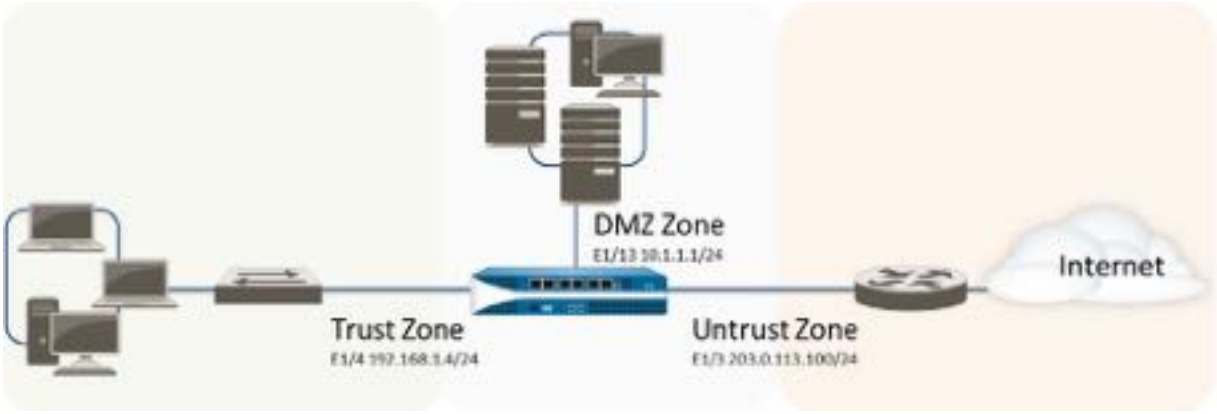
 © Palo Alto Networks, Inc.

Configure NAT

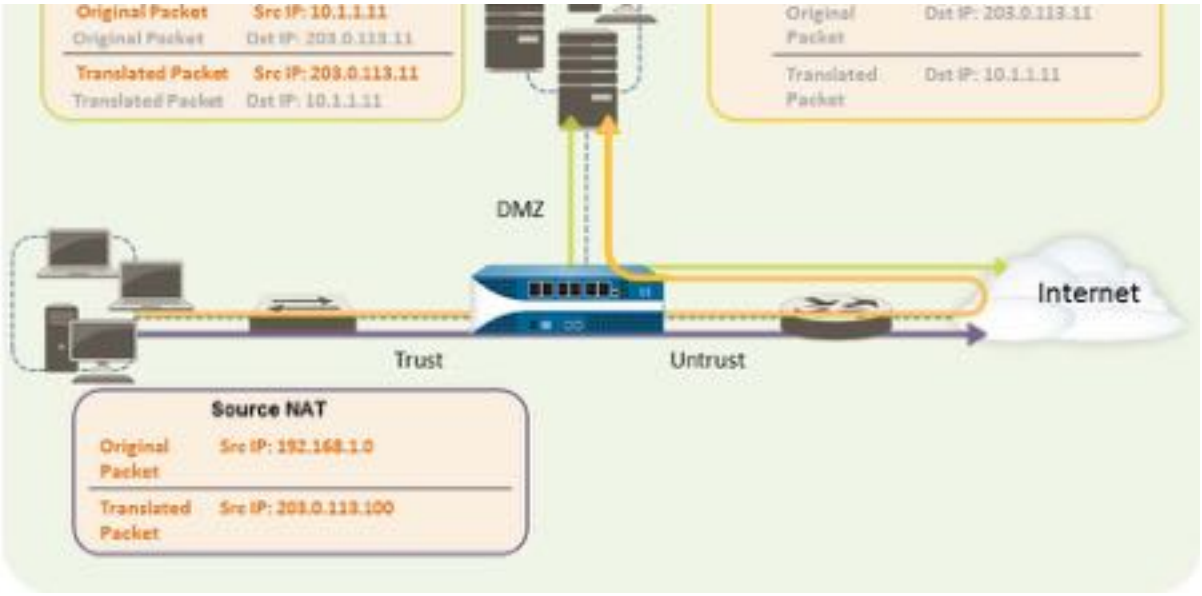
Perform the following tasks to configure various aspects of NAT. In addition to the examples below, there are examples in the section NAT Configuration Examples.

- 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:



Based on this topology, 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**. 2. Enter a **Name** and optional **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**.  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. On the **General** tab, enter a descriptive **Name** for the policy. 3. (Optional) Enter a tag, which is a keyword or phrase that allows you to sort or filter policies. 4. For **NAT Type**, select **ipv4** (default). 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. 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. 7. For **Address Type**, there are two choices. You could select **Translated Address** and then click **Add**. Select the address object you just created. 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. 8. Click **OK** to save the NAT policy.
Step 3	Save the configuration.	Click **Commit**.

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. 2. Use the `show session id <id_number>` to view more details about a session. 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**. 2. Enter a **Name** and optional **Description** for the object. 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. 4. Click **OK**.
Step 2 Create the NAT policy.	1. Select **Policies > NAT** and click **Add**. 2. On the **General** tab, 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 **Destination Address Translation** 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**.

 © Palo Alto Networks, Inc.

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.

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**. 2. Enter a **Name** and optional **Description** for the object. 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. 4. Click **OK**.  If you did not already create an address object for the *public* address of your web server, you should create that object now.
Step 2	Create the NAT policy.	1. Select **Policies > NAT** and click **Add**. 2. On the **General** tab, 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**.

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. Edit the Session Settings section. 2. In the **NAT Oversubscription Rate** drop-down, select **1x**, **2x**, **4x**, or **8x**, depending on which ratio you want. The **Platform Default** setting applies the default oversubscription setting for the platform. If you want no oversubscription, select **1x**. 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**. 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. For **Source Address**, click **Add** and enter the host address. Click **OK**. 5. On the **Translated Packet** tab, select **None** from the **Translation Type** drop-down in the Source Address Translation section of the screen. 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.

 © Palo Alto Networks, Inc.

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	

❑ Reserve dynamic IP NAT addresses for a firewall.	Enter the following commands:
	`admin@PA-3020# set setting nat reserve-ip yes`
	`admin@PA-3020# set setting nat reserve-time <1-604800 secs>`

❑ Reserve dynamic IP NAT addresses for a virtual system.	Enter the following commands:
	`admin@PA-3020# set vsys <vsysid> setting nat reserve-ip yes`
	`admin@PA-3020# 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.

NAT Configuration Examples

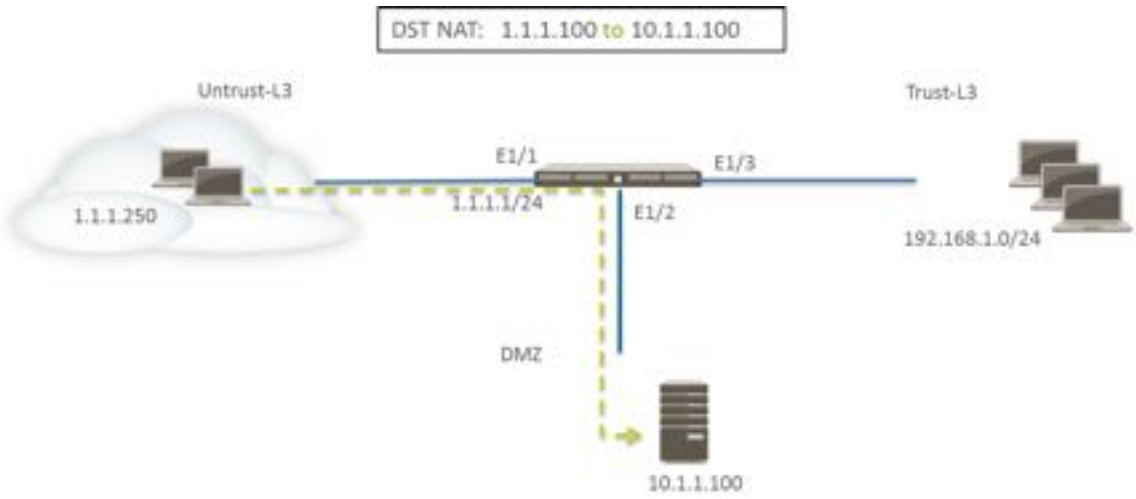
- Destination NAT Example—One-to-One Mapping
- Destination NAT with Port Translation Example
- Destination NAT Example—One-to-Many Mapping
- Source and Destination NAT Example
- Virtual Wire Source NAT Example
- Virtual Wire Static NAT Example
- Virtual Wire Destination NAT Example

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 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).

 © Palo Alto Networks, Inc.

- 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.
- The firewall performs a security policy lookup to see if the traffic is permitted from zone Untrust-L3 to DMZ.

  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.

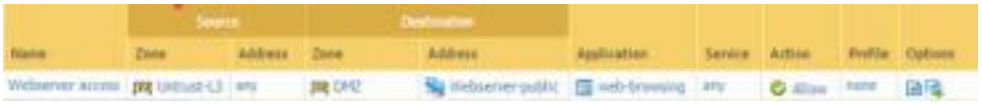
- The firewall forwards the packet to the server out egress interface Ethernet1/2. The destination address is changed to 10.1.1.100 as the packet leaves the firewall.

For this example, address objects are configured for webserver-private (10.1.1.100) and Webserver-public (1.1.1.100). The configured NAT rule would look like this:



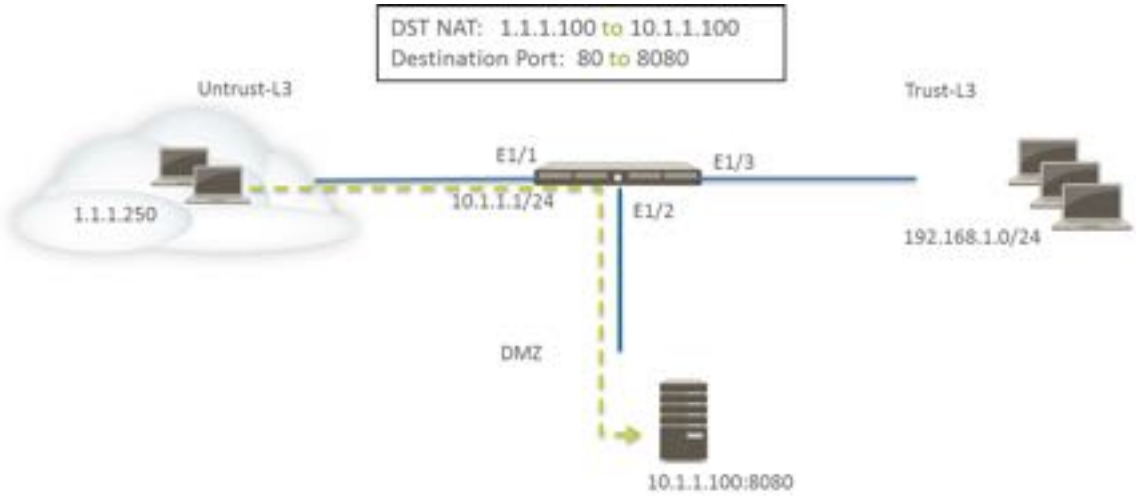
The direction of the NAT rules is based on the result of route lookup.

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. Address objects are configured for webserver-private (10.1.1.100) and Servers-public (1.1.1.100).



The following NAT and security rules must be configured on the firewall:

	Name	Original Packet						Translated Packet	
		Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation
1	Dst NAT-webserver	Untrust-L3	Untrust-L3	any	any	Servers-public	any	none	address: webserver-private port: 8080

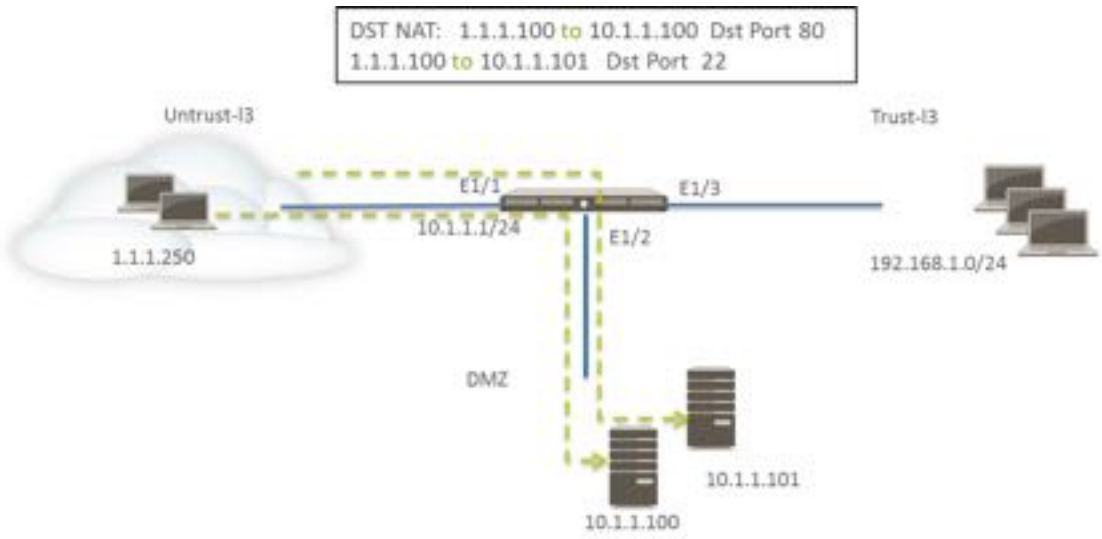
	Source		Destination			
Name	Zone	Address	Zone	Address	Application	Service
Webserver access	Untrust-L3	any	DMZ	Servers-public	web-browsing	any

Use the `show session all` CLI command to verify the translation.

 © Palo Alto Networks, Inc.

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 corresponding address objects are created:

- Servers-public: 1.1.1.100
- SSH-server: 10.1.1.101
- webserver-private: 10.1.1.100

The NAT rules would look like this:

	Name	Original Packet						Translated Packet	
		Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation
1	Dst NAT-webserver	Untrust-L3	Untrust-L3	any	any	Servers-public	service-http	none	address: webserver-private
2	Dst NAT-SSH	Untrust-L3	Untrust-L3	any	any	Servers-public	custom-ssh	none	address: SSH-server

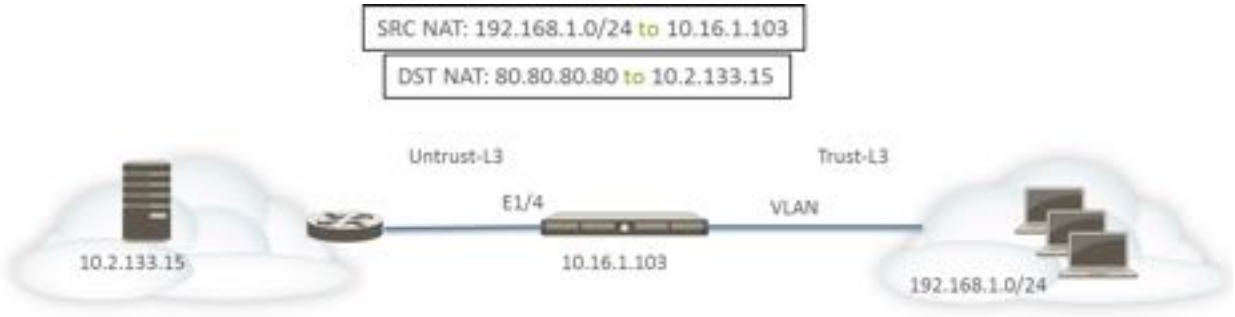
The security rules would look like this:

	Name	Source		Destination		Application	Service	Action
		Zone	Address	Zone	Address			
1	Webserver access	Untrust-L3	any	DMZ	Servers-public	web-browsing	application-default	Allow
2	SSH access	Untrust-L3	any	DMZ	Servers-public	ssh	application-default	Allow

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.

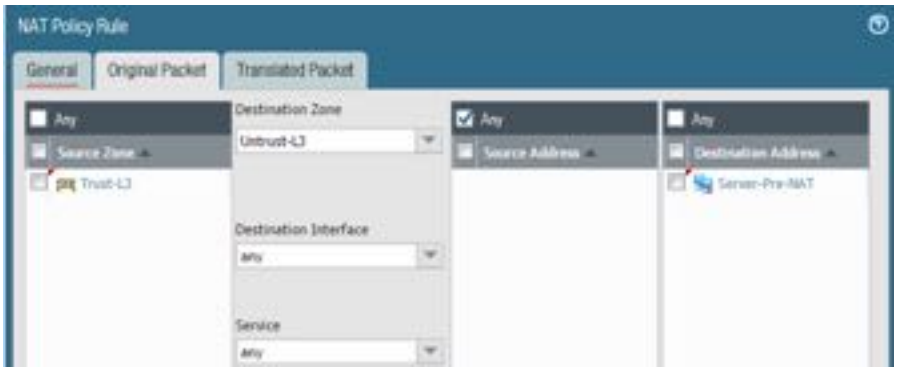
- Source NAT—The source addresses in the packets from the clients in the Trust-L3 zone to the server in the Untrust-L3 zone are translated from the private addresses in the network 192.168.1.0/24 to the IP address of the egress interface on the firewall (10.16.1.103). Dynamic IP and Port translation 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).



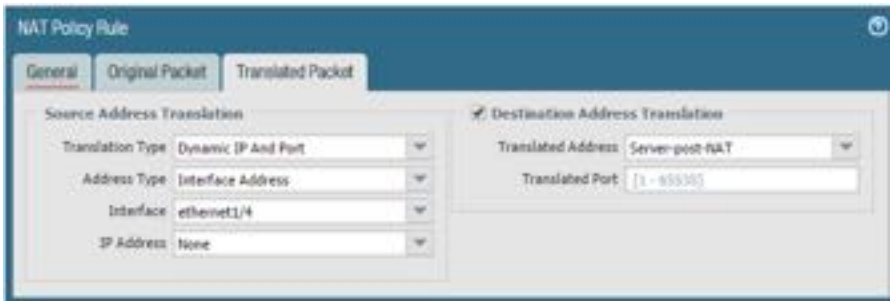
The following address objects are created for destination NAT.

- 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.



 © Palo Alto Networks, Inc.



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.

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. The firewall will not proxy ARP for NAT addresses. Proper routing must be configured on the upstream and downstream routers in order for the packets to be translated in virtual wire mode. 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. See Proxy ARP for NAT Address Pools for more explanation about proxy ARP.

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:



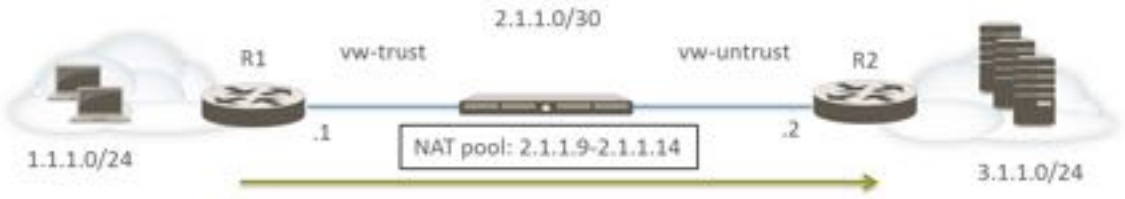
Route on R1:

Destination	Next Hop
3.1.1.0/24	2.1.1.2

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.



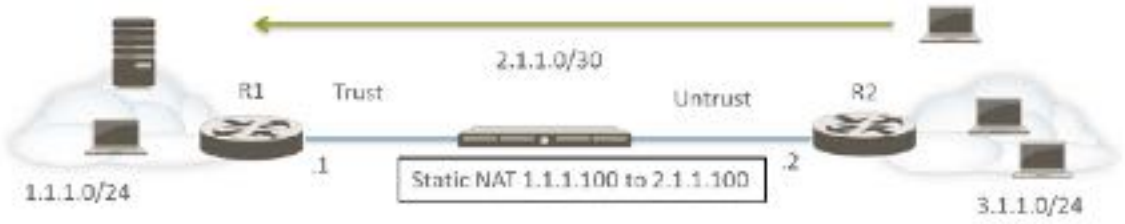
All connections from the clients in subnet 1.1.1.0/24 will arrive at router R2 with a translated source address in the range 2.1.1.9-2.1.1.14. The response from servers will be directed to these addresses. In order for source NAT to work, you must configure proper routing on router R2, so that packets 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 Trust to the virtual wire zone named 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 Untrust zone to the Trust zone. Clients on the 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:

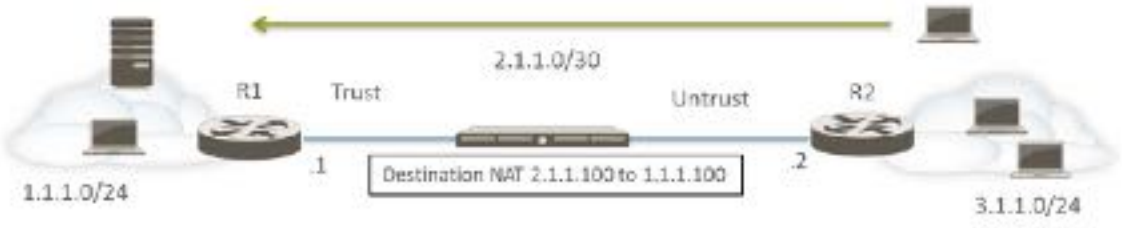
Destination	Next Hop
2.1.1.100/32	2.1.1.1

 © Palo Alto Networks, Inc.



Virtual Wire Destination NAT Example

Clients in the 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 Untrust zone to the Trust zone.



Route on R2:

Destination	Next Hop
2.1.1.100/32	2.1.1.1



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

 © Palo Alto Networks, Inc.

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-3060 firewall, 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.

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:

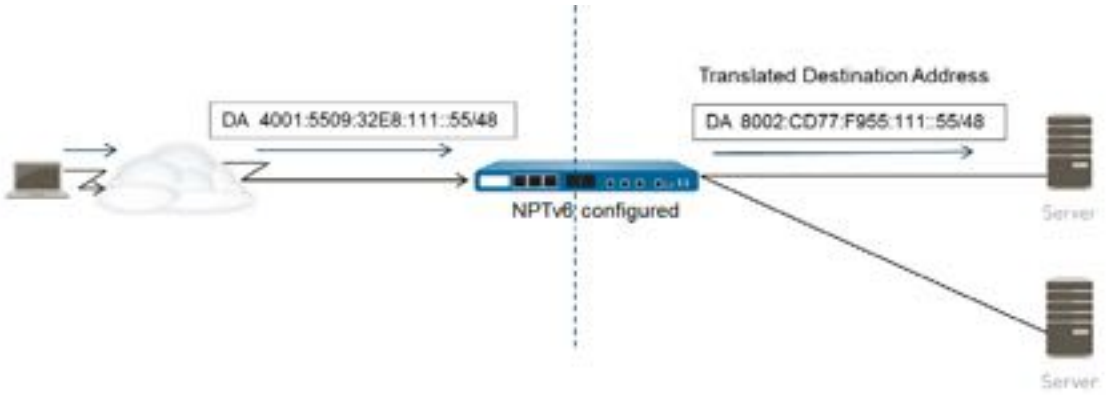
- **Prevents asymmetrical routing**—Asymmetric routing can occur if a Provider Independent address space (/48, for example) is advertised by multiple data centers to the global Internet. By using NPTv6, you can advertise more specific routes from regional firewalls, and the return traffic will arrive 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.

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.

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

 © Palo Alto Networks, Inc.

▲ 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** option 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:

- 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.
- If the source IPv6 address is 0, that means the packet is a Duplicate Address Detection packet, and the firewall ignores the ND solicitation.
- 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.
- 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.

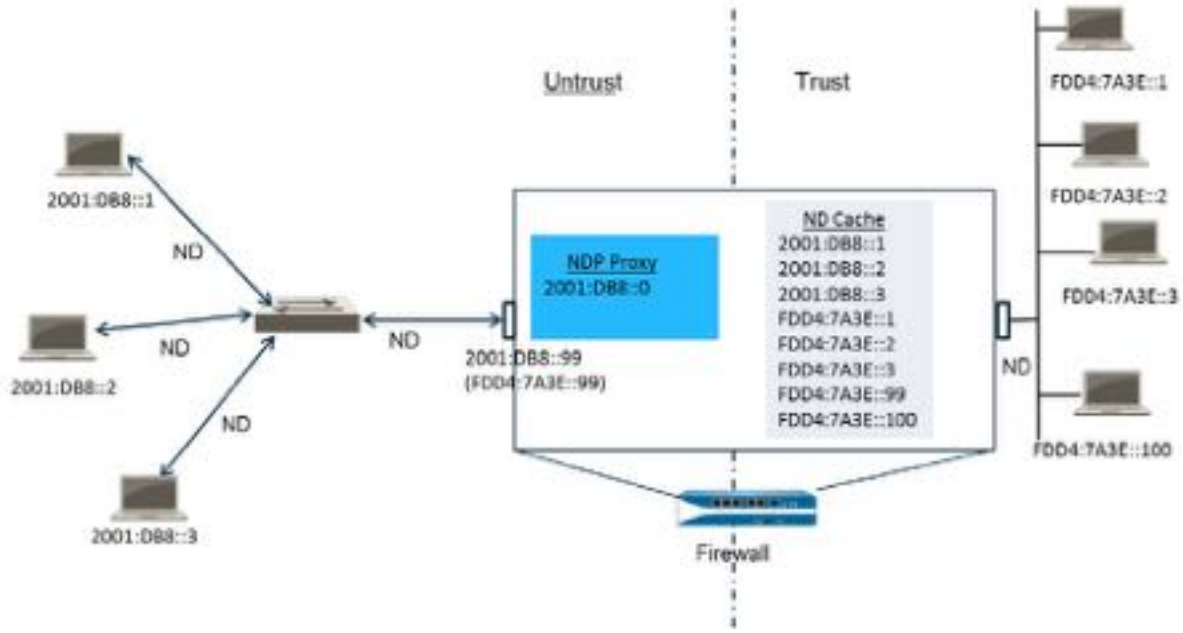
In order to successfully support NDP, the firewall does not perform NDP Proxy for the following:

- 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 illustrates how NPTv6 and NDP Proxy function together.



- The ND Cache in NPTv6 Example
- The NDP Proxy in NPTv6 Example
- The NPTv6 Translation in NPTv6 Example
- Neighbors in the ND Cache are Not Translated

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.

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**. 2. On the **General** tab, enter a descriptive **Name** for the NPTv6 policy rule. 3. (Optional) Enter a **Description** and **Tag**. 4. For **NAT Type**, select **NPTv6**.

 © Palo Alto Networks, Inc.

Configure an NPTv6 Policy (Continued)		
Step 2	Specify the match criteria for incoming packets; packets that match all of the criteria are subject to the NPTv6 translation. 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. 2. Enter the **Destination Zone** to which the policy applies. 3. (Optional) Select a **Destination Interface**. 4. (Optional) Select a **Service** to restrict what type of packets are translated. 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**: • 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. • The IPv6 address cannot have an interface identifier (host) portion defined. • The range of supported prefix lengths is /32 to /64. • The **Source Address** and **Destination Address** cannot both be set to **Any**. 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.
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**. 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. 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. 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. If you enable **Bi-directional** translation, it is very important to make sure you have Security policy rules in place to control the traffic in both directions. Without such policy rules, **Bi-directional** translation allows packets to be automatically translated in both directions, which you might not want. 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. 5. Click **OK**.

Configure an NPTv6 Policy (Continued)		
Step 4	Configure NDP Proxy. 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. 2. On the **Advanced** > **NDP Proxy** tab, select **Enable NDP Proxy** and click **Add**. 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.  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. 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**.

© Palo Alto Networks, Inc.

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.

For information about ECMP path selection when an HA peer fails, see ECMP in Active/Active HA Mode.

The following sections describe ECMP and how to configure it.

▲ ECMP Load-Balancing Algorithms
▲ ECMP Platform, Interface, and IP Routing Support
▲ Configure ECMP on a Virtual Router
▲ Enable ECMP for Multiple BGP Autonomous Systems
▲ Verify ECMP

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.

  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).



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.

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.

 © Palo Alto Networks, Inc.

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.

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.

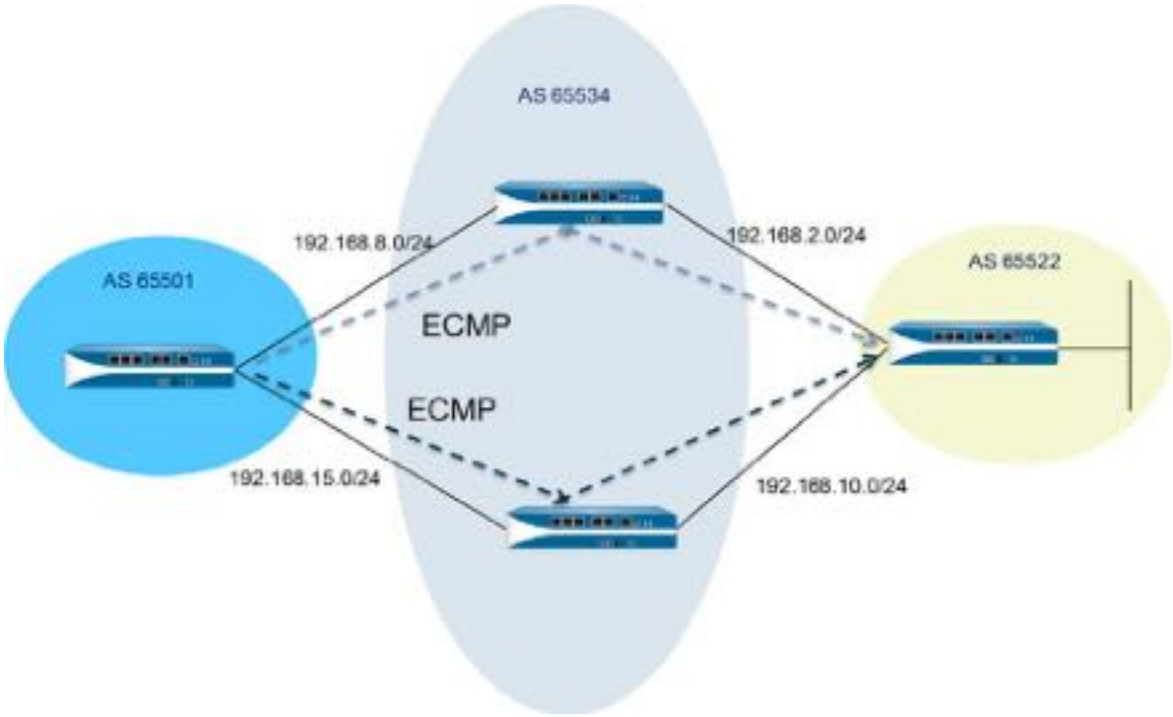
Configure ECMP on a Virtual Router		
Step 1	Enable ECMP for a virtual router.	1. Select **Network > Virtual Routers** and select the virtual router on which to enable ECMP. 2. Select **Router Settings > ECMP** and select **Enable**.
Step 2	(Optional) Enable symmetric return of packets from server to client.	(Optional) Select **Symmetric Return** 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: • **IP Modulo** (default)—Uses a hash of the source and destination IP addresses in the packet header to determine which ECMP route to use. • **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. • **Balanced Round Robin**—Uses round robin among the ECMP paths and re-balances paths when the number of paths changes. • **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	(**IP Hash** only) Configure IP Hash options.	If you selected **IP Hash** as the **Method**: 1. Select **Use Source/Destination Ports** if you want to use source or destination port numbers in the **IP Hash** calculation. 2. Enter a **Hash Seed** value (an integer with a maximum of nine 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). The higher the weight, the more often that equal-cost path will be selected for a new session. Give higher speed links a higher weight than a slower links so that more of the ECMP traffic goes over the faster link. 1. Create an ECMP group by clicking **Add** and selecting an **Interface** from the drop-down. 2. **Add** the other interfaces in the ECMP group. 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**. 2. At the ECMP Configuration Change prompt, click **Yes** to restart the virtual router. Restarting the virtual router might cause existing sessions to be terminated. This message displays only if you are modifying an existing virtual router with ECMP.
Step 8	Save the configuration.	**Commit** the configuration.

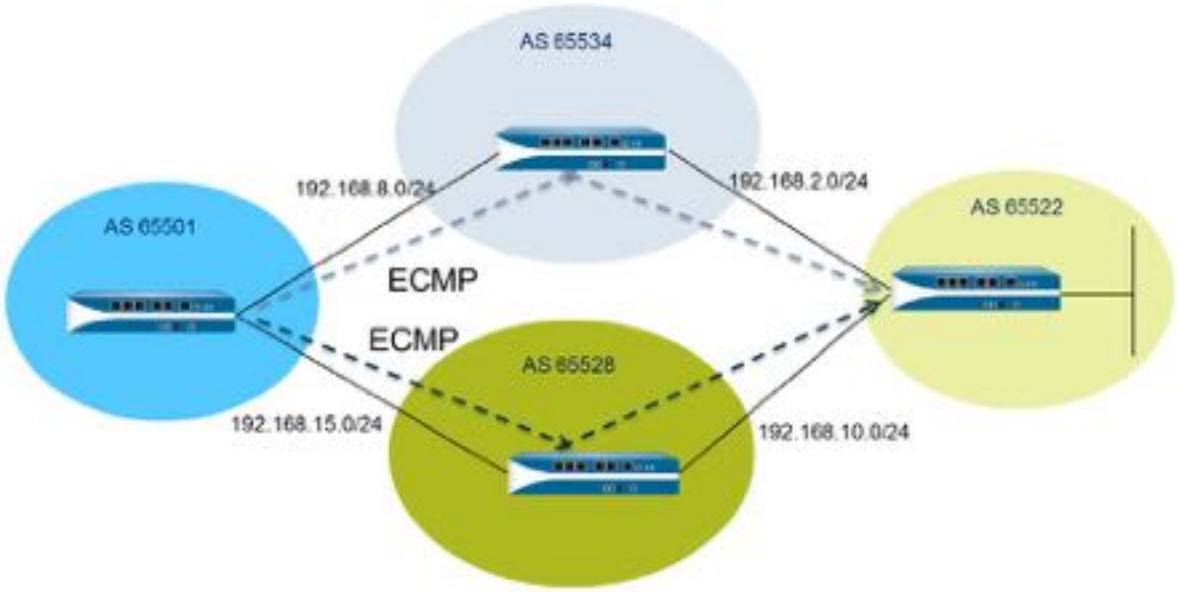
Enable ECMP for Multiple BGP Autonomous Systems

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.

 © Palo Alto Networks, Inc.



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.

Enable ECMP for BGP Autonomous Systems (Continued)		
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. 2. Select **BGP > Advanced** and select **ECMP Multiple AS Support**.
Step 3	Save the configuration.	Click **OK** and **Commit** the configuration.

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. To verify ECMP, use the following procedure to look at the FIB and confirm that some routes are equal-cost multiple paths.

Confirm That Routes Are Equal-Cost Multiple Paths	
Step 1	Select **Network > Virtual Routers**.
Step 2	In the row of the virtual router for which you enabled ECMP, click **More Runtime Stats**.
Step 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. An asterisk (*****) denotes the preferred path for the ECMP group.

 © Palo Alto Networks, Inc.

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

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	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.

 © Palo Alto Networks, Inc.

The following table lists the optional TLVs that the Palo Alto Networks firewall supports:

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: • An L3 interface is advertised with router (bit 6) capability and the “other” bit (bit 1). • An L2 interface is advertised with MAC Bridge (bit 3) capability and the “other” bit (bit 1). • A virtual wire interface is advertised with Repeater (bit 2) capability and the “other” bit (bit 1).
Management Address	8	One or more IP addresses used for firewall management, as follows: • IP address of the management (MGT) interface • IPv4 and/or IPv6 address of the interface • Loopback address • User-defined address entered in the management address field If no management IP address is provided, the default is the MAC address of the transmitting interface. 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). 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. 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.

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 **Enable**.
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. 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. 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. 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. 5. Click **OK**.

 © Palo Alto Networks, Inc.

Configure LLDP (Continued)		
Step 3	Create an LLDP profile. For descriptions of the optional TLVs, see Supported TLVs in LLDP.	1. Select **Network > Network Profiles > LLDP Profile** and click **Add**. 2. Enter a **Name** for the LLDP profile. 3. For **Mode**, select **transmit-receive** (default), **transmit-only**, or **receive-only**. 4. Select **SNMP Syslog Notification** 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**. 5. For Optional TLVs, select the TLVs you want transmitted: • **Port Description** • **System Name** • **System Description** • **System Capabilities** 6. (Optional) Select **Management Address** to add one or more management addresses and **Add** a **Name**. 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. 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. 9. Click **OK**. 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. 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. 2. Select **Advanced > LLDP**. 3. Select **Enable LLDP** to assign an LLDP profile to the interface. 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. If you want to create a new profile, click **LLDP Profile** and follow the instructions steps above. 5. Click **OK**.
Step 5	Save the configuration.	Click **Commit**.

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**. • On the LLDP General screen, **Enable** indicates whether LLDP is enabled or not. – If LLDP is enabled, the configured global settings (Transmit Interval, Transmit Delay, Hold Time Multiple, and Notification Interval) are displayed. – If LLDP is not enabled, the default values of the global settings are displayed. For descriptions of these values, see (Optional) Change LLDP global settings.
Step 2 View the LLDP status information.	1. Select the **Status** tab. 2. (Optional) Enter a filter to restrict the information that is displayed. Interface Information: • **Interface**—Name of the interfaces that have LLDP profiles assigned to them. • **LLDP**—LLDP status: enabled or disabled. • **Mode**— LLDP mode of the interface: Tx/Rx, Tx Only, or Rx Only. • **Profile**—Name of the profile assigned to the interface. Transmission Information: • **Total Transmitted**—Count of LLDPDUs transmitted out the interface. • **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. Received Information: • **Total Received**—Count of LLDP frames received on the interface. • **Dropped TLV**—Count of LLDP frames discarded upon receipt. • **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. • **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. • **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. 2. (Optional) Enter a filter to restrict the information being displayed. • **Local Interface**—Interface on the firewall that detected the neighboring device. • **Remote Chassis ID**—Chassis ID of the peer. The MAC address will be used. • **Port ID**—Port ID of the peer. • **Name**—Name of peer. • **More info**—Provides the following remote peer details, which are based on the Mandatory and Optional TLVs: – Chassis Type: MAC address. – MAC Address: MAC address of the peer. – System Name: Name of the peer. – System Description: Description of the peer. – Port Description: Port description of the peer. – Port Type: Interface name. – Port ID: The firewall uses the interface's ifname. – System Capabilities: Capabilities of the system. O=Other, P=Repeater, B=Bridge, W=Wireless-LAN, R=Router, T=Telephone – Enabled Capabilities: Capabilities enabled on the peer. – 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. 2. Click **Clear LLDP Statistics** at the bottom of the screen.

BFD

The firewall supports Bidirectional Forwarding Detection (BFD), a protocol that recognizes a failure in the bidirectional path between two routing peers. BFD failure detection is extremely fast, providing for a faster failover than can be achieved by link monitoring or frequent dynamic routing health checks, such as Hello packets or heartbeats. Mission-critical data centers and networks that require high availability and extremely fast failover need the extremely fast failure detection that BFD provides.

- BFD Overview
- Configure BFD
- Reference: BFD Details

BFD Overview

When you enable BFD, BFD establishes a session from one endpoint (the firewall) to its BFD peer at the endpoint of a link using a three-way handshake. Control packets perform the handshake and negotiate the parameters configured in the BFD profile, including the minimum intervals at which the peers can send and receive control packets. BFD control packets for both IPv4 and IPv6 are transmitted over UDP port 3784. BFD control packets for multihop support are transmitted over UDP port 4784. BFD control packets transmitted over either port are encapsulated in the UDP packets.

After the BFD session is established, the Palo Alto Networks implementation of BFD operates in asynchronous mode, meaning both endpoints send each other control packets (which function like Hello packets) at the negotiated interval. If a peer does not receive a control packet within the detection time (calculated as the negotiated transmit interval multiplied by a Detection Time Multiplier), the peer considers the session down. (The firewall does not support demand mode, in which control packets are sent only if necessary rather than periodically.)

When you enable BFD for a static route and a BFD session between the firewall and the BFD peer fails, the firewall removes the failed route from the RIB and FIB tables and allows an alternate path with a lower priority to take over. When you enable BFD for a routing protocol, BFD notifies the routing protocol to switch to an alternate path to the peer. Thus, the firewall and BFD peer reconverge on a new path.

A BFD profile allows you to Configure BFD settings and apply them to one or more routing protocols or static routes on the firewall. If you enable BFD without configuring a profile, the firewall uses its default BFD profile (with all of the default settings). You cannot change the default BFD profile.

When an interface is running multiple protocols that use different BFD profiles, BFD uses the profile having the lowest **Desired Minimum Tx Interval**. See BFD for Dynamic Routing Protocols.

Active/passive HA peers synchronize BFD configurations and sessions; active/active HA peers do not.

BFD is standardized in RFC 5880. PAN-OS does not support all components of RFC 5880; see Non-Supported RFC Components of BFD.

PAN-OS also supports RFC 5881, Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop). In this case, BFD tracks a single hop between two systems that use IPv4 or IPv6, so the two systems are directly connected to each other. BFD also tracks multiple hops from peers connected by BGP. PAN-OS follows BFD encapsulation as described in RFC 5883, Bidirectional Forwarding Detection (BFD) for Multihop Paths. However, PAN-OS does not support authentication.

- BFD Platform, Interface, and Client Support

 © Palo Alto Networks, Inc.

▲ Non-Supported RFC Components of BFD
▲ BFD for Static Routes
▲ BFD for Dynamic Routing Protocols

BFD Platform, Interface, and Client Support

PAN-OS supports BFD on PA-3000 Series, PA-5000 Series, PA-7000 Series, and VM-Series firewalls. Each platform supports a maximum number of BFD sessions, as listed in the Product Selection tool.

BFD runs on physical Ethernet, Aggregated Ethernet (AE), VLAN, and tunnel interfaces (site-to-site VPN and LSVPN), and on Layer 3 subinterfaces.

Supported BFD clients are:

- Static routes (IPv4 and IPv6) consisting of a single hop
- OSPFv2 and OSPFv3 (interface types include broadcast, point-to point, and point-to-multipoint)
- BGP IPv4 (IBGP, EBGP) consisting of a single hop or multiple hops
- RIP (single hop)

Non-Supported RFC Components of BFD

- Demand mode
- Authentication
- Sending or receiving Echo packets; however, the firewall will pass Echo packets that arrive on a virtual wire or tap interface. (BFD Echo packets have the same IP address for the source and destination.)
- Poll sequences
- Congestion control

BFD for Static Routes

To use BFD on a static route, both the firewall and the peer at the opposite end of the static route must support BFD sessions. A static route can have a BFD profile only if the **Next Hop** type is **IP Address**.

If an interface is configured with more than one static route to a peer (the BFD session has the same source IP address and same destination IP address), a single BFD session automatically handles the multiple static routes. This behavior reduces BFD sessions. If the static routes have different BFD profiles, the profile with the smallest **Desired Minimum Tx Interval** takes effect.

In a deployment where you want to configure BFD for a static route on a DHCP or PPPoE client interface, you must perform two commits. Enabling BFD for a static route requires that the **Next Hop** type must be **IP Address**. But at the time of a DHCP or PPPoE interface commit, the interface IP address and next hop IP address (default gateway) are unknown.

You must first enable a DHCP or PPPoE client for the interface, perform a commit, and wait for the DHCP or PPPoE server to send the firewall the client IP address and default gateway IP address. Then you can configure the static route (using the default gateway address of the DHCP or PPPoE client as the next hop), enable BFD, and perform a second commit.

BFD for Dynamic Routing Protocols

In addition to BFD for static routes, the firewall supports BFD for the BGP, OSPF, and RIP routing protocols.



The Palo Alto Networks implementation of multihop BFD follows the encapsulation portion of RFC 5883, Bidirectional Forwarding Detection (BFD) for Multihop Paths but does not support authentication. A workaround is to configure BFD in a VPN tunnel for BGP. The VPN tunnel can provide authentication without the duplication of BFD authentication.

When you enable BFD for OSPFv2 or OSPFv3 broadcast interfaces, OSPF establishes a BFD session only with its Designated Router (DR) and Backup Designated Router (BDR). On point-to-point interfaces, OSPF establishes a BFD session with the direct neighbor. On point-to-multipoint interfaces, OSPF establishes a BFD session with each peer.

The firewall does not support BFD on an OSPF or OSPFv3 virtual link.

Each routing protocol can have independent BFD sessions on an interface. Alternatively, two or more routing protocols (BGP, OSPF, and RIP) can share a common BFD session for an interface.

When you enable BFD for multiple protocols on the same interface, and the source IP address and destination IP address for the protocols are also the same, the protocols share a single BFD session, thus reducing both dataplane overhead (CPU) and traffic load on the interface. If you configure different BFD profiles for these protocols, only one BFD profile is used: the one that has the lowest **Desired Minimum Tx Interval**. If the profiles have the same **Desired Minimum Tx Interval**, the profile used by the first created session takes effect. In the case where a static route and OSPF share the same session, because a static session is created right after a commit, while OSPF waits until an adjacency is up, the profile of the static route takes effect.

The benefit of using a single BFD session in these cases is that this behavior uses resources more efficiently. The firewall can use the saved resources to support more BFD sessions on different interfaces or support BFD for different source IP and destination IP address pairs.

IPv4 and IPv6 on the same interface always create different BFD sessions, even though they can use the same BFD profile.



If you implement both BFD for BGP and HA path monitoring, Palo Alto Networks recommends you not implement BGP Graceful Restart. When the BFD peer's interface fails and path monitoring fails, BFD can remove the affected routes from the routing table and synchronize this change to the passive HA firewall before Graceful Restart can take effect. If you decide to implement BFD for BGP, Graceful Restart for BGP, and HA path monitoring, you should configure BFD with a larger Desired Minimum Tx Interval and larger Detection Time Multiplier than the default values.

 © Palo Alto Networks, Inc.

Configure BFD

This task assumes you have performed the following prerequisites:

- Configured a virtual router.
- Configured one or more static routes if you are applying BFD to static routes.
- Configured a routing protocol (BGP, OSPF, OSPFv3, or RIP) if you are applying BFD to a routing protocol.



The effectiveness of your BFD implementation depends on a variety of factors, such as traffic loads, network conditions, how aggressive your BFD settings are, and how busy the dataplane is.

Configure BFD	
Step 1 Create a BFD profile.  If you change a setting in a BFD profile that an existing BFD session is using and you commit the change, before the firewall deletes that BFD session and recreates it with the new setting, the firewall sends a BFD packet with the local state set to admin down. The peer device may or may not flap the routing protocol or static route, depending on the peer's implementation of RFC 5882, Section 3.2.	1. Select **Network > Network Profiles > BFD Profile** and **Add** a **Name** for the BFD profile. The name is case-sensitive and must be unique on the firewall. Use only letters, numbers, spaces, hyphens, and underscores. 2. Select the **Mode** in which BFD operates: • **Active**—BFD initiates sending control packets to peer (default). At least one of the BFD peers must be Active; both can be Active. • **Passive**—BFD waits for peer to send control packets and responds as required. 3. Enter the **Desired Minimum Tx Interval (ms)**. This is the minimum interval, in milliseconds, at which you want the BFD protocol (referred to as BFD) to send BFD control packets; you are thus negotiating the transmit interval with the peer. Minimum on PA-7000 and PA-5000 Series firewalls is 50; minimum on PA-3000 Series firewall is 100; minimum on VM-Series firewall is 200. Maximum is 2000; default is 1000.  If you have multiple routing protocols that use different BFD profiles on the same interface, configure the BFD profiles with the same **Desired Minimum Tx Interval**. 4. Enter the **Required Minimum Rx Interval (ms)**. This is the minimum interval, in milliseconds, at which BFD can receive BFD control packets. Minimum on PA-7000 and PA-5000 Series firewalls is 50; minimum on PA-3000 Series firewall is 100; minimum on VM-Series firewall is 200. Maximum is 2000; default is 1000. 5. Enter the **Detection Time Multiplier**. The transmit interval (negotiated from the **Desired Minimum Tx Interval**) multiplied by the **Detection Time Multiplier** equals the detection time. If BFD does not receive a BFD control packet from its peer before the detection time expires, a failure has occurred. Range is 2-50; default is 3. For example, a transmit interval of 300 ms x 3 (Detection Time Multiplier) = 900 ms detection time.  When configuring a BFD profile, take into consideration that the firewall is a session-based device typically at the edge of a network or data center and may have slower links than a dedicated router. Therefore, the firewall likely needs a longer interval and a higher multiplier than the fastest settings allowed. A detection time that is too short can cause false failure detections when the issue is really just traffic congestion.

 © Palo Alto Networks, Inc.

Configure BFD (Continued)		
		6. Enter the **Hold Time (ms)**. This is the delay, in milliseconds, after a link comes up before BFD transmits BFD control packets. **Hold Time** applies to BFD Active mode only. If BFD receives BFD control packets during the **Hold Time**, it ignores them. Range is 0-120000. The default is 0, which means no transmit **Hold Time** is used; BFD sends and receives BFD control packets immediately after the link is established. 7. (Optional) For a BGP IPv4 implementation only, configure hop-related settings for the BFD profile: • Select **Multihop** to enable BFD over BGP multihop. • Enter the **Minimum Rx TTL**.This is the minimum Time-to-Live value (number of hops) BFD will accept (receive) in a BFD control packet when BGP supports multihop BFD. (Range is 1-254; there is no default). The firewall drops the packet if it receives a smaller TTL than its configured **Minimum Rx TTL**.For example, if the peer is 5 hops away, and the peer transmits a BFD packet with a TTL of 100 to the firewall, and if the **Minimum Rx TTL** for the firewall is set to 96 or higher, the firewall drops the packet. 8. Click **OK**.
Step 2	(Optional) Enable BFD for a static route. Both the firewall and the peer at the opposite end of the static route must support BFD sessions.	1. Select **Network > Virtual Routers** and select the virtual router where the static route is configured. 2. Select the **Static Routes** tab. 3. Select the **IPv4** or **IPv6** tab. 4. Select the static route where you want to apply BFD. 5. Select an **Interface** (even if you are using a DHCP address). The **Interface** setting cannot be **None**. 6. For **Next Hop**, select **IP Address** and enter the IP address if not already specified. 7. For **BFD Profile**, select one of the following: • **default**—Uses only default settings. • A BFD profile you configured—See Create a BFD profile. • **New BFD Profile**—Allows you to Create a BFD profile. Selecting **None (Disable BFD)** disables BFD for this static route. 8.Click **OK**. A BFD column on the **IPv4** or **IPv6** tab indicates the BFD profile configured for the static route.

Configure BFD (Continued)	
Step 3 (Optional) Enable BFD for all BGP interfaces or for a single BGP peer.  If you enable or disable BFD globally, all interfaces running BGP will be taken down and brought back up with the BFD function. This can disrupt all BGP traffic. When you enable BFD on the interface, the firewall stops the BGP connection to the peer to program BFD on the interface. The peer device sees the BGP connection drop, which can result in a reconvergence. Enable BFD for BGP interfaces during an off-peak time when a reconvergence will not impact production traffic.  If you implement both BFD for BGP and HA path monitoring, Palo Alto Networks recommends you not implement BGP Graceful Restart. When the BFD peer's interface fails and path monitoring fails, BFD can remove the affected routes from the routing table and synchronize this change to the passive HA firewall before Graceful Restart can take effect. If you decide to implement BFD for BGP, Graceful Restart for BGP, and HA path monitoring, you should configure BFD with a larger Desired Minimum Tx Interval and larger Detection Time Multiplier than the default values.	1. Select **Network > Virtual Routers** and select the virtual router where BGP is configured. 2. Select the **BGP** tab. 3. (Optional) To apply BFD to all BGP interfaces on the virtual router, in the **BFD** drop-down, select one of the following and click **OK**: • **default**—Uses only default settings. • A BFD profile you configured—See Create a BFD profile. • **New BFD Profile**—Allows you to Create a BFD profile.  Selecting **None (Disable BFD)** disables BFD for all BGP interfaces on the virtual router; you cannot enable BFD for a single BGP interface. 4. (Optional) To enable BFD for a single BGP peer interface (thereby overriding the **BFD** setting for BGP as long as it is not disabled), perform the following tasks: a. Select the **Peer Group** tab. b. Select a peer group. c. Select a peer. d. In the **BFD** drop-down, select one of the following: – **default**—Uses only default settings. – **Inherit-vr-global-setting** (default)—The BGP peer inherits the BFD profile that you selected globally for BGP for the virtual router. – A BFD profile you configured—See Create a BFD profile.  Selecting **Disable BFD** disables BFD for the BGP peer. e. Click **OK**. 5. Click **OK**. A BFD column on the BGP - Peer Group/Peer list indicates the BFD profile configured for the interface.

 © Palo Alto Networks, Inc.

Configure BFD (Continued)	
Step 4 (Optional) Enable BFD for OSPF or OSPFv3 globally or for an OSPF interface.	1. Select **Network > Virtual Routers** and select the virtual router where OSPF or OSPFv3 is configured. 2. Select the **OSPF** or **OSPFv3** tab. 3. (Optional) In the **BFD** drop-down, select one of the following to enable BFD for all OSPF or OSPFv3 interfaces and click **OK**: • **default**—Uses only default settings. • A BFD profile you configured—See Create a BFD profile. • **New BFD Profile**—Allows you to Create a BFD profile.  Selecting **None (Disable BFD)** disables BFD for all OSPF interfaces on the virtual router; you cannot enable BFD for a single OSPF interface. 4. (Optional) To enable BFD on a single OSPF peer interface (and thereby override the **BFD** setting for OSPF, as long as it is not disabled), perform the following tasks: a. Select the **Areas** tab and select an area. b. On the **Interface** tab, select an interface. c. In the **BFD** drop-down, select one of the following to configure BFD for the specified OSPF peer: – **default**—Uses only default settings. – **Inherit-vr-global-setting** (default)—OSPF peer inherits the **BFD** setting for OSPF or OSPFv3 for the virtual router. – A BFD profile you configured—See Create a BFD profile.  Selecting **Disable BFD** disables BFD for the OSPF or OSPFv3 interface. d. Click **OK**. 5. Click **OK**. A BFD column on the OSPF **Interface** tab indicates the BFD profile configured for the interface.

Configure BFD (Continued)		
Step 5	(Optional) Enable BFD for RIP globally or for a single RIP interface.	1. Select **Network > Virtual Routers** and select the virtual router where RIP is configured. 2. Select the **RIP** tab. 3. (Optional) In the **BFD** drop-down, select one of the following to enable BFD for all RIP interfaces on the virtual router and click **OK**: • **default**—Uses only default settings. • A BFD profile you configured—See Create a BFD profile. • **New BFD Profile**—Allows you to Create a BFD profile. Selecting **None (Disable BFD)** disables BFD for all RIP interfaces on the virtual router; you cannot enable BFD for a single RIP interface. 4. (Optional) To enable BFD for a single RIP interface (and thereby override the **BFD** setting for RIP, as long as it is not disabled), perform the following tasks: a. Select the **Interfaces** tab and select an interface. b. In the **BFD** drop-down, select one of the following: – **default**—Uses only default settings). – **Inherit-vr-global-setting** (default)—RIP interface inherits the BFD profile that you selected for RIP globally for the virtual router. – A BFD profile you configured—See Create a BFD profile. Selecting **None (Disable BFD)** disables BFD for the RIP interface. c. Click **OK**. 5. Click **OK**. The BFD column on the **Interface** tab indicates the BFD profile configured for the interface.
Step 6	Save the configuration.	Click **Commit**.
Step 7	View BFD summary and details.	1. Select **Network > Virtual Routers**, find the virtual router you are interested in, and click **More Runtime Stats**. 2. Select the **BFD Summary Information** tab to see summary information, such as BFD state and run-time statistics. 3. (Optional) Select **details** in the row of the interface you are interested in to view Reference: BFD Details.
Step 8	Monitor BFD profiles referenced by a routing configuration; monitor BFD statistics, status, and state.	Use the following CLI operational commands: • `show routing bfd active-profile [<name>]` • `show routing bfd details [interface <name>] [local-ip <ip>] [multihop] [peer-ip <ip>] [session-id] [virtual-router <name>]` • `show routing bfd drop-counters session-id <session-id>` • `show counter global | match bfd`
Step 9	(Optional) Clear BFD transmit, receive, and drop counters.	• `clear routing bfd counters session-id all | <1-1024>`

 © Palo Alto Networks, Inc.

Configure BFD (Continued)	
Step 10 (Optional) Clear BFD sessions for debugging.	• `clear routing bfd session-state session-id all | <1-1024>`

Reference: BFD Details

To see the following information for a virtual router, you can View BFD summary and details.

Name	Value (Example)	Description
Session ID	1	ID number of the BFD session.
Interface	ethernet1/12	Interface you selected where BFD is running.
Protocol	STATIC(IPV4) OSPF	Static route (IP address family of static route) and/or dynamic routing protocol that is running BFD on the interface.
Local IP Address	10.55.55.2	IP address of interface.
Neighbor IP Address	10.55.55.1	IP address of BFD neighbor.
BFD Profile	default *(This BFD session has multiple BFD profiles. Lowest 'Desired Minimum Tx Interval (ms)' is used to select the effective profile.)	Name of BFD profile applied to the interface. Because the sample interface has both a static route and OSPF running BFD with different profiles, the firewall uses the profile with the lowest '**Desired Minimum Tx Interval**.' In this example, the profile used is the default profile.
State (local/remote)	up/up	BFD states of the local and remote BFD peers. Possible states are admin down, down, init, and up.
Up Time	2h 36m 21s 419ms	Length of time BFD has been up (hours, minutes, seconds, and milliseconds).
Discriminator (local/remote)	1391591427/ 1	Discriminators for local and remote BFD peers.
Mode	Active	Mode in which BFD is configured on the interface: Active or Passive.
Demand Mode	Disabled	PAN-OS does not support BFD Demand Mode, so it is always in Disabled state.
Multihop	Disabled	BFD multihop: Enabled or Disabled.
Multihop TTL		TTL of multihop; range is 1-254. Field is empty if Multihop is disabled.
Local Diag Code	0 (No Diagnostic)	Diagnostic codes indicating the reason for the local system's last change in state: 0—No Diagnostic 1—Control Detection Time Expired 2—Echo Function Failed 3—Neighbor Signaled Session Down 4—Forwarding Plane Reset 5—Path Down 6—Concatenated Path Down 7—Administratively Down 8—Reverse Concatenated Path Down

 © Palo Alto Networks, Inc.

Name	Value (Example)	Description
Last Received Remote Diag Code	0 (No Diagnostic)	Diagnostic code last received from BFD peer.
Transmit Hold Time	0ms	Hold time (in milliseconds) after a link comes up before BFD transmits BFD control packets. A hold time of 0ms means to transmit immediately. Range is 0-120000ms.
Received Min Rx Interval	1000ms	Minimum Rx interval received from the peer; the interval at which the BFD peer can receive control packets. Maximum is 2000ms.
Negotiated Transmit Interval	1000ms	Transmit interval (in milliseconds) that the BFD peers have agreed to send BFD control packets to each other. Maximum is 2000ms.
Received Multiplier	3	Detection time multiplier value received from the BFD peer. The Transmit Time multiplied by the Multiplier equals the detection time. If BFD does not receive a BFD control packet from its peer before the detection time expires, a failure has occurred. Range is 2-50.
Detect Time (exceeded)	3000ms(0)	Calculated detection time (Negotiated Transmit Interval multiplied by Multiplier) and the number of milliseconds the detection time is exceeded.
Tx Control Packets (last)	9383 (420ms ago)	Number of BFD control packets transmitted (and length of time since BFD transmitted the most recent control packet).
Rx Control Packets (last)	9384 (407ms ago)	Number of BFD control packets received (and length of time since BFD received the most recent control packet).
Agent Data Plane	Slot 1 - DP 0	On PA-7000 Series firewalls, the dataplane CPU that is assigned to handle packets for this BFD session.
Errors	0	Number of BFD errors.
Last Packet Causing State Change		
Version	1	BFD version.
Poll Bit	0	BFD poll bit; 0 indicates not set.
Desired Min Tx Interval	1000ms	Desired minimum transmit interval of last packet causing state change.
Required Min Rx Interval	1000ms	Required minimum receive interval of last packet causing state change.
Detect Multiplier	3	Detect Multiplier of last packet causing state change.
My Discriminator	1	Remote discriminator. A discriminator is a unique, nonzero value the peers use to distinguish multiple BFD sessions between them.
Your Discriminator	1391591427	Local discriminator. A discriminator is a unique, nonzero value the peers use to distinguish multiple BFD sessions between them.
Diagnostic Code	0 (No Diagnostic)	Diagnostic code of last packet causing state change.

Name	Value (Example)	Description
Length	24	Length of BFD control packet in bytes.
Demand Bit	0	PAN-OS does not support BFD Demand mode, so Demand Bit is always set to 0 (disabled).
Final Bit	0	PAN-OS does not support the Poll Sequence, so Final Bit is always set to 0 (disabled).
Multipoint Bit	0	This bit is reserved for future point-to-multipoint extensions to BFD. It must be zero on both transmit and receipt.
Control Plane Independent Bit	1	• If set to 1, the transmitting system's BFD implementation does not share fate with its control plane (i.e., BFD is implemented in the forwarding plane and can continue to function through disruptions in the control plane). In PAN-OS, this bit is always set to 1. • If set to 0, the transmitting system's BFD implementation shares fate with its control plane.
Authentication Present Bit	0	PAN-OS does not support BFD Authentication, so the Authentication Present Bit is always set to 0.
Required Min Echo Rx Interval	0ms	PAN-OS does not support the BFD Echo function, so this will always be 0ms.

 © 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
- Best Practice Internet Gateway Security Policy
- 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 an External Dynamic 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

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. For more details, see DoS Protection Profiles.

 © Palo Alto Networks, Inc.

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, individual 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.

All traffic passing through the firewall is matched against a session and each session is matched against a Security policy rule. When a session match occurs, the firewall applies the matching Security policy rule 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 policy 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.

- Components of a Security Policy Rule
- Security Policy Actions
- Create a Security Policy Rule

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: • **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 rule 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. • **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. • **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 *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. For more details, see Security Policy Actions.

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).

 © Palo Alto Networks, Inc.

Optional Field	Description (Continued)
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. Although you can manually configure URL categories on your firewall, to take advantage of the dynamic URL categorization updates available on the Palo Alto Networks firewalls, you must purchase a URL filtering license.  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 Policy 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.  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 firewall 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.

Security Policy Actions

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. 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.
	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.

Create a Security Policy Rule

Create a Security Policy Rule		
Step 1	(Optional) Delete the default Security policy rule.	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 conventions.
Step 2	Add a rule.	1. Select **Policies > Security** and click **Add**. 2. Enter a descriptive **Name** for the rule in the **General** tab. 3. Select a **Rule Type**.
Step 3	Define the matching criteria for the source fields in the packet.	1. In the **Source** tab, select a **Source Zone**. 2. Specify a **Source IP Address** or leave the value set to **any**. 3. Specify a Source **User** or leave the value set to **any**.

 © Palo Alto Networks, Inc.

Create a Security Policy Rule (Continued)		
Step 4	Define the matching criteria for the destination fields in the packet.	4. In the **Destination** tab, set the **Destination Zone**. 5. Specify a **Destination IP Address** or leave the value set to **any**. As a best practice, consider using address objects in the **Destination Address** field to enable access to specific servers or groups of servers only, particularly for services such as DNS and SMTP that are commonly exploited. By restricting users to specific destination server addresses you can prevent data exfiltration and command and control traffic from establishing communication through techniques such as DNS tunneling.
Step 5	Specify the application the rule will allow or block. As a best practice, always use application-based security policy rules instead of port based rules and always set the Service to application-default unless you are using a more restrictive list of ports than the standard ports for an application.	1. In the **Applications** tab, **Add** the **Application** to safely enable. You can select multiple applications, or use application groups or application filters. 2. In the **Service/URL Category** tab, keep the **Service** set to **application-default** to ensure that any applications the rule allows are only allowed on their standard ports.
Step 6	(Optional) Specify a URL category as match criteria for the rule.	In the **Service/URL Category** tab, select the **URL Category**. If you select a URL category, only web traffic will match the rule and only if the traffic is to the specified category.
Step 7	Define what action you want the firewall to take for traffic that matches the rule.	In the **Actions** tab, select an **Action**. See Security Policy Actions for a description of each action.
Step 8	Configure the log settings.	• By default, the rule is set to **Log at Session End**. You can clear this setting if you don't want any logs generated when traffic matches this rule, or select **Log at Session Start** for more detailed logging. • Select a **Log Forwarding** profile.
Step 9	Attach security profiles to enable the firewall to scan all allowed traffic for threats. See Create Best Practice Security Profiles to learn how to create security profiles that protect your network from both known and unknown threats.	In the **Actions** tab, select **Profiles** from the **Profile Type** drop-down and then select the individual security profiles to attach to the rule. Alternatively, select **Group** from the **Profile Type** drop-down and select a security **Group Profile** to attach.
Step 10	Save the policy rule to the running configuration on the firewall.	Click **Commit**.

Create a Security Policy Rule (Continued)	

Step 11 To verify that you have set up your basic policies effectively, test whether your security policy rules are being evaluated and determine which security policy rule applies to a traffic flow.	To verify the policy rule that matches a flow, use the following CLI command:
	`test security-policy-match source <IP_address> destination <IP_address> destination port <port_number> protocol <protocol_number>`
	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 in the data center with the IP address 208.90.56.11 when it accesses the Microsoft update server:
	`test security-policy-match source 208.80.56.11 destination 176.9.45.70 destination-port 80 protocol 6`
	
	`"Updates-DC to Internet" {`
	`from data_center_applications;`
	`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;`

 © 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. 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. 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).  To view the standard ports used by an application, in **Objects > Applications** search for the application and click the link. A succinct description displays.

Security Profiles

While security policy rules enable you to allow or block traffic on your network, security profiles help you define an *allow but scan* rule, which scans 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 a Basic Security Policy 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.



For recommendations on the best-practice settings for security profiles, see Create Best Practice Security Profiles.

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:

Action	Description
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).

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 rule:

- **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.

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.

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.

 © Palo Alto Networks, Inc.

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.

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.

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.

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 prevents 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.

 © Palo Alto Networks, Inc.

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.

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.



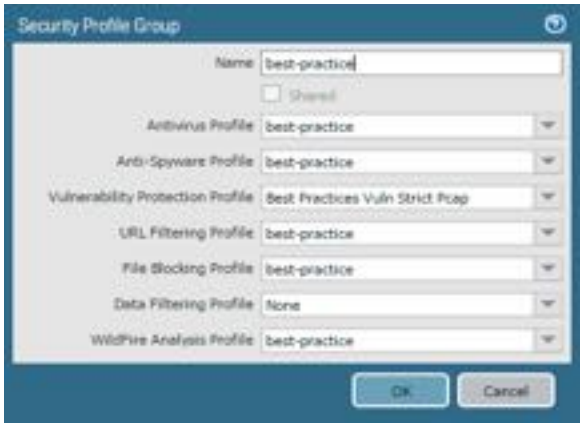
For recommendations on the best-practice settings for security profiles, see Create Best Practice Security Profiles.

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. If you name the group `default`, the firewall will automatically attach it to any new rules you create. This is a time saver if you have a preferred set of security profiles that you want to make sure get attached to every new rule.	1. Select **Objects > Security Profile Groups** and **Add** a new security profile group. 2. Give the profile group a descriptive **Name**, for example, Threats. 3. If the firewall is in Multiple Virtual System Mode, enable the profile to be **Shared** by all virtual systems. 4. Add existing profiles to the group.  5. Click **OK** to save the 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. 2. Select the **Actions** tab. 3. In the Profile Setting section, select **Group** for the **Profile Type**. 4. In the **Group Profile** drop-down, select the group you created (for example, select the best-practice group):  5. Click **OK** to save the policy and **Commit** your changes.
Step 3	Save your changes.	Click **Commit**.

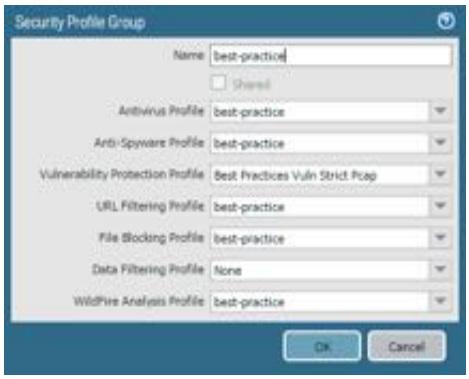
 © Palo Alto Networks, Inc.

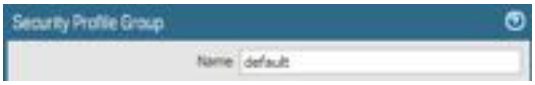
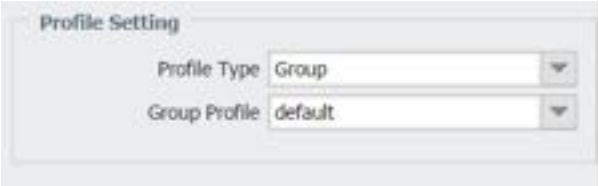
Set Up or Override a Default Security Profile Group

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. 2. Give the profile group a descriptive **Name**, for example, Threats. 3. If the firewall is in Multiple Virtual System Mode, enable the profile to be **Shared** by all virtual systems. 4. Add existing profiles to the group. For details on creating profiles, see Security Profiles.  5. Click **OK** to save the profile group. 6. Add the security profile group to a security policy. 7. **Add** or modify a security policy rule and select the **Actions** tab. 8. Select **Group** for the **Profile Type**. 9. In the **Group Profile** drop-down, select the group you created (for example, select the Threats group):  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. 2. **Name** the security profile group `default`:  3. Click **OK** and **Commit**. 4. Confirm that the default security profile group is included in new security policies by default: a. Select **Policies > Security** and **Add** a new security policy. b. Select the **Actions** tab and view the **Profile Setting** fields:  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.

Best Practice Internet Gateway Security Policy

One of the cheapest and easiest ways for an attacker to gain access to your network is through users accessing the internet. By successfully exploiting an endpoint, an attacker can take hold in your network and begin to move laterally towards the end goal, whether that is to steal your source code, exfiltrate your customer data, or take down your infrastructure. To protect your network from cyberattack and improve your overall security posture, implement a best practice internet gateway security policy. A best practice policy allows you to safely enable applications, users, and content by classifying all traffic, across all ports, all the time.

The following topics describe the overall process for deploying a best practice internet gateway security policy and provide detailed instructions for creating it.

- What Is a Best Practice Internet Gateway Security Policy?
- Why Do I Need a Best Practice Internet Gateway Security Policy?
- How Do I Deploy a Best Practice Internet Gateway Security Policy?
- Identify Whitelist Applications
- Create User Groups for Access to Whitelist Applications
- Decrypt Traffic for Full Visibility and Threat Inspection
- Create Best Practice Security Profiles
- Define the Initial Internet Gateway Security Policy
- Monitor and Fine Tune the Policy Rulebase
- Remove the Temporary Rules
- Maintain the Rulebase

What Is a Best Practice Internet Gateway Security Policy?

A best practice internet gateway security policy has two main security goals:

- **Minimize the chance of a successful intrusion**—Unlike legacy port-based security policies that either block everything in the interest of network security, or enable everything in the interest of your business, a best practice security policy leverages App-ID, User-ID, and Content-ID to ensure safe enablement of applications across all ports, for all users, all the time, while simultaneously scanning all traffic for both known and unknown threats.
- **Identify the presence of an attacker**—A best practice internet gateway security policy provides built-in mechanisms to help you identify gaps in the rulebase and detect alarming activity and potential threats on your network.

To achieve these goals, the best practice internet gateway security policy uses application-based rules to allow access to whitelisted applications by user, while scanning all traffic to detect and block all known threats, and send unknown files to WildFire to identify new threats and generate signatures to block them:



The best practice policy is based on the following methodologies. The best practice methodologies ensure detection and prevention at multiple stages of the attack life cycle.

Best Practice Methodology	Why is this important?
Inspect All Traffic for Visibility	Because you cannot protect against threats you cannot see, you must make sure you have full visibility into all traffic across all users and applications all the time. To accomplish this: • Deploy GlobalProtect to extend the next-generation security platform to users and devices no matter where they are located. • Enable SSL decryption so the firewall can inspect encrypted traffic (SSL/TLS traffic flows account for 40% or more of the total traffic on a typical network today). • Enable User-ID to map application traffic and associated threats to users/devices. The firewall can then inspect all traffic—inclusive of applications, threats, and content—and tie it to the user, regardless of location or device type, port, encryption, or evasive techniques employed using the native App-ID, Content-ID, and User-ID technologies. Complete visibility into the applications, the content, and the users on your network is the first step toward informed policy control.
Reduce the Attack Surface	After you have context into the traffic on your network—applications, their associated content, and the users who are accessing them—create application-based Security policy rules to allow those applications that are critical to your business and additional rules to block all high-risk applications that have no legitimate use case. To further reduce your attack surface, attach File Blocking and URL Filtering profiles to all rules that allow application traffic to prevent users from visiting threat-prone web sites and prevent them from uploading or downloading dangerous file types (either knowingly or unknowingly).
Prevent Known Threats	Enable the firewall to scan all allowed traffic for known threats by attaching security profiles to all allow rules to detect and block network and application layer vulnerability exploits, buffer overflows, DoS attacks, and port scans, known malware variants, (including those hidden within compressed files or compressed HTTP/HTTPS traffic). To enable inspection of encrypted traffic, enable SSL decryption.

 © Palo Alto Networks, Inc.

Best Practice Methodology	Why is this important?
Detect Unknown Threats	Forward all unknown files to WildFire for analysis. WildFire identifies unknown or targeted malware (also called *advanced persistent threats* or *APTs*) hidden within files by directly observing and executing unknown files in a virtualized sandbox environment in the cloud or on the WF-500 appliance. WildFire monitors more than 250 malicious behaviors and, if it finds malware, it automatically develops a signature and delivers it to you in as little as five minutes (and now that unknown threat is a known threat).

Why Do I Need a Best Practice Internet Gateway Security Policy?

Unlike legacy port-based security policies that either block everything in the interest of network security, or enable everything in the interest of your business, a best practice security policy allows you to safely enable applications by classifying all traffic, across all ports, all the time, including encrypted traffic. By determining the business use case for each application, you can create security policy rules to allow and protect access to relevant applications. Simply put, a best practice security policy is a policy that leverages the next-generation technologies—App-ID, Content-ID, and User-ID—on the Palo Alto Networks enterprise security platform to:

- Identify applications regardless of port, protocol, evasive tactic or encryption
- Identify and control users regardless of IP address, location, or device
- Protect against known and unknown application-borne threats
- Provide fine-grained visibility and policy control over application access and functionality

A best practice security policy uses a layered approach to ensure that you not only safely enable sanctioned applications, but also block applications with no legitimate use case. To mitigate the risk of breaking applications when moving from a port-based enforcement to an application-based enforcement, the best-practice rulebase provides built-in mechanisms to help you identify gaps in the rulebase and detect alarming activity and potential threats on your network. These temporary best practice rules ensure that applications your users are counting on don't break, while allowing you to monitor application usage and craft appropriate rules. You may find that some of the applications that were being allowed through existing port-based policy rules are not necessarily applications that you want to continue to allow or that you want to limit to a more granular set of users.

Unlike a port-based policy, a best-practice security policy is easy to administer and maintain because each rule meets a specific goal of allowing an application or group of applications to a specific user group based on your business needs. Therefore, you can easily understand what traffic the rule enforces by looking at the match criteria. Additionally, a best-practice security policy rulebase leverages tags and objects to make the rulebase more scannable and easier to keep synchronized with your changing environment.

How Do I Deploy a Best Practice Internet Gateway Security Policy?

Moving from a port-based security policy to an application-based security policy may seem like a daunting task. However, the security risks of sticking with a port-based policy far outweigh the effort required to implement an application-based policy. And, while legacy port-based security policies may have hundreds, if not thousands of rules (many of which nobody in the organization knows the purpose), a best practice policy has a streamlined set of rules that align with your business goals, simplifying administration and reducing the chance of error. Because the rules in an application-based policy align with your business goals and acceptable use policies, you can quickly scan the policy to understand the reason for each and every rule.

As with any technology, there is usually a gradual approach to a complete implementation, consisting of carefully planned deployment phases to make the transition as smooth as possible, with minimal impact to your end users. Generally, the workflow for implementing a best practice internet gateway security policy is:

- **Assess your business and identify what you need to protect**—The first step in deploying a security architecture is to assess your business and identify what your most valuable assets are as well as what the biggest threats to those assets are. For example, if you are a technology company, your intellectual property is your most valuable asset. In this case, one of your biggest threats would be source code theft.
- Segment Your Network Using Interfaces and Zones—Traffic cannot flow between zones unless there is a security policy rule to allow it. One of the easiest defenses against lateral movement of an attacker that has made its way into your network is to define granular zones and only allow access to the specific user groups who need to access an application or resource in each zone. By segmenting your network into granular zones, you can prevent an attacker from establishing a communication channel within your network (either via malware or by exploiting legitimate applications), thereby reducing the likelihood of a successful attack on your network.
- Identify Whitelist Applications—Before you can create an internet gateway best practice security policy, you must have an inventory of the applications you want to allow on your network, and distinguish between those applications you administer and officially sanction and those that you simply want users to be able to use safely. After you identify the applications (including general types of applications) you want to allow, you can map them to specific best practice rules.
- Create User Groups for Access to Whitelist Applications—After you identify the applications you plan to allow, you must identify the user groups that require access to each one. Because compromising an end user's system is one of the cheapest and easiest ways for an attacker to gain access to your network, you can greatly reduce your attack surface by only allowing access to applications to the user groups that have a legitimate business need.
- Decrypt Traffic for Full Visibility and Threat Inspection—You can't inspect traffic for threats if you can't see it. And today SSL/TLS traffic flows account for 40% or more of the total traffic on a typical network. This is precisely why encrypted traffic is a common way for attackers to deliver threats. For example, an attacker may use a web application such as Gmail, which uses SSL encryption, to email an exploit or malware to employees accessing that application on the corporate network. Or, an attacker may compromise a web site that uses SSL encryption to silently download an exploit or malware to site visitors. If you are not decrypting traffic for visibility and threat inspection, you are leaving a very large surface open for attack.
- Create Best Practice Security Profiles—Command and control traffic, CVEs, drive-by downloads of malicious content, APTs are all delivered via legitimate applications. To protect against known and unknown threats, you must attach stringent security profiles to all Security policy allow rules.
- Define the Initial Internet Gateway Security Policy—Using the application and user group inventory you conducted, you can define an initial policy that allows access to all of the applications you want to whitelist by user or user group. The initial policy rulebase you create must also include temporary rules

 © Palo Alto Networks, Inc.

to prevent other applications you might not have known about from breaking and to identify policy gaps and security holes in your existing design.

- Monitor and Fine Tune the Policy Rulebase—After the temporary rules are in place, you can begin monitoring traffic that matches to them so that you can fine tune your policy. Because the temporary rules are designed to uncover unexpected traffic on the network, such as traffic running on non-default ports or traffic from unknown users, you must assess the traffic matching these rules and adjust your application allow rules accordingly.
- Remove the Temporary Rules—After a monitoring period of several months, you should see less and less traffic hitting the temporary rules. When you reach the point where traffic no longer hits the temporary rules, you can remove them to complete your best practice internet gateway security policy.
- Maintain the Rulebase—Due to the dynamic nature of applications, you must continually monitor your application whitelist and adapt your rules to accommodate new applications that you decide to sanction as well to determine how new or modified App-IDs impact your policy. Because the rules in a best practice rulebase align with your business goals and leverage policy objects for simplified administration, adding support for a new sanctioned application or new or modified App-ID oftentimes is as simple as adding or removing an application from an application group or modifying an application filter.

Identify Whitelist Applications

The application whitelist includes not only the applications you provision and administer for business and infrastructure purposes, but also other applications that your users may need to use in order to get their jobs done, and applications you may choose to allow for personal use. Before you can begin creating your best practice internet gateway security policy, you must create an inventory of the applications you want to whitelist.

▲ Map Applications to Business Goals for a Simplified Rulebase
▲ Use Temporary Rules to Tune the Whitelist
▲ Application Whitelist Example

Map Applications to Business Goals for a Simplified Rulebase

As you inventory the applications on your network, consider your business goals and acceptable use policies and identify the applications that correspond to each. This will allow you to create a goal-driven rulebase. For example, one goal might be to allow all users on your network to access data center applications. Another goal might be to allow the sales and support groups access your customer database. You can then create a whitelist rule that correspond to each goal you identify and group all of the applications that align with the goal into a single rule. This approach allows you to create a rulebase with a smaller number of individual rules, each with a clear purpose.

In addition, because the individual rules you create align with your business goals, you can use application objects to group the whitelist to further simplify administration of the best practice rulebase:

- Create application groups for sanctioned applications—Because you will know exactly what applications you require and sanction for official use, create application groups that explicitly include only those applications. Using application groups also simplifies the administration of your policy because it allows you to add and remove sanctioned applications without requiring you to modify individual policy rules. Generally, if the applications that map to the same goal have the same requirements for enabling access

(for example, they all have a destination address that points to your data center address group, they all allow access to any known user, and you want to enable them on their default ports only) you would add them to the same application group.

- Create application filters to allow general types of applications—Besides the applications you officially sanctioned, you will also need to decide what additional applications you will want to allow your users to access. Application filters allow you to safely enable certain categories of applications using application filters (based on category, subcategory, technology, risk factor, or characteristic). Separate the different types of applications based on business and personal use. Create separate filters for each type of application to make it easier to understand each policy rule at a glance.

Use Temporary Rules to Tune the Whitelist

Although the end-goal of a best-practice application-based policy is to use positive enforcement to safely enable your whitelist applications, the initial rulebase requires some additional rules designed to ensure that you have full visibility into the all applications in use on your network so that you can properly tune it. The initial rulebase you create will have the following types of rules:

- Whitelist rules for the applications you officially sanction and deploy.
- Whitelist rules for safely enabling access to general types of applications you want to allow per your acceptable use policy.
- Blacklist rules that block applications that have no legitimate use case. You need these rules so that the temporary rules that “catch” applications that haven’t yet been accounted for in your policy don’t let anything bad onto your network.
- Temporary allow rules to give you visibility into all of the applications running on your network so that you can tune the rulebase.

The temporary rules are a very important part of the initial best practice rulebase. Not only will they give you visibility into applications you weren’t aware were running on your network (and prevent legitimate applications you didn’t know about from breaking), but they will also help you identify things such as unknown users and applications running on non-standard ports. Because attackers commonly use standard applications on non-standard ports as an evasion technique, allowing applications on any port opens the door for malicious content. Therefore, you must identify any legitimate applications running on non-standard ports (for example, internally developed applications) so that you can either modify what ports are used or create a custom applications to enable them.

Application Whitelist Example

Keep in mind that you do not need to capture every application that might be in use on your network in your initial inventory. Instead you should focus here on the applications (and general types of applications) that you want to allow. Temporary rules in the best practice rulebase will catch any additional applications that may be in use on your network so that you are not inundated with complaints of broken applications during your transition to application-based policy. The following is an example application whitelist for an enterprise gateway deployment.

 © Palo Alto Networks, Inc.

Application Type	Best Practice for Securing

Sanctioned Applications	These are the applications that your IT department administers specifically for business use within your organization or to provide infrastructure for your network and applications. For example, in an internet gateway deployment these applications fall into the following categories:
	• **Infrastructure Applications**—These are the applications that you must allow to enable networking and security, such as ping, NTP, SMTP, and DNS.
	• **IT Sanctioned Applications**—These are the applications that you provision and administer for your users. These fall into two categories:
	  • **IT Sanctioned On-Premise Applications**—These are the applications you install and host in your data center for business use. With IT sanctioned on-premise applications, the application infrastructure and the data reside on enterprise-owned equipment. Examples include Microsoft Exchange and active sync, as well as authentication tools such as Kerberos and LDAP.
	  • **IT Sanctioned SaaS Applications**—SaaS applications are those where the software and infrastructure are owned and managed by the application service provider, but where you retain full control of the data, including who can create, access, share, and transfer it (for example, Salesforce, Box, and GitHub).
	• **Administrative Applications**—These are applications that only a specific group of administrative users should have access to in order to administer applications and support users (for example, remote desktop applications).

General Types of Applications	Besides the applications you officially sanction and deploy, you will also want to allow your users to safely use other types of applications: • **General Business Applications**—For example, allow access to software updates, and web services, such as WebEx, Adobe online services, and Evernote. • **Personal Applications**—For example, you may want to allow your users to browse the web or safely use web-based mail, instant messaging, or social networking applications. The recommended approach here is to begin with wide application filters so you can gain an understanding of what applications are in use on your network. You can then decide how much risk you are willing to assume and begin to pare down the application whitelist. For example, suppose you find that Box, Dropbox, and Office 365 file-sharing applications are all on use on your network. Each of these applications has an inherent risk associated with it, from data leakage to risks associated with transfer of malware-infected files. The best approach would be to officially sanction a single file-sharing application and then begin to phase out the others by slowly transitioning from an allow policy to an alert policy, and finally, after giving users ample warning, a block policy for all file sharing applications except the one you choose to sanction. In this case, you might also choose to enable a small group of users to continue using an additional file-sharing application as needed to perform job functions with partners.
Custom Applications Specific to Your Environment	If you have proprietary applications on your network or applications that you run on non-standard ports, it is a best practice to create custom applications for them. This way you can allow the application as a sanctioned application and lock it down to its default port. Otherwise you would either have to open up additional ports (for applications running on non-standard ports), or allow unknown traffic (for proprietary applications), neither of which are recommended in a best practice Security policy.

Create User Groups for Access to Whitelist Applications

Safely enabling applications means not only defining the list of applications you want to allow, but also enabling access only for those users who have a legitimate business need. For example, some applications, such as SaaS applications that enable access to Human Resources services (such as Workday or Service Now) must be available to any known user on your network. However, for more sensitive applications you can reduce your attack surface by ensuring that only users who need these applications can access them. For example, while IT support personnel may legitimately need access to remote desktop applications, the majority of your users do not. Limiting user access to applications prevents potential security holes for an attacker to gain access to and control over systems in your network.

To enable user-based access to applications:

- Enable User-ID in zones from which your users initiate traffic.
- For each application whitelist rule you define, identify the user groups that have a legitimate business need for the applications allowed by the rule. Keep in mind that because the best practice approach is to map the application whitelist rules to your business goals (which includes considering which users have a business need for a particular type of application), you will have a much smaller number of rules to manage than if you were trying to map individual port-based rules to users.
- If you don't have an existing group on your AD server, you can alternatively create custom LDAP groups to match the list of users who need access to a particular application.

Decrypt Traffic for Full Visibility and Threat Inspection

The best practice security policy dictates that you decrypt all traffic except sensitive categories, which include Health, Finance, Government, Military, and Shopping.

Use decryption exceptions only where required, and be precise to ensure that you are limiting the exception to a specific application or user based on need only:

- If decryption breaks an important application, create an exception for the specific IP address, domain, or common name in the certificate associated with the application.
- If a specific user needs to be excluded for regulatory or legal reasons, create an exception for just that user.

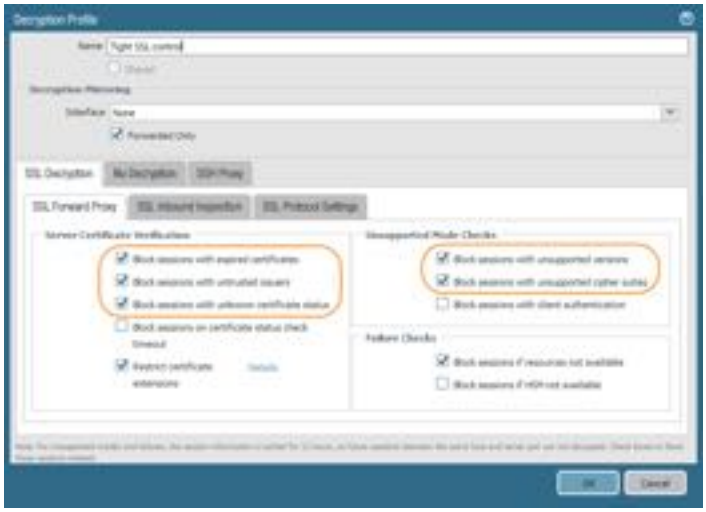
To ensure that certificates presented during SSL decryption area valid, configure the firewall to perform CRL/OCSP checks.

Best practice Decryption policy rules include a strict Decryption Profile. Before you configure SSL Forward Proxy, create a best practice Decryption Profile (**Objects > Decryption Profile**) to attach to your Decryption policy rules:

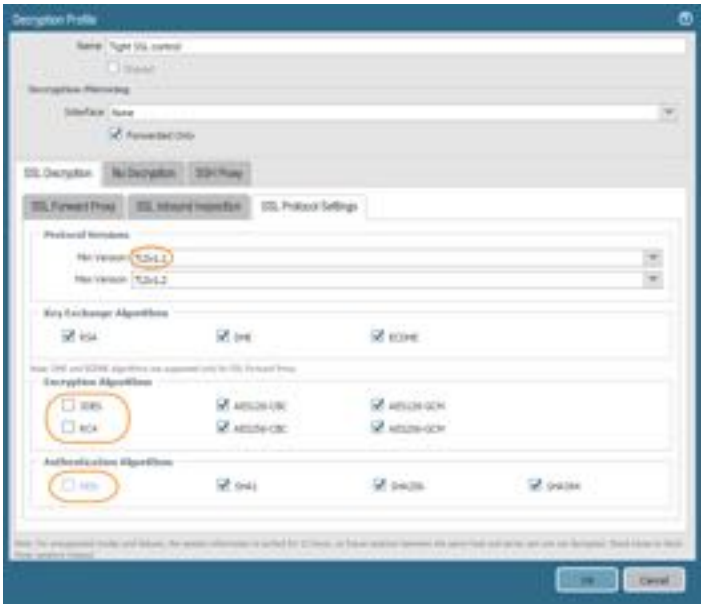
 © Palo Alto Networks, Inc.

Best Practice Decryption Profile

Step 1 Configure the **SSL Decryption > SSL Forward Proxy** settings to block exceptions during SSL negotiation and block sessions that can't be decrypted:



Step 2 Configure the **SSL Decryption > SSL Protocol Settings** to block use of vulnerable SSL/TLS versions (TLS 1.0 and SSLv3) and to avoid weak algorithms (MD5, RC4, and 3DES):



Best Practice Decryption Profile (Continued)

Step 3 For traffic that you are not decrypting, configure the **No Decryption** settings to to block encrypted sessions to sites with expired certificates or untrusted issuers:

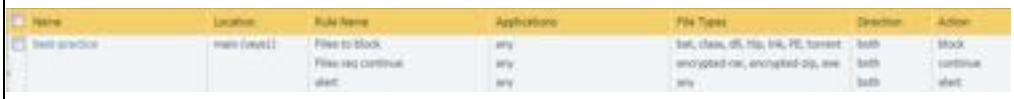
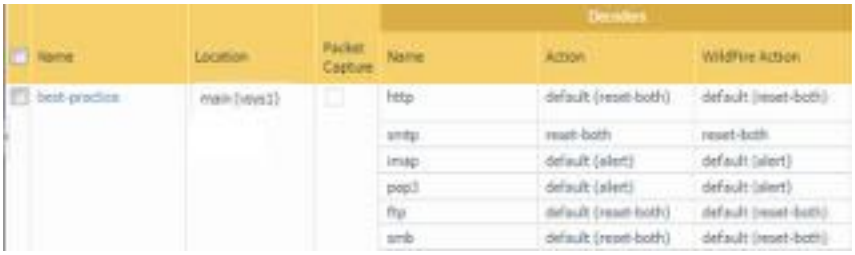


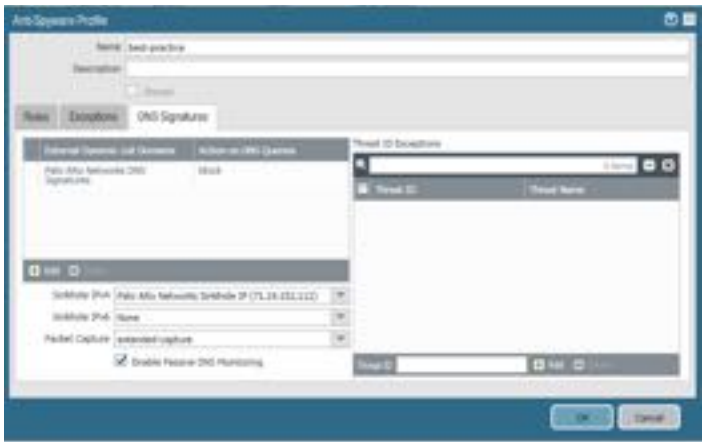
Create Best Practice Security Profiles

Most malware sneaks onto the network in legitimate applications or services. Therefore, to safely enable applications you must scan all traffic allowed into the network for threats. To do this, attach security profiles to all Security policy rules that allow traffic so that you can detect threats—both known and unknown—in your network traffic. The following are the recommended best practice settings for each of the security profiles that you should attach to every Security policy rule.

Consider adding the best practice security profiles to a default security profile group so that it will automatically attach to any new Security policy rules you create.

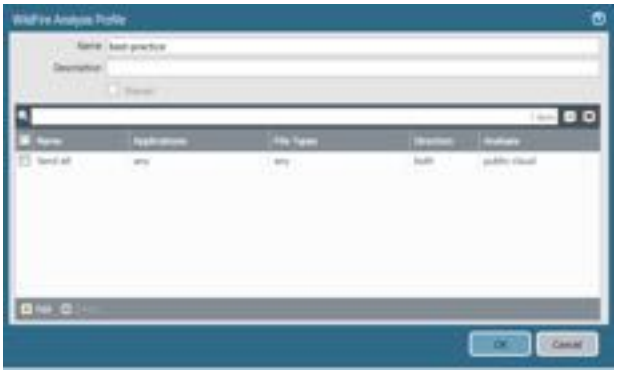
 © Palo Alto Networks, Inc.

Security Profile	Best Practice Settings
File Blocking	Create a File Blocking profile that blocks files that are commonly included in malware attack campaigns or that have no real use case for upload/download. Currently, these include batch files, DLLs, Java class files, help files, Windows shortcuts (.lnk), and BitTorrent files as well as Windows Portable Executable (PE) files, which include .exe, .cpl, .dll, .ocx, .sys, .scr, .drv, .efi, .fon, and .pif files. You can allow download/upload of executables and archive files (.zip and .rar), but force users to click continue before transferring a file to give them pause. Finally, alert on all other file types for visibility into what other file transfers are happening so that you can determine if you need to make policy changes.  **Why do I need this profile?** There are many ways for attackers to deliver malicious files: As attachments or links in corporate email or in webmail, links or IMs in social media, Exploit Kits, through file sharing applications (such as FTP, Google Drive, or Dropbox), or on USB drives. Attaching a File Blocking profile reduces your attack surface by preventing these types of attacks. **What if I can't block all of the recommended file types?** If you cannot block all PE files per the recommendation, make sure you send all unknown files to WildFire for analysis. Additionally, set the Action to continue to prevent drive-by downloads. A drive-by download is when an end user downloads content that installs malicious files, such as Java applets or executables, without knowing they are doing it. Drive-by downloads can occur when users visit web sites, view email messages, or click into pop-up windows meant to deceive them. Educate your users that if they are prompted to continue with a file transfer they didn't knowingly initiate, they may be subject to a malicious download.
Antivirus	Attach an Antivirus profile to all allowed traffic to detect and prevent viruses and malware from being transferred over the HTTP, SMTP, IMAP, POP3, FTP, and SMB protocols. The best practice Antivirus profile uses the default action when it detects traffic that matches either an Antivirus signature or a WildFire signature. The default action differs for each protocol and follows the most up-to-date recommendation from Palo Alto Networks for how to best prevent malware in each type of protocol from propagating. By default, the firewall alerts on viruses found in SMTP traffic. However, if you don't have a dedicated Antivirus gateway solution in place for your SMTP traffic, define a stricter action for this protocol to protect against infected email content. Use the reset-both action to return a 541 response to the sending SMTP server to prevent it from resending the blocked message.  **Why do I need this profile?** By attaching Antivirus profiles to all Security rules you can block known malicious files (malware, ransomware bots, and viruses) as they are coming into the network. Common ways for users to receive malicious files include malicious attachments in email, links to download malicious files, or silent compromise with Exploit Kits that exploit a vulnerability and then automatically deliver malicious payloads to the end user.

Security Profile	Best Practice Settings
Vulnerability Protection	Attach a Vulnerability Protection profile to all allowed traffic to protect against buffer overflows, illegal code execution, and other attempts to exploit client- and server-side vulnerabilities. The best practice profile is a clone of the predefined Strict profile, with packet capture settings enabled to help you track down the source of any potential attacks.  Why do I need this profile? Without strict vulnerability protection, attackers can leverage client- and server-side vulnerabilities to compromise end-users. For example, an attacker could leverage a vulnerability to install malicious code on client systems or use an Exploit Kit (Angler, Nuclear, Fiesta, KaiXin) to automatically deliver malicious payloads to the end user. Vulnerability Protection profiles also prevent an attacker from using vulnerabilities on internal hosts to move laterally within your network.
Anti-Spyware	Attach an Anti-Spyware profile to all allowed traffic to detect command and control traffic (C2) initiated from spyware installed on a server or endpoint and prevents compromised systems from establishing an outbound connection from your network. The best practice Anti-Spyware profile resets the connection when the firewall detects a medium, high, or critical severity threat and blocks or sinkholes any DNS queries for known malicious domains.  To create this profile, clone the predefined strict profile and make sure to enable DNS sinkhole and packet capture to help you track down the endpoint that attempted to resolve the malicious domain. For the best possible protection, enable passive DNS monitoring, which 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. 

 © Palo Alto Networks, Inc.

Security Profile	Best Practice Settings
URL Filtering	As a best practice, use PAN-DB URL filtering to prevent access to web content that is at high-risk for being malicious. Attach a URL Filtering profile to all rules that allow access to web-based applications to protect against URLs that have been observed hosting malware or exploitive content. The best practice URL Filtering profile sets all known dangerous URL categories to block. These include malware, phishing, dynamic DNS, unknown, proxy-avoidance-and-anonymizers, questionable, extremism, copyright-infringement, and parked. Failure to block these dangerous categories puts you at risk for exploit infiltration, malware download, command and control activity, and data exfiltration. In addition to blocking known bad categories, you should also alert on all other categories so that you have visibility into the sites your users are visiting. If you need to phase in a block policy, set categories to continue and create a custom response page to educate users on your acceptable use policies and alert them to the fact that they are visiting a site that may pose a threat. This will pave the way for you to outright block the categories after a monitoring period.  What if I can't block all of the recommended categories? If you find that users need access to sites in the blocked categories, consider creating an allow list for just the specific sites, if you feel the risk is justified. Allowing traffic to a recommended block category poses the following risks: • **malware**—Sites known to host malware or used for command and control (C2) traffic. May also exhibit Exploit Kits. • **phishing**—Known to host credential phishing pages or phishing for personal identification. • **dynamic-dns**—Hosts and domain names for systems with dynamically assigned IP addresses and which are oftentimes used to deliver malware payloads or C2 traffic. Also, dynamic DNS domains do not go through the same vetting process as domains that are registered by a reputable domain registration company, and are therefore less trustworthy. • **unknown**—Sites that have not yet been identified by PAN-DB, perhaps because they were just registered. However, oftentimes these are sites that are generated by domain generation algorithms and are later found to exhibit malicious behavior. • **proxy-avoidance-and-questionable**—URLs and services often used to bypass content filtering products. • **questionable**—Domains with illegal content, such as content that allows illegal download of software or other intellectual property. • **parked**—Domains registered by individuals, oftentimes later found to be used for credential phishing. These domains may be similar to legitimate domains, for example, pal0alto0netw0rks.com, with the intent of phishing for credentials or personal identify information. Or, they may be domains that an individual purchases rights to in hopes that it may be valuable someday, such as panw.net.

Security Profile	Best Practice Settings
WildFire Analysis	While the rest of the best practice security profiles significantly reduce the attack surface on your network by detecting and blocking known threats, the threat landscape is ever changing and the risk of unknown threats lurking in the files we use daily—PDFs, Microsoft Office documents (.doc and .xls files)—is ever growing. And, because these unknown threats are increasingly sophisticated and targeted, they often go undetected until long after a successful attack. To protect your network from unknown threats, you must configure the firewall to forward files to WildFire for analysis. Without this protection, attackers have free reign to infiltrate your network and exploit vulnerabilities in the applications your employees use everyday. Because WildFire protects against unknown threats, it is your greatest defense against advanced persistent threats (APTs). The best practice WildFire Analysis profile sends all files in both directions (upload and download) to WildFire for analysis. Specifically, make sure you are sending all PE files (if you're not blocking them per the file blocking best practice), Adobe Flash and Reader files (PDF, SWF), Microsoft Office files (PowerPoint, Excel, Word, RTF), Java files (Java, .CLASS), and Android files (.APK). 

Define the Initial Internet Gateway Security Policy

The overall goal of a best practice internet gateway security policy is to use positive enforcement of whitelist applications. However, it takes some time to identify exactly what applications are running on your network, which of these applications are critical to your business, and who the users are that need access to each one. The best way to accomplish the end goal of a policy rulebase that includes only application allow rules is to create an initial policy rulebase that liberally allows both the applications you officially provision for your users as well as other general business and, if appropriate, personal applications. This initial policy also includes additional rules that explicitly block bad applications as well as some temporary allow rules that are designed to help you refine your policy and prevent applications your users may need from breaking while you transition to the best practices.

The following topics describe how to create the initial rulebase and describe why each rule is necessary and what the risks are of not following the best practice recommendation:

- Step 1: Create the Application Whitelist Rules
- Step 2: Create the Application Block Rules
- Step 3: Create the Temporary Tuning Rules
- Step 4: Enable Logging for Traffic that Doesn't Match Any Rules

 © Palo Alto Networks, Inc.

Step 1: Create the Application Whitelist Rules

After you Identify Whitelist Applications you are ready to create the first part of the best practice internet gateway security policy rulebase: the application whitelist rules. Every whitelist rule you create must allow traffic based on application (not port) and, with the exception of certain infrastructure applications that require user access before the firewall can identify the user, must only allow access to known users. Whenever possible, Create User Groups for Access to Whitelist Applications so that you can limit user access to the specific users or user groups who have a business need to access the application.

When creating the application whitelist rules, make sure to place more specific rules above more general rules. For example, the rules for all of your sanctioned and infrastructure applications would come before the rules that allow general access to certain types of business and personal applications. This first part of the rulebase includes the allow rules for the applications you identified as part of your application whitelist:

- Sanctioned applications you provision and administer for business and infrastructure purposes
- General business applications that your users may need to use in order to get their jobs done
- General applications you may choose to allow for personal use

Every application whitelist rule also requires that you attach the best practice security profiles to ensure that you are scanning all allowed traffic for known and unknown threats. If you have not yet created these profiles, see Create Best Practice Security Profiles. And, because you can't inspect what you can't see, you must also make sure you have configured the firewall to Decrypt Traffic for Full Visibility and Threat Inspection.

Create the Application Whitelist Rules

Step 1 Allow access to your corporate DNS servers.

Why do I need this rule?	Rule Highlights
❑ Access to DNS is required to provide network infrastructure services, but it is commonly exploited by attackers. ❑ Allowing access only on your internal DNS server reduces your attack surface.	• Because this rule is very specific, place it at the top of the rulebase. • Create an address object to use for the destination address to ensure that users only access the DNS server in your data center. • Because users will need access to these services before they are logged in, you must allow access to any user.



Step 2 Allow access to other required IT infrastructure resources.

Why do I need this rule?	Rule Highlights
❑ Enable the applications that provide your network infrastructure and management functions, such as NTP, OCSP, STUN, and ping. ❑ While DNS traffic allowed in the preceding rule is restricted to the destination address in the data center, these applications may not reside in your data center and therefore require a separate rule.	• Because these applications run on the default port, allow access to any user (users may not yet be a known-user because of when these services are needed), and all have a destination address of any, contain them in a single application group and create a single rule to enable access to all of them. • Users may not have logged in yet at the time they need access to the infrastructure applications, so make sure this rule allows access to any user.

Create the Application Whitelist Rules (Continued)



Step 3 Allow access to IT sanctioned SaaS applications.

Why do I need this rule?

- ❑ With SaaS applications, your proprietary data is in the cloud. This rule ensures that only your known users have access to these applications (and the underlying data).
- ❑ Scan allowed SaaS traffic for threats.

Rule Highlights

- Group all sanctioned SaaS applications in an application group.
- SaaS applications should always run on the application default port.
- Restrict access to your known users. See Create User Groups for Access to Whitelist Applications.



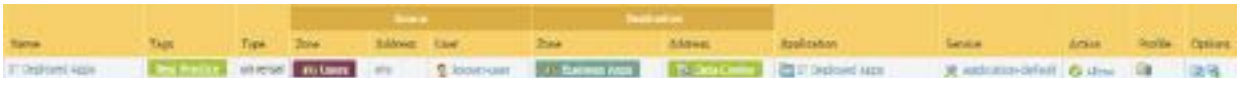
Step 4 Allow access to IT provisioned on-premise applications.

Why do I need this rule?

- ❑ Business-critical data center applications are often leveraged in attacks during the exfiltration stage, using applications such as FTP, or in the lateral movement stage by exploiting application vulnerabilities.
- ❑ Many data center applications use multiple ports; setting the Service to application-default safely enables the applications on their standard ports. You should not allow applications on non-standard ports because it is often associated with evasive behavior.

Rule Highlights

- Group all data center applications in an application group.
- Create an address group for your data center server addresses.
- Restrict access to your known users. See Create User Groups for Access to Whitelist Applications.



Step 5 Allow access to applications your administrative users need.

Why do I need this rule?

- ❑ To reduce your attack surface, Create User Groups for Access to Whitelist Applications.
- ❑ Because administrators often need access to sensitive account data and remote access to other systems (for example RDP), you can greatly reduce your attack surface by only allowing access to the administrators who have a business need.

Rule Highlights

- This rule restricts access to users in the IT_admins group.
- Create custom applications for internal applications or applications that run on non-standard ports so that you can enforce them on their default ports rather than opening additional ports on your network.
- If you have different user groups for different applications, create separate rules for granular control.

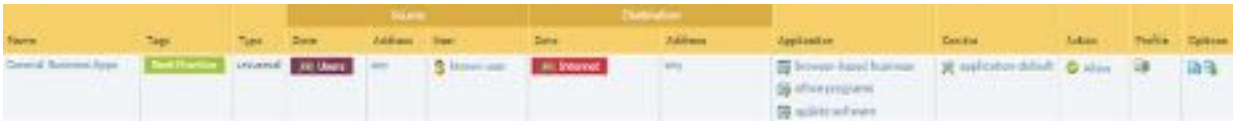


 © Palo Alto Networks, Inc.

Create the Application Whitelist Rules (Continued)

Step 6 Allow access to general business applications.

Why do I need this rule?	Rule Highlights
❑ Beyond the applications you sanction for use and administer for your users, there are a variety of applications that users may commonly use for business purposes, for example to interact with partners, such as WebEx, Adobe online services, or Evernote, but which you may not officially sanction. ❑ Because malware often sneaks in with legitimate web-based applications, this rule allows you to safely allow web browsing while still scanning for threats. See Create Best Practice Security Profiles.	• Restrict access to your known users. See Create User Groups for Access to Whitelist Applications. • For visibility, create separate application filters for each type of application you want to allow. • Attach the best practice security profiles to ensure that all traffic is free of known and unknown threats. See Create Best Practice Security Profiles.



Step 7 (Optional) Allow access to personal applications.

Why do I need this rule?	Rule Highlights
❑ As the lines blur between work and personal devices, you want to ensure that all applications your users access are safely enabled and free of threats. ❑ By using application filters, you can safely enable access to personal applications when you create this initial rulebase. After you assess what applications are in use, you can use the information to decide whether to remove the filter and allow a smaller subset of personal applications appropriate for your acceptable use policies.	• Restrict access to your known users. See Create User Groups for Access to Whitelist Applications. • For visibility, create separate application filters for each type of application you want to allow. • Scan all traffic for threats by attaching your best practice security profile group. See Create Best Practice Security Profiles.



Create the Application Whitelist Rules (Continued)	
Step 8 Allow general web browsing.	
Why do I need this rule? ❑ While the previous rule allowed access to personal applications (many of them browser-based), this rule allows general web browsing. ❑ General web browsing is more risk-prone than other types of application traffic. You must Create Best Practice Security Profiles and attach them to this rule in order to safely enable web browsing. ❑ Because threats often hide in encrypted traffic, you must Decrypt Traffic for Full Visibility and Threat Inspection if you want to safely enable web browsing.	**Rule Highlights** • This rule uses the same best practice security profiles as the rest of the rules, except for the File Blocking profile, which is more stringent because general web browsing traffic is more vulnerable to threats. • This rule allows only known users to prevent devices with malware or embedded devices from reaching the internet. • Use application filters to allow access to general types of applications. • Make sure you also explicitly allow SSL as an application here if you want to allow users to be able to browse to HTTPS sites. that are excluded from decryption.
	

Step 2: Create the Application Block Rules

Although the overall goal of your security policy is to safely enable applications using application whitelist rules (also known as *positive enforcement*), the initial best practice rulebase must also include rules to help you find gaps in your policy and identify possible attacks. Because these rules are designed to catch things you didn't know were running on your network, they allow traffic that could also pose security risks on your network. Therefore, before you can create the temporary rules, you must create rules that explicitly blacklist applications designed to evade or bypass security or that are commonly exploited by attackers, such as public DNS and SMTP, encrypted tunnels, remote access, and non-sanctioned file-sharing applications.

Each of the tuning rules you will define in Step 3: Create the Temporary Tuning Rules are designed to identify a specific gap in your initial policy. Therefore some of these rules will need to go above the application block rules and some will need to go after.

 © Palo Alto Networks, Inc.

Create the Application Block Rules

Step 1 Block applications that do not have a legitimate use case.

Why do I need this rule?	Rule Highlights
❑ Block nefarious applications such as encrypted tunnels and peer-to-peer file sharing, as well as web-based file sharing applications that are not IT sanctioned. ❑ Because the tuning rules that follow are designed to allow traffic with malicious intent or legitimate traffic that is not matching your policy rules as expected, these rules could also allow risky or malicious traffic into your network. This rule prevents that by blocking traffic that has no legitimate use case and that could be used by an attacker or a negligent user.	• Use the **Drop** Action to silently drop the traffic without sending a signal to the client or the server. • Enable logging for traffic matching this rule so that you can investigate misuse of applications and potential threats on your network. • Because this rule is intended to catch malicious traffic, it matches to traffic from any user running on any port.



Step 2 Block public DNS and SMTP applications.

Why do I need this rule?	Rule Highlights
❑ Block public DNS/SMTP applications to avoid DNS tunneling, command and control traffic, and remote administration.	• Use the **Reset both client and server** Action to send a TCP reset message to both the client-side and server-side devices. • Enable logging for traffic matching this rule so that you can investigate a potential threat on your network.



Step 3: Create the Temporary Tuning Rules

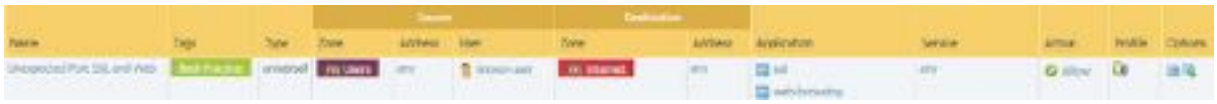
The temporary tuning rules are explicitly designed to help you monitor the initial best practice rulebase for gaps and alert you to alarming behavior. For example, you will create temporary rules to identify traffic that is coming from unknown user or applications running on unexpected ports. By monitoring the traffic matching on the temporary rules you can also gain a full understanding of all of the applications in use on your network (and prevent applications from breaking while you transition to a best practice rulebase). You can use this information to help you fine tune your whitelist, either by adding new whitelist rules to allow applications you weren't aware were needed or to narrow your whitelist rules to remove application filters and instead allow only specific applications in a particular category. When traffic is no longer hitting these rules you can Remove the Temporary Rules.

Some of the temporary tuning rules must go above the rules to block bad applications and some must go after to ensure that targeted traffic hits the appropriate rule, while still ensuring that bad traffic is not allowed onto your network.

Create Temporary Tuning Rules

Step 1 Allow web-browsing and SSL on non-standard ports for known users to determine if there are any legitimate applications running on non-standard ports.

Why do I need this rule?	Rule Highlights
❑ This rule helps you determine if you have any gaps in your policy where users are unable to access legitimate applications because they are running on non-standard ports. ❑ You must monitor all traffic that matches this rule. For any traffic that is legitimate, you should tune the appropriate allow rule to include the application, perhaps creating a custom application where appropriate.	• Unlike the whitelist rules that allow applications on the default port only, this rule allows web-browsing and SSL traffic on any port so that you can find gaps in your whitelist. • Because this rule is intended to find gaps in policy, limit it to known users on your network. See Create User Groups for Access to Whitelist Applications. • Make sure you also explicitly allow SSL as an application here if you want to allow users to be able to browse to HTTPS sites that aren't decrypted (such as financial services and healthcare sites). • You must add this rule above the application block rules or no traffic will hit this rule.



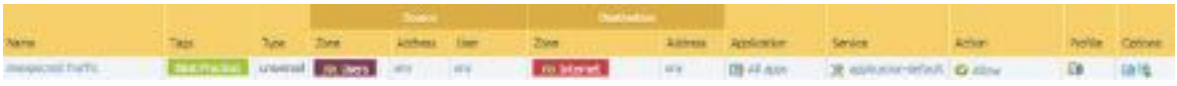
Step 2 Allow web-browsing and SSL traffic on non-standard ports from unknown users to highlight all unknown users regardless of port.

Why do I need this rule?	Rule Highlights
❑ This rule helps you determine whether you have gaps in your User-ID coverage. ❑ This rule also helps you identify compromised or embedded devices that are trying to reach the internet. ❑ It is important to block non-standard port usage, even for web-browsing traffic, because it is usually an evasion technique.	• While the majority of the application whitelist rules apply to known users or specific user groups, this rule explicitly matches traffic from **unknown** users. • Note that this rule must go above the application block rules or traffic will never hit it. • Because it is an allow rule, you must attach the best practice security profiles to scan for threats.



Step 3 Allow all applications on the application-default port to identify unexpected applications.

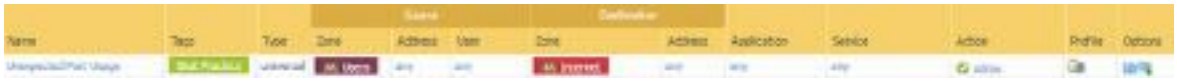
Why do I need this rule?	Rule Highlights
❑ This rule provides visibility into applications that you weren't aware were running on your network so that you can fine-tune your application whitelist. ❑ Monitor all traffic matching this rule to determine whether it represents a potential threat, or whether you need to modify your whitelist rules to allow the traffic.	• Because this rule allows all applications, you must add it after the application block rules to prevent bad applications from running on your network. • If you are running PAN-OS 7.0.x or earlier, to appropriately identify unexpected applications, you must use an application filter that includes all applications, instead of setting the rule to allow **any** application.



 © Palo Alto Networks, Inc.

Create Temporary Tuning Rules	
Step 4	Allow any application on any port to identify applications running where they shouldn't be.

Why do I need this rule?	Rule Highlights
❑ This rule helps you identify legitimate, known applications running on unknown ports. ❑ This rule also helps you identify unknown applications for which you need to create a custom application to add to your application whitelist. ❑ Any traffic matching this rule is actionable and requires that you track down the source of the traffic and ensure that you are not allowing any unknown tcp, udp or non-syn-tcp traffic.	• Because this is a very general rule that allows any application from any user on any port, it must come at the end of your rulebase. • Enable logging for traffic matching this rule so that you can investigate for misuse of applications and potential threats on your network or identify legitimate applications that require a custom application.



Step 4: Enable Logging for Traffic that Doesn't Match Any Rules

Traffic that does not match any of the rules you defined will match the predefined interzone-default rule at the bottom of the rulebase and be denied. For visibility into the traffic that is not matching any of the rules you created, enable logging on the interzone-default rule:

Enable Logging for Traffic That Doesn't Match Any Rules	
Step 1	Select the interzone-default row in the rulebase and click **Override** to enable editing on this rule.
Step 2	Select the **interzone-default** rule name to open the rule for editing.
Step 3	On the **Actions** tab, select **Log at Session End** and click **OK**.
Step 4	Create a custom report to monitor traffic that hits this rule. 1. Select **Monitor > Manage Custom Reports**. 2. **Add** a report and give it a descriptive **Name**. 3. Set the **Database** to **Traffic Summary**. 4. Select the **Scheduled** check box. 5. Add the following to the Selected Columns list: **Rule**, **Application**, **Bytes**, **Sessions**. 6. Set the desired **Time Frame**, **Sort By** and **Group By** fields. 7. Define the query to match traffic hitting the interzone-default rule: `(rule eq 'interzone-default')`
Step 5	**Commit** the changes you made to the rulebase.

Monitor and Fine Tune the Policy Rulebase

A best practice security policy is iterative. It is a tool for safely enabling applications, users, and content by classifying all traffic, across all ports, all the time. As soon as you Define the Initial Internet Gateway Security Policy, you must begin to monitor the traffic that matches the temporary rules designed to identify policy gaps and alarming behavior and tune your policy accordingly. By monitoring traffic hitting these rules, you can make appropriate adjustments to your rules to either make sure all traffic is hitting your whitelist application allow rules or assess whether particular applications should be allowed. As you tune your rulebase, you should see less and less traffic hitting these rules. When you no longer see traffic hitting these rules, it means that your positive enforcement whitelist rules are complete and you can Remove the Temporary Rules.

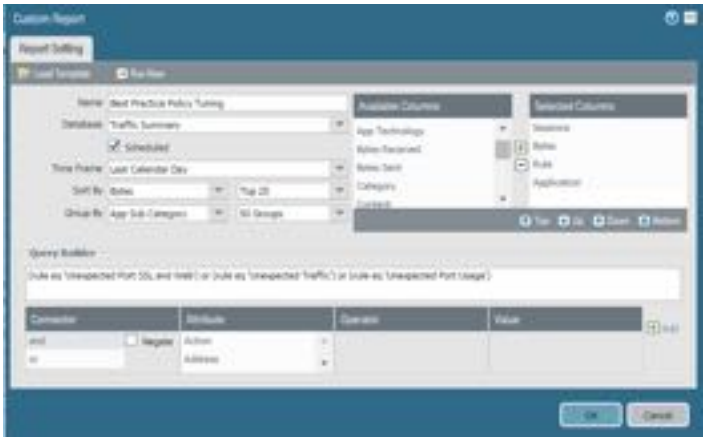


Because new App-IDs are added in weekly content releases, you should review the impact the changes in App-IDs have on your policy.

Identify Policy Gaps

Step 1 Create custom reports that let you monitor traffic that hits the rules designed to identify policy gaps.

1. Select **Monitor > Manage Custom Reports**.
2. **Add** a report and give it a descriptive **Name** that indicates the particular policy gap you are investigating, such as Best Practice Policy Tuning.
3. Set the **Database** to **Traffic Summary**.
4. Select the **Scheduled** check box.
5. Add the following to the Selected Columns list: **Rule**, **Application**, **Bytes**, **Sessions**.
6. Set the desired **Time Frame**, **Sort By** and **Group By** fields.
7. Define the query to match traffic hitting the rules designed to find policy gaps and alarming behavior. You can create a single report that details traffic hitting any of the rules (using the **or** operator), or create individual reports to monitor each rule. Using the rule names defined in the example policy, you would enter the corresponding queries:
   - `(rule eq 'Unexpected Port SSL and Web')`
   - `(rule eq 'Unknown User SSL and Web')`
   - `(rule eq 'Unexpected Traffic')`
   - `(rule eq 'Unexpected Port Usage')`



 © Palo Alto Networks, Inc.

Identify Policy Gaps (Continued)	
Step 2	Review the report regularly to make sure you understand why traffic is hitting each of the best practice policy tuning rules and either update your policy to include legitimate applications and users, or use the information in the report to assess the risk of that application usage and implement policy reforms. 

Remove the Temporary Rules

After several months of monitoring your initial internet gateway best practice security policy, you should see less and traffic hitting the temporary rules as you make adjustments to the rulebase. When you no longer see any traffic hitting these rules, you have achieved your goal of transitioning to a fully application-based Security policy rulebase. At this point, you can finalize your policy rulebase by removing the temporary rules, which includes the rules you created to block bad applications and the rules you created for tuning the rulebase.

Remove the Temporary Rules	
Step 1	Select **Policies > Security**.
Step 2	Select the rule and click **Delete**. Alternatively, **Disable** the rules for a period of time before deleting them. This would allow you to **Enable** them again if traffic logs show traffic matching the interzone-default rule.
Step 3	**Commit** the changes.

Maintain the Rulebase

Because applications are always evolving, your application whitelist will need to evolve also. Each time you make a change in what applications you sanction, you must make a corresponding policy change. As you do this, instead of just adding a new rule like you would do with a port-based policy, instead identify and modify the rule that aligns with the business use case for the application. Because the best practice rules leverage policy objects for simplified administration, adding support for a new application or removing an application from your whitelist typically means modifying the corresponding application group or application filter accordingly.

Additionally, installing new App-IDs included in a content release version can sometimes cause a change in policy enforcement for applications with new or modified App-IDs. Therefore, 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 policy rules and install new content, and allows for a seamless shift in policy enforcement. Alternatively, you can 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.

Maintain the Best Practice Rulebase	
Step 1	Before installing a new content release version, review the new App-IDs to determine if there is policy impact.
Step 2	Disable new App-IDs introduced in a content release, in order 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.
Step 3	Tune security policy rules to account for App-ID changes included in a content release or to add new sanctioned applications to or remove applications from your application whitelist rules.

 © 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, 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 > Security**. 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**.



- 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 Use Global Find to Search the Firewall or Panorama Management Server 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: • Select **Move > Move to other vsys** (for policy rules). • Click **Move** (for objects). • Click **Clone** (for policy rules or objects).
Step 4	In the **Destination** drop-down, select the new virtual system or **Shared**.
Step 5	(Policy rules only) Select the **Rule order**: • **Move top** (default)—The rule will come before all other rules. • **Move bottom**—The rule will come after all other rules. • **Before rule**—In the adjacent drop-down, select the rule that comes after the Selected Rules. • **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 them for easy scanning. You can create tags for the following objects: address objects, address groups, zones, service groups, and policy rules.

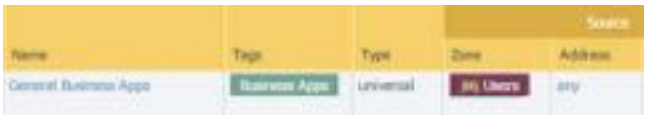
The firewall and Panorama support both static tags and dynamic tags. Dynamic tags are registered from a variety of sources and are not displayed with the static tags because dynamic tags are not part of the firewall/Panorama configuration. See Register IP Addresses and Tags Dynamically for information on registering tags dynamically. The tags discussed in this section are statically added and are part of the configuration.

You can apply one or more tags to objects and to policy rules, up to a maximum of 64 tags per object. Panorama supports a maximum of 10,000 tags, which you can apportion across Panorama (shared and device groups) and the managed firewalls (including firewalls with multiple virtual systems).

- Create and Apply Tags
- Modify Tags
- Use the Tag Browser

Create and Apply Tags

Create and Apply Tags	
Step 1 Create tags. 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**. 2. On Panorama or a multiple virtual system firewall, select the **Device Group** or the **Virtual System** to to make the tag available. 3. Click **Add** and enter a **Name** to identify the tag, or select a zone name from the drop-down to create a tag for a zone. The maximum length is 127 characters. 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. 5. (Optional) Assign one of the 17 predefined colors to the tag. By default, **Color** is **None**.  6. Click **OK** and **Commit** to save the changes.

Create and Apply Tags (Continued)	
Step 2 Apply tags to policy.	1. Select **Policies** and any rulebase under it. 2. Click **Add** to create a policy rule and use the tagged objects you created in Step 1. 3. Verify that the tags are in use. 
Step 3 Apply tags to an address object, address group, service, or service group.	1. Create the object. For example to create a service group, select **Objects > Service Groups > Add**. 2. Select a tag from the **Tags** drop-down or enter a name in the field to create a new tag. To edit a tag or add color to the tag, see Modify Tags.

Modify Tags

Modify Tags

• Select **Objects > Tags** to perform any of the following operations with tags:
  • Click the link in the **Name** column to edit the properties of a tag.
  • Select a tag in the table, and click **Delete** to remove the tag from the firewall.
  • 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.

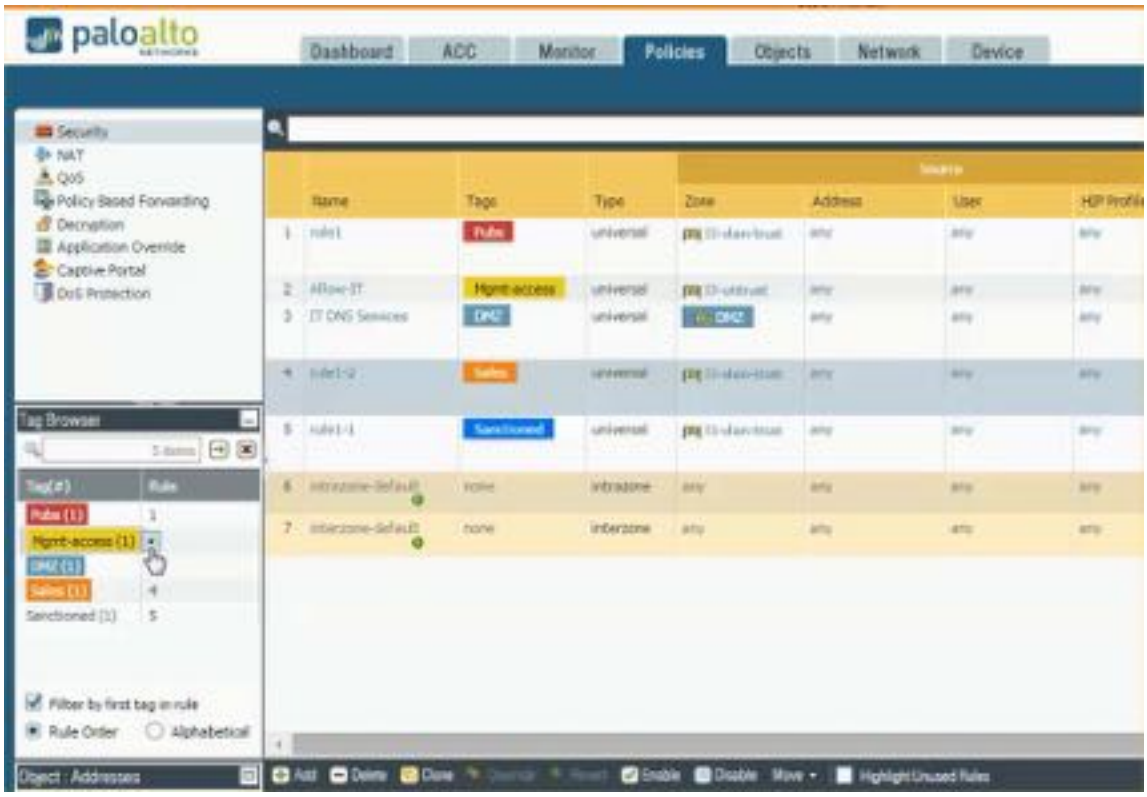
Use the Tag Browser

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.

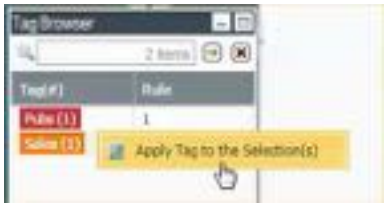
For firewalls that are managed by Panorama, the tags applied to pre-rules and post-rules that have been pushed from Panorama, display in a green background and are demarcated with green lines so that you can identify these tags from the local tags on the firewall.

 © Palo Alto Networks, Inc.


paloalto
Dashboard
ACC
Monitor
Policies
Objects
Network
Device
Security
NAT
QoS
Policy Based Forwarding
Decryption
Application Override
Captive Portal
DoS Protection
Name
Tags
Type
Zone
Address
User
Tag Browser
Tag(#)
Rule
Filter by first tag in rule
Rule Order
Alphabetical
Object : Addresses
Add
Delete
Clone
Enable
Disable
Move
Highlight Unused Rules

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**. 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. 3. **Rule**—Lists the rule number or range of numbers associated with the tags. 4. Sort the tags. • **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. • **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. • **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. 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. 5. **Clear**—Clears the filter on the currently selected tags in the search bar.  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. 7. Expand or collapse the tag browser. 

 © Palo Alto Networks, Inc.

Use the Tag Browser (Continued)	
• Tag a rule.	1. Select a rule on the right pane. 2. Do one of the following: • Select a tag in the tag browser and select **Apply the Tag to the Selection(s)** from the drop-down.  • 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. 3. **Commit** the changes.
• View rules that match the selected tags. 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. • 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. 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.

Use the Tag Browser (Continued)	
• 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)**. 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.
• 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. 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.

 © Palo Alto Networks, Inc.

Use an External Dynamic List in Policy

An *external dynamic list* (formerly called a dynamic block list) is a text file that you host on an external web server so that the firewall can import objects—IP addresses, URLs, domains—to enforce policy on the entries in the list. As you update the list, the firewall dynamically imports the list at the configured interval and enforces policy without the need to make a configuration change or a commit on the firewall.

- External Dynamic List
- Formatting Guidelines for an External Dynamic List
- Enforce Policy on Entries in an External Dynamic List
- View the List of Entries in an External Dynamic List
- Retrieve an External Dynamic List from the Web Server

External Dynamic List

An external dynamic list is a text file that is hosted on an external web server so that the firewall can import objects—IP addresses, URLs, domains—included in the list and enforce policy. To enforce policy on the entries included in the external dynamic list, you must reference the list in a supported policy rule or profile. As you modify the list, the firewall dynamically imports the list at the configured interval and enforces policy without the need to make a configuration change or a commit on the firewall. If the web server is unreachable, the firewall will use the last successfully retrieved list for enforcing policy until the connection is restored with the web server. To retrieve the external dynamic list, the firewall uses the interface attached to the service route that it uses to access the Palo Alto Updates service.

The firewall supports three types of external dynamic lists:

- **IP Address**—The firewall typically enforces 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 or destination IP addresses that emerge ad hoc, you can use an external dynamic list of type IP address as a source or destination address object in policy rules, and configure the firewall to deny or allow access to the IP addresses (IPv4 and IPv6 address, IP range and IP subnets) included in the list. The firewall treats an external dynamic list of type IP address as an address object; all the IP addresses included in a list are handled as one address object.
- **URL**—An external dynamic list of type URL gives you the agility to protect your network from new sources of threat or malware. The firewall handles an external dynamic list with URLs like a custom URL category and you can use this list in two ways:
  - As a match criteria in Security policy rules, Decryption policy rules, and QoS policy rules to allow, deny, decrypt, not decrypt, or allocate bandwidth for the URLs in the custom category.
  - In a URL Filtering profile where you can define more granular actions, such as continue, alert, or override, before you attach the profile to a Security policy rule.
- **Domain**—An external dynamic list of type domain allows you to import custom domain names into the firewall to enforce policy using an Anti-Spyware profile. This capability is very useful if you subscribe to third-party threat intelligence and want to protect your network from new sources of threat or malware as soon as you learn of a malicious domain. For each domain you include in the external dynamic list, the firewall creates a custom DNS-based spyware signature so that you can enable DNS sinkholing. The

DNS-based spyware signature is of type spyware with medium severity and each signature is named `Custom Malicious DNS Query <domain name>`. For details, see Configure DNS Sinkholing for a List of Custom Domains.

On each firewall platform, you can configure a maximum of 30 unique sources for external dynamic lists; these limits are not applicable to Panorama. When using Panorama to manage a firewall that is enabled for multiple virtual systems, if you exceed the limit for the firewall, a commit error displays on Panorama. A source is a URL that includes the IP address or hostname, the path, and the filename for the external dynamic list. The firewall matches the URL (complete string) to determine whether a source is unique.

While the firewall does not impose a limit on the number of lists of a specific type, the following limits are enforced:

- IP address—The PA-5000 Series and the PA-7000 Series firewalls support a maximum of 150,000 total IP addresses; all other platforms support a maximum of 50,000 total IP addresses. No limits are enforced for the number of IP addresses per list. When the maximum supported IP address limit is reached on the firewall, the firewall generates a syslog message.
- URL and domain—A maximum of 50,000 URLs and 50,000 domains are supported on each platform, with no limits enforced on the number of entries per list.



When parsing the list, the firewall skips entries that do not match the list type, and ignores entries that exceed the maximum number supported for the platform.

Formatting Guidelines for an External Dynamic List

An external dynamic list of one type —IP address, URL or Domain—must include entries of that type only.

- ▲ IP Address List
- ▲ Domain List
- ▲ URL List

IP Address List

The external dynamic 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].`

Enter each IP address/range/subnet in a new line; URLs or domains are not supported in this list. A subnet or an IP address range, such as 92.168.20.0/24 or 192.168.20.40-192.168.20.50, count as one IP address entry and not as multiple IP addresses. 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 IP address list:

```
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
```

 © Palo Alto Networks, Inc.

For an IP address that is blocked, you can display a notification page only if the protocol is HTTP.

Domain List

Enter each domain name in a new line; URLs or IP addresses are not supported in this list. Do not prefix the domain name with the protocol, http:// or https://. Wildcards are not supported.

An example list of domains:

```
www.example.com
baddomain.com
qqq.abcedfg.au
```

URL List

See Block and Allow Lists.

Enforce Policy on Entries in an External Dynamic List

Enforce Policy on Entries in an External Dynamic List		
Step 1	Create the external dynamic list and host it on a web server so that the firewall can retrieve the list for policy evaluation.	Create a text file and enter the URLs, domains, or IP addresses in the file. To prevent commit errors and invalid entries, do not prefix http:// or https:// to any of the entries. See Formatting Guidelines for an External Dynamic List.

Enforce Policy on Entries in an External Dynamic List (Continued)	
Step 2 Configure the firewall to access the external dynamic list.	1. Select **Objects > External Dynamic Lists**. 2. Click **Add** and enter a descriptive **Name** for the list. 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. 4. (Panorama only) Select **Disable override** to ensure that a firewall administrator cannot override settings locally on a firewall that inherits this configuration through a Device Group commit from Panorama. 5. In the Type drop-down, select the list type, for example, **URL List**. Ensure that the list only includes entries for the list type. See Verify whether entries in the external dynamic list were ignored or skipped. 6. Enter the **Source** for the list you just created on the web server. The source must include the full path to access the list. For example, https://1.2.3.4/EDL_IP_2015. 7. Click **Test Source URL** to verify that the firewall (not available on Panorama) can connect to the web server.  If the web server is unreachable after the connection is established, the firewall uses the last successfully retrieved list for enforcing policy until the connection is restored with the web server. 8. (Optional) Specify the **Repeat** frequency at which the firewall retrieves the list. By default, the firewall retrieves the list once every hour and commits the changes.  The interval is relative to the last commit. So, for the five-minute interval, the commit occurs in 5 minutes if the last commit was an hour ago. To retrieve the list immediately, see Retrieve an External Dynamic List from the Web Server. 9. Click **OK**. 10. Use the external dynamic list in a security profile or directly in a policy rule, as supported. See the following: • Use an External Dynamic List in a URL Filtering Profile. • Configure DNS Sinkholing for a List of Custom Domains • Use an external dynamic list of Type URL as Match Criteria in a Security Policy Rule. • Use an external dynamic list of Type IP as a Source or Destination Address Object in a Security Policy Rule.

 © Palo Alto Networks, Inc.

Enforce Policy on Entries in an External Dynamic List (Continued)	

Step 3 Use an external dynamic list of Type URL as Match Criteria in a Security Policy Rule.
You can also Use an External Dynamic List in a URL Filtering Profile.
	1. Select **Policies > Security**.
	2. Click **Add** and enter a descriptive **Name** for the rule.
	3. In the **Source** tab, select the **Source Zone**.
	4. In the **Destination** tab, select the **Destination Zone**.
	5. In the **Service/URL Category** tab, click **Add** to select the appropriate external dynamic list from the URL Category list.
	6. In the **Actions** tab, set the **Action Setting** to **Allow** or **Deny**.
	7. Click **OK** and **Commit**.
	8. Verify whether entries in the external dynamic list were ignored or skipped.
	Use the following CLI command on a firewall to review the details for a list.
	`request system external-list show type <domain | ip | url> name_of_ list`
	`For example:`
	`request system external-list show type url EBL_ISAC_Alert_List`
	9. Test that the policy action is enforced.
	a. Attempt to access a URL that is included in the external dynamic list.
	b. Verify that the action you defined is enforced in the browser.
	c. To monitor the activity on the firewall:
	d. Select **ACC** and add a URL Domain as a global filter to view the Network Activity and Blocked Activity for the URL you accessed.
	e. Select **Monitor > Logs > URL Filtering** to access the detailed log view.

Enforce Policy on Entries in an External Dynamic List (Continued)	
Step 4 Use an external dynamic list of Type IP as a Source or Destination Address Object in a Security Policy Rule. This capability is useful if you deploy new servers and want to allow access to the newly deployed servers without requiring a firewall commit.	1. Select **Policies > Security**. 2. Click **Add** and give the rule a descriptive name in the **General** tab. 3. In the **Source** tab, select the **Source Zone** and optionally select the external dynamic list as the Source Address. 4. In the **Destination** tab, select the **Destination Zone** and optionally select the external dynamic list as the Destination Address. 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** or **Deny**. Create separate external dynamic lists if you want to specify allow and deny actions for specific IP addresses. 7. Leave all the other options at the default values. 8. Click **OK** to save the changes. 9. **Commit** the changes. 10. Test that the policy action is enforced. a. Access a IP address that is included in the external dynamic list and verify that action you defined is enforced. b. Select **Monitor > Logs > Traffic** and view the log entry for the session. c. To verify the policy rule that matches a flow, use the following CLI command: `test security-policy-match source <IP_address> destination <IP_address> destination port <port_number> protocol <protocol_number>`

View the List of Entries in an External Dynamic List

View the List of Entries in an External Dynamic List	
Step 1	Log in to the CLI on the firewall.

Step 2	Enter the following command to view the list of entries that the firewall has retrieved from the web server:
	`request system external-list show name <name>`
	For example, for a list named case DBL_2014 of type IP address, 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 an External Dynamic List from the Web Server

You can configure the firewall to retrieve the External Dynamic 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, use the following process:

Retrieve an External Dynamic List	
Step 1	To retrieve the list on demand, select **Objects > External Dynamic Lists**.
Step 2	Select the list that you want to refresh, and click **Import Now**. The job to import the list will be added to queue.
Step 3	To view the status of the job in the Task Manager, see Manage and Monitor Administrative Tasks.

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 Environments for the list.

When monitoring ESXi hosts that are part of the VM-Series NSX edition solution, use Dynamic Address Groups instead of using VM Information Sources to learn about changes in the virtual environment. For the VM-Series NSX edition solution, the NSX Manager provides Panorama with information on the NSX security group to which an IP address belongs. The information from the NSX Manager provides the full context for defining the match criteria in a Dynamic Address Group because it uses the service profile ID as a distinguishing attribute and allows you to properly enforce policy when you have overlapping IP addresses across different NSX security groups. Up to a maximum of 32 tags (from vCenter server and NSX Manager) that can be registered to an IP address.

Set up the VM Monitoring Agent	
Step 1 Enable the VM Monitoring Agent. You can configure up to 10 VM information sources 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: • In an active/passive setup, only the active firewall monitors the VM sources. • In 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 .

 © Palo Alto Networks, Inc.

Set up the VM Monitoring Agent (Continued)	
Step 2 Verify the connection status.	Verify that the connection **Status** displays as connected.  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).

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 firewall.

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-300, 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-200, VM-100	2,500

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.
  - 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.

 © Palo Alto Networks, Inc.

- Validate that the members of the dynamic address groups are updated as new FTP or web servers are deployed. This ensures 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. Watch the tutorial to see a big picture view of the feature.	1. Log in to the web interface of the firewall. 2. Select **Object > Address Groups**. 3. Click **Add** and enter a **Name** and a **Description** for the address group. 4. Select **Type** as **Dynamic**. 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**.  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. Watch 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 just created. 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 to create another policy rule. 8. Click **Commit**.

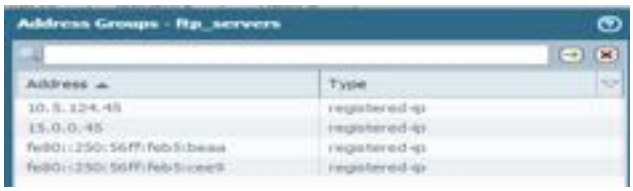
This example shows how to create two policies: one for all access to FTP servers and the other for access to web servers.



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` `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`
	`registered IP Tags`
	`---------- ----------`
	`fe80::20c:29ff:fe69:2f76 "state.poweredOn"`
	`10.1.22.100 "state.poweredOn"`
	`2001:1890:12f2:11:20c:29ff:fe69:2f76 "state.poweredOn"`
	`fe80::20c:29ff:fe69:2f80 "state.poweredOn"`
	`192.168.1.102 "state.poweredOn"`
	`10.1.22.105 "state.poweredOn"`
	`2001:1890:12f2:11:2cf8:77a9:5435:c0d "state.poweredOn"`
	`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`
	`vlanId.4095`
	`vswitch.vSwitch1`
	`host-ip.10.1.5.22`
	`portgroup.TOBEUSED`
	`hostname.panserver22`
	`portgroup.VM Network 2`
	`datacenter.ha-datacenter`
	`vlanId.0`
	`state.poweredOn`
	`vswitch.vSwitch0`
	`vmname.Ubuntu22-100`
	`vmname.win2k8-22-105`
	`resource-pool.Resources`
	`vswitch.vSwitch2`
	`guestos.Ubuntu Linux 32-bit`
	`guestos.Microsoft Windows Server 2008 32-bit`
	`annotation.`
	`version.vmx-08`
	`portgroup.VM Network`
	`vm-info-source.vmware1`
	`uuid.564d362c-11cd-b27f-271f-c361604dfad7`
	`uuid.564dd337-677a-eb8d-47db-293bd6692f76`
	`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:
	`show log iptag datasource_type equal unknown`
	• To view tags registered from the XML API:
	`show log iptag datasource_type equal xml-api`
	• To view tags registered from VM Information sources:
	`show log iptag datasource_type equal vm-monitor`
	• To view tags registered from the Windows User-ID agent:
	`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 Enable User-ID.



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. 2. Select **Use X-Forwarded-For Header in User-ID**.
Step 2	Remove XFF values from outgoing web requests.	1. Select **Strip X-Forwarded-For Header**. 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**). 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**. 2. Select an existing profile or **Add** a new profile and enter a descriptive **Name**.  You can't enable XFF logging in the default URL Filtering profile. 3. In the **Categories** tab, Define how to control access to web content. 4. Select the **Settings** tab and select **X-Forwarded-For**. 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. 2. Select the **Actions** tab, set the **Profile Type** to **Profiles**, and select the **URL Filtering** profile you just created. 3. Click **OK** and **Commit**.
Step 3 Verify the firewall is logging XFF values.	1. Select **Monitor > Logs > URL Filtering**. 2. Display the XFF values in one of the following ways: • 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. • To display the XFF values for all logs—Open the drop-down in any column header, select **Columns**, and select **X-Forwarded-For**. 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

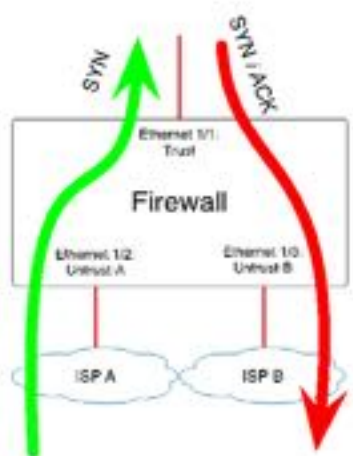
PBF

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
- Path Monitoring for PBF
- Service Versus Applications in PBF

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 for PBF

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. 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 ☐ Disable this rule if nexthop/monitor ip is unreachable	If rule is disabled when the monitored IP address is unreachable ☑ 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.

 © Palo Alto Networks, Inc.

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.



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. 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. You can specify the source and destination addresses using an IP address, an address object, or a FQDN. For the next hop, however, you must specify an IP address.	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, select the following: a. Select the **Type**—**Zone** or **Interface**— to which the forwarding policy will be applied, and the relevant zone or interface. If you want to enforce symmetric return, you must select a source interface. PBF is only supported on Layer 3 interfaces; loopback interfaces do not support PBF. 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. 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. 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. c. (Optional) **Add** and select the **Source User** or groups of users to whom the policy applies. 4. In the **Destination/Application/Service** tab, select the following: 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. b. Select the Application(s) or Service(s) that you want to control using PBF. 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.

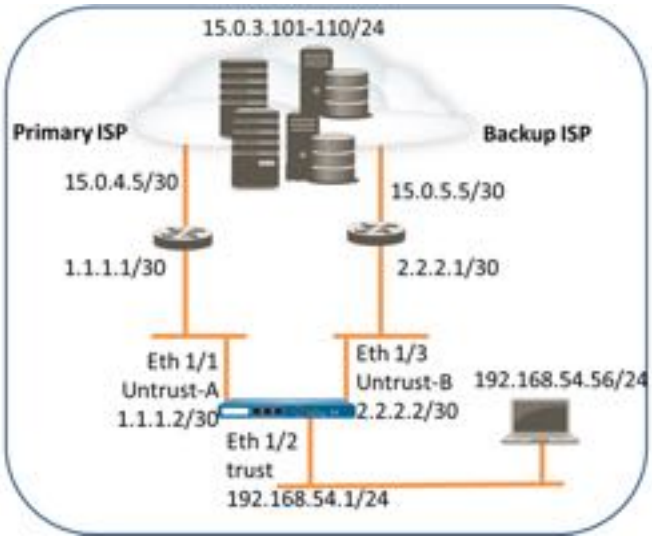
Create a PBF Rule (Continued)		
Step 2	Specify how to forward traffic that matches the rule. If you are configuring PBF in a multi-VSYS environment, you must create separate PBF rules for each virtual system (and create the appropriate Security policy rules to enable the traffic).	1. In the **Forwarding** tab, select the following: a. Set the **Action.** The options are as follows: – **Forward**—Directs the packet to a specific **Egress Interface**. Enter the **Next Hop** IP address for the packet (you cannot use a domain name for the next hop). – **Forward To VSYS**—(On a firewall enabled for multiple virtual systems) Select the virtual system to which to forward the packet. – **Discard**—Drop the packet. – **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. To trigger the specified action at a daily, weekly or non-recurring frequency, create and attach a **Schedule**. (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. b. (Optional, required for asymmetric routing environments) Select **Enforce Symmetric Return** and enter one or more IP addresses in the **Next Hop Address List** (you cannot use an FQDN as the next hop). You can add up to 8 next-hop IP addresses; tunnel and PPoE interfaces are not available as a next-hop IP address. 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 3	Save the policies to the running configuration on the firewall.	Click **Commit**. 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. 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. The interface configuration on the firewall used in this example is as follows: • Ethernet 1/1 connected to the primary ISP: – Zone: ISP-East – IP Address: 1.1.1.2/30 – Virtual Router: Default • Ethernet 1/3 connected to the backup ISP: – Zone: ISP-West – IP Address: 2.2.2.2/30 – Virtual Router: Default • Ethernet 1/2 is the ingress interface, used by the network clients to connect to the internet: – Zone: Trust – IP Address: 192.168. 54.1/24 – Virtual Router: Default 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. 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. 3. Select the **IP Address** radio button and set the **Next Hop** IP address for your router that connects to the backup internet gateway (you cannot use a domain name for the next hop). In this example, 2.2.2.1. 4. Specify a cost metric for the route. In this example, we use 10.  5. Click **OK** twice to save the virtual router configuration.

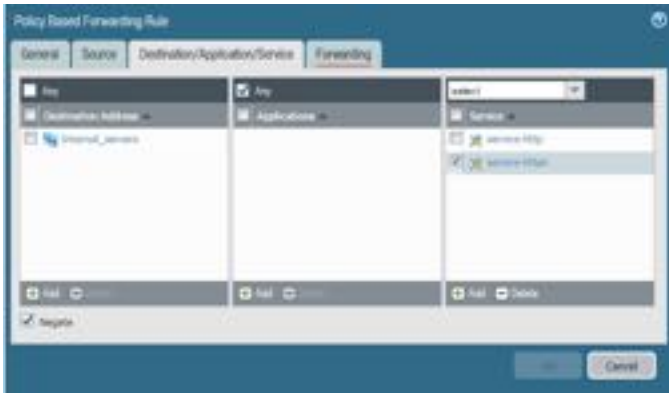
PBF for Outbound Access with Dual ISPs (Continued)

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 (you cannot use a FQDN for the next hop).



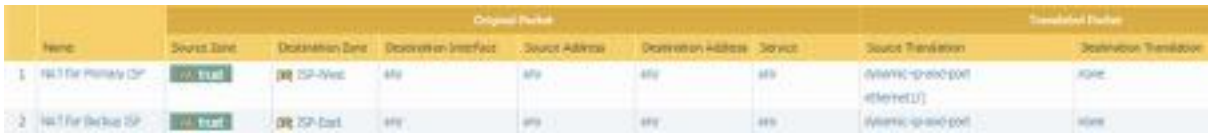
© Palo Alto Networks, Inc.

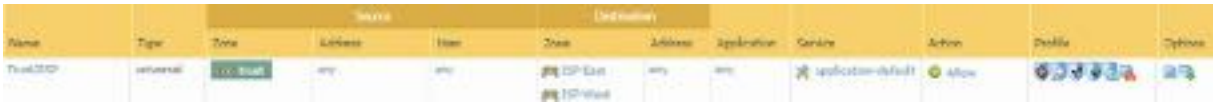
PBF for Outbound Access with Dual ISPs (Continued)

	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.



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**: ethernet1/2 • **IP Address**: 2.2.2.2/30



PBF for Outbound Access with Dual ISPs (Continued)	
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. 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 ISP-East and ISP-West. 5. In the **Service/ URL Category** tab, leave the default **application-default**. 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. Under **Options**, verify that logging is enabled at the end of a session. Only traffic that matches a security rule is logged.
	
Step 6 Save the policies to the running configuration on the firewall.	Click **Commit**.

 © Palo Alto Networks, Inc.

PBF for Outbound Access with Dual ISPs (Continued)

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 following CLI command:

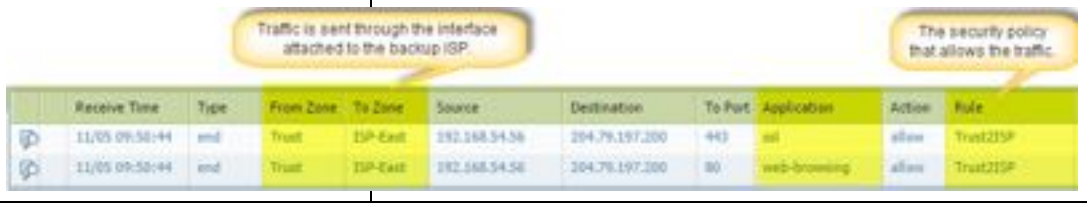
```
admin@PA-NGFW> show pbf rule all
Rule       ID    Rule State Action   Egress IF/VSYS
NextHop
========== === ========== ====== ==============
Use ISP-Pr 1 Active     Forward ethernet1/1 1.1.1.1
```

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 following CLI command:

```
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.



PBF for Outbound Access with Dual ISPs (Continued)

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)
```

 © 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
- Service Routes for Virtual Systems
- Configure Virtual Systems
- Configure Inter-Virtual System Communication within the Firewall
- Configure a Shared Gateway
- Customize Service Routes for a Virtual System
- Virtual System Functionality with Other Features

Virtual Systems Overview

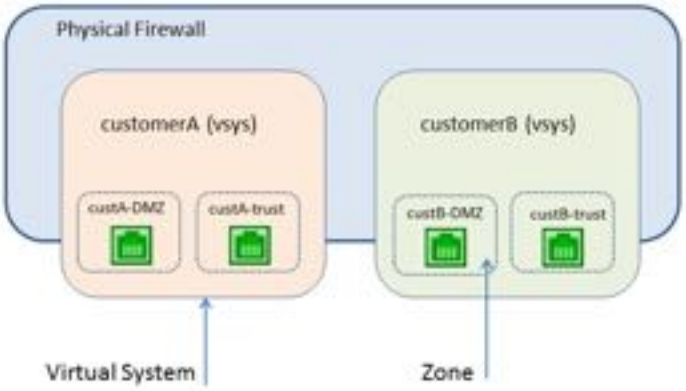
Virtual systems are separate, logical firewall instances within a single physical Palo Alto Networks firewall. Rather than using multiple firewalls, managed service providers and enterprises can use a single pair of firewalls (for high availability) and enable virtual systems on them. Each virtual system (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
- Certificate management

 © Palo Alto Networks, Inc.

- 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 firewall.
- **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.

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

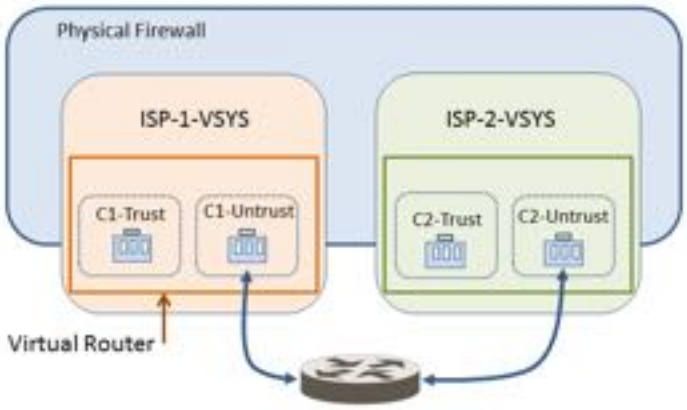
There are two typical scenarios where communication between virtual systems (inter-vsys traffic) is desirable. In a multi-tenancy environment, communication between virtual systems can occur by having traffic leave the firewall, go through the Internet, and re-enter the firewall. In a single organization 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.



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.

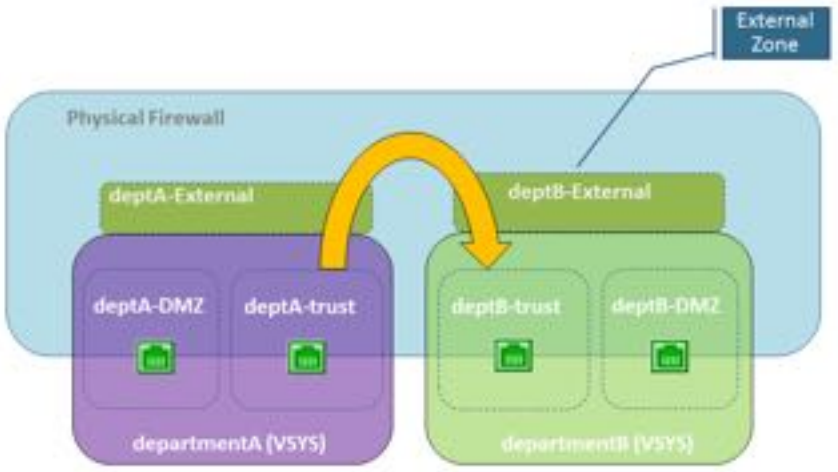
Because external zones do not have interfaces or IP addresses associated with them, some zone protection profiles are not supported on external zones.

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.

 © Palo Alto Networks, Inc.

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.



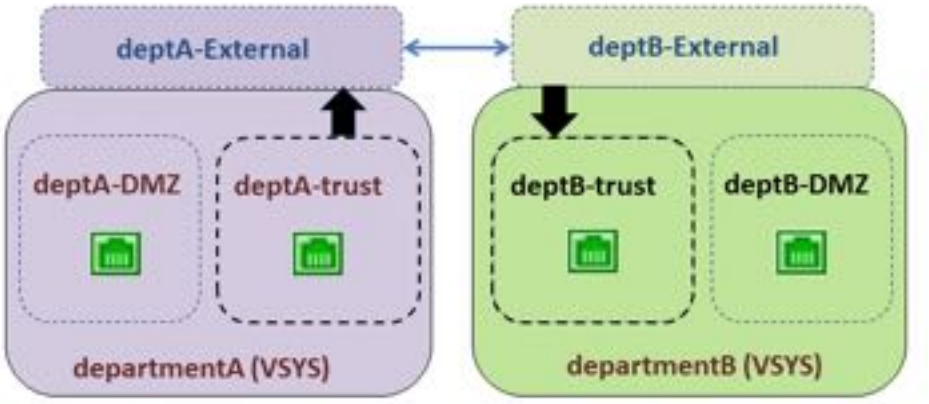
To create external zones, the firewall administrator must configure the virtual systems so that they are *visible* to each other. External zones do not have security policies between them because their virtual systems are visible to each other.

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.



- 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

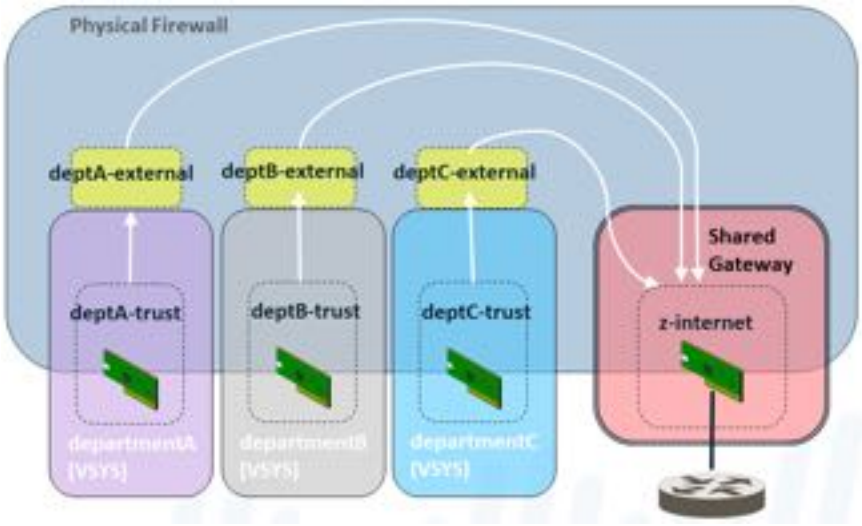
External Zones and Shared Gateway

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 firewall 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

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 firewall.

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.

 © 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**. 2. Select the **Multi Virtual System Capability** check box and click **OK**. This action triggers a commit if you approve it. 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). 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. 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. 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. 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.	1. On the **General** tab, select a **DNS Proxy** object if you want to apply DNS proxy rules to the interface. 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. 3. Do any of the following, based on the deployment type(s) you need in the virtual system: • In the **VLANs** field, click **Add** to enter the VLAN(s) to assign to the vsys. • In the **Virtual Wires** field, click **Add** to enter the virtual wire(s) to assign to the vsys. • In the **Virtual Routers** field, click **Add** to enter the virtual router(s) to assign to the vsys. 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. In a multi-tenancy scenario where strict administrative boundaries are required, no virtual systems would be checked. 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 firewall resources.	1. On the **Resource** tab, optionally set limits for a virtual system. There are no default values. • **Sessions Limit**—Range is 1-262144. • **Security Rules**—Range is 0-2500. • **NAT Rules**—Range is 0-3000. • **Decryption Rules**—Range is 0-250. • **QoS Rules**—Range is 0-1000. • **Application Override Rules**—Range is 0-250. • **Policy Based Forwarding Rules**—Range is 0-500. • **Captive Portal Rules**—Range is 0-1000. • **DoS Protection Rules**—Range is 0-1000. • **Site to Site VPN Tunnels**—Range is 0-1024. • **Concurrent SSL VPN Tunnels**—Range is 0-1024. 2. Click **OK**.
Step 5	Save the configuration.	Click **Commit** and **OK**. The virtual system is now an object accessible from the **Objects** tab.
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. 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**. 2. For **Interfaces,** click **Add** and from the drop-down, select the interfaces that belong to the virtual router. 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.

 © Palo Alto Networks, Inc.

Configure a Virtual System		
Step 8	Configure the security policy rules that allow or deny traffic to and from the zones in the virtual system.	See Create a Security Policy Rule.
Step 9	Save the configuration.	Click **Commit** and **OK**.  After creating a virtual system, you can use the CLI to commit a configuration for only a specific virtual system: `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:
		`set system setting target-vsys <vsys-id>`
		`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**. 2. For **Location**, select the virtual system for which you are creating an external zone. 3. For **Type**, select **External**. 4. For **Virtual Systems**, click **Add** and enter the virtual system that the external zone can reach. 5. **Zone Protection Profile**—Optionally select a zone protection profile (or configure one later) that provides flood, reconnaissance, or packet-based attack protection. 6. **Log Setting**—Optionally select a log forwarding profile for forwarding zone protection logs to an external system. 7. Optionally select the **Enable User Identification** check box to enable User-ID for the external zone. 8. Click **OK**.
Step 2	Configure the Security policy rules to allow or deny traffic from the internal zones to the external zone of the virtual system, and vice versa.	• See Create a Security Policy Rule. • 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**. 2. Enter a helpful **Name**, preferably including the **ID** of the gateway. 3. In the **DNS Proxy** field, select a DNS proxy object if you want to apply DNS proxy rules to the interface. 4. **Add** an **Interface** that connects to the outside world. 5. Click **OK**.
Step 2	Configure the zone for the shared gateway. 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**. 2. For **Location**, select the shared gateway for which you are creating a zone. 3. For **Type**, select **Layer3**. 4. **Zone Protection Profile**—Optionally select a zone protection profile (or configure one later) that provides flood, reconnaissance, or packet-based attack protection. 5. **Log Setting**—Optionally select a log forwarding profile for forwarding zone protection logs to an external system. 6. Optionally select the **Enable User Identification** check box to enable User-ID for the shared gateway. 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 Administrative Access Per Virtual System or Firewall

Customize Service Routes to Services for Virtual Systems

When you enable Multi Virtual System Capability, any virtual system that does not have specific service routes configured inherits the global service and service route settings for the firewall. You can instead configure a virtual system to use a different service route, as described in the following workflow.



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.

 © 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. 2. Click the **Service Route Configuration** link. 3. Select one of the radio buttons: • **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. • **Customize**—Allows you to specify a source interface and source address for each service. 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**. • 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. • **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. • 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. 5. Click **OK**. 6. Repeat steps 4 and 5 to configure source addresses for other external services. 7. Click **OK**.
Step 2 Save the configuration.	Click **Commit** and **OK**. 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

If you have enabled multi virtual system capability on your PA-7000 Series firewall, you can configure logging for different virtual systems as described in the following workflow. 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. 2. Enter the **Interface Name**. 3. For **Interface Type**, select **Log Card** from the drop-down. 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**. 2. For **Interface Name**, after the period, enter the subinterface assigned to the tenant's virtual system. 3. For **Tag**, enter a VLAN tag value.  Make the tag the same as the subinterface number for ease of use, but it could be a different number. 4. (Optional) Enter a **Comment**. 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. 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: • 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]). • For the IPv6 section, enter the **IPv6 Address** assigned to the subinterface. Enter the **IPv6 Default Gateway**. 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 Administrative Access Per Virtual System or Firewall

If you have a superuser administrative account, you can create and configure granular permissions for a vsysadmin or device admin role.

Create an Admin Role Profile Per Virtual System or Firewall		
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**. 2. Enter a **Name** and optional **Description** of the profile. 3. For **Role**, specify which level of control the profile affects: • **Device**—The profile allows the management of the global settings and any virtual systems. • **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.) 4. On the **Web UI** tab for the Admin Role Profile, scroll down to **Device**, and leave the green check mark (Enable). • 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.) – **Management**—Allows an admin with this profile to configure settings on the **Management** tab. – **Operations**—Allows an admin with this profile to configure settings on the **Operations** tab. – **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**. – **Content-ID**—Allows an admin with this profile to configure settings on the **Content-ID** tab. – **WildFire**—Allows an admin with this profile to configure settings on the **WildFire** tab. – **Session**—Allows an admin with this profile to configure settings on the **Session** tab. – **HSM**—Allows an admin with this profile to configure settings on the **HSM** tab. 5. Click **OK**. 6. (Optional) Repeat the entire step to create another Admin Role profile with different permissions, as necessary.

Create an Admin Role Profile Per Virtual System or Firewall (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. 2. (Optional) Select an **Authentication Profile**. 3. (Optional) Select **Use only client certificate authentication (Web)** to have bi-directional authentication; to get the server to authenticate the client. 4. Enter a **Password** and **Confirm Password**. 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. 6. For **Administrator Type**, select **Role Based**. 7. For **Profile**, select the profile that you just created. 8. (Optional) Select a **Password Profile**. 9. Click **OK**.
Step 3	Save the configuration.	Click **Commit** and **OK**.

 © Palo Alto Networks, Inc.

Virtual System Functionality with Other Features

Many firewall 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.

Certifications

The following topics describe how to configure the firewall to support the Common Criteria and the Federal Information Processing Standard 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
- FIPS-CC Security Functions

Enable FIPS and Common Criteria Support

Use the following procedure to enable FIPS-CC mode on a software version that supports Common Criteria and the Federal Information Processing Standards 140-2 (FIPS 140-2). When you enable FIPS-CC mode, all FIPS and CC functionality is included.



When you enable FIPS-CC mode, the firewall will reset to the factory default settings; all configuration will be removed.

Enable FIPS-CC Mode	
Step 1	Boot the firewall into maintenance mode as follows: 1. Establish a serial connection to the console port on the firewall. 2. Enter the following CLI command: `debug system maintenance-mode` 3. Press Enter to continue.  You can also reboot the firewall and enter `maint` at the maintenance mode prompt.
Step 2	Select **Set FIPS-CC Mode** from the menu.
Step 3	Select **Enable FIPS-CC Mode** from the menu.
Step 4	When prompted, select **Reboot**. After successfully switching to FIPS-CC mode, the following status displays: `FIPS-CC mode enabled successfully`. In addition, the following changes will take place: • **FIPS-CC** will display at all times in the status bar at the bottom of the web interface. • The console port functions as a status output port only. • The default admin login credentials change to admin/paloalto.

 © Palo Alto Networks, Inc.

FIPS-CC Security Functions

When FIPS-CC 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.
- Unapproved FIPS/CC 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 FIPS-CC mode is enabled.
- Telnet, TFTP, and HTTP management connections are unavailable.
- High availability (HA) port encryption is required.

 © Palo Alto Networks, Inc.

PAN-OS® New Features Guide

Version 8.1

paloaltonetworks.com/documentation

techDOCS

Contact Information

Corporate Headquarters:
Palo Alto Networks
3000 Tannery Way
Santa Clara, CA 95054
www.paloaltonetworks.com/company/contact-support

About the Documentation

- For the most recent version of this guide or for access to related documentation, visit the Technical Documentation portal www.paloaltonetworks.com/documentation.
- To search for a specific topic, go to our search page www.paloaltonetworks.com/documentation/document-search.html.
- Have feedback or questions for us? Leave a comment on any page in the portal, or write to us at documentation@paloaltonetworks.com.

Copyright

Palo Alto Networks, Inc.
www.paloaltonetworks.com

© 2018-2020 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

Last Revised

June 17, 2020

Table of Contents

Upgrade to PAN-OS 8.1 7
- Upgrade/Downgrade Considerations 9
- Upgrade the Firewall to PAN-OS 8.1 16
  - Determine the Upgrade Path to PAN-OS 8.1 16
  - Upgrade Firewalls Using Panorama 17
  - Upgrade a Standalone Firewall to PAN-OS 8.1 22
  - Upgrade an HA Firewall Pair to PAN-OS 8.1 24
- Downgrade from PAN-OS 8.1 29
  - Downgrade a Firewall to a Previous Feature Release 29
  - Downgrade a Firewall to a Previous Maintenance Release 30
  - Downgrade a Windows Agent from PAN-OS 8.1 30

App-ID Features 33
- SaaS Application Hosting Characteristics 35
- Simplified App-ID 37
  - New App-ID Threshold 37
  - New App-ID Characteristic 38
  - Extended Policy Impact Review for Content Releases 40
  - Coverage Change Details for Modified App-IDs 41
- HTTP Header Insertion and Modification 42
- Service-Based Session Timeouts 43

Virtualization Features 45
- VM-50 Lite 47
- Integration with Azure Security Center 48
- Support for Azure Application Insights 50
- Bootstrapping Enhancements for VM-Series on Azure 52
- VM-Series Firewall on Google Cloud Platform 55
  - Deploy the VM-Series Firewall from Google Cloud Platform Marketplace 55
  - Enable Google Stackdriver Monitoring 56
  - Enable VM-Series Firewall to Track Changes on Google Cloud Platform VMs 57

Decryption Features 59
- Decryption Broker 61
- Automatic SAN Support for SSL Decryption 62
- HSM Client Upgrade and SafeNet HSM Cluster Support 63
- ECDSA Certificate Support for SSL Decryption with HSMs 64
- Decryption Port Mirroring Support Extension 65

Panorama Features 67
- Device Monitoring on Panorama 69
- Configuration Reusability for Templates and Template Stacks 71
- Panorama Virtual Appliance and Virtual Dedicated Log Collector on AWS 73
- Panorama Virtual Appliance and Virtual Dedicated Log Collector on AWS GovCloud 75
- Panorama Virtual Appliance and Virtual Dedicated Log Collector on Azure 77

Panorama Virtual Appliance and Virtual Dedicated Log Collector on Google Cloud Platform....78
Panorama Virtual Appliance and Virtual Dedicated Log Collector on KVM....80
Panorama Virtual Appliance and Virtual Dedicated Log Collector on Hyper-V....82
Virtual Dedicated Log Collector on ESXi....84
Streamlined Panorama Deployment for Application and Threat Content Updates....85
Content Update Revert from Panorama....87
Direct Query of PA-7000 Series Firewalls from Panorama....88
Panorama Interconnect....89

Content Inspection Features....91
SCTP Security....93
Rapid Deployment of the Latest Threat Prevention....95
Tools to Avoid or Mitigate Content Update Issues....96
Option to Hold Web Requests During URL Category Lookup....98

Authentication Features....99
Extensible Authentication Protocol (EAP) Support for RADIUS....101
Authentication Using Custom Certificates for WildFire, PAN-DB, and Between Log Collectors....104

GlobalProtect Features....107
Optimized Split Tunneling for GlobalProtect....109
Kerberos Authentication Support for macOS....114
SAML SSO for GlobalProtect on Chromebooks....119
OPSWAT SDK V4 Support....121
GlobalProtect Tunnel Preservation on User Logout....123
Automatic Launching of Web Browser in Captive Portal Environment....125

Management Features....127
Rule Usage Tracking....129
Configuration Table Export....131
Reporting Engine Enhancements....132
Enhanced Application Logging....134

Networking Features....135
Tunnel Content Inspection Logging....137
Dynamic IP Address Support for Destination NAT....138
FQDN Support for IKE Gateway Peer IP Address....140
Configuration Capacity Improvements....141
Refresh of Default Trusted CAs....143
ARP Cache Timeout....144

User-ID Features....145
Support for Multiple Username Formats....147

WildFire Features....151

WildFire Appliance-to-Appliance Encryption..... 153

Upgrade to PAN-OS 8.1

> Upgrade/Downgrade Considerations
> Upgrade the Firewall to PAN-OS 8.1
> Downgrade from PAN-OS 8.1

© 2020 Palo Alto Networks, Inc.

Upgrade/Downgrade Considerations

The following table lists the new features that have upgrade or downgrade impacts. Make sure you understand all potential changes before you upgrade to or downgrade from a PAN-OS 8.1 release. For additional information about PAN-OS 8.1 releases, refer to the PAN-OS 8.1 Release Notes.

For M-100 appliances running in Panorama mode, Palo Alto Networks recommends upgrading the memory to 32GB to avoid the risk of running out of memory for management and log collection tasks. See M-100 Memory Upgrade Guide *for more information.*

Table 1: PAN-OS 8.1 Upgrade/Downgrade Considerations

Feature	Upgrade Considerations	Downgrade Considerations
Support for Third-Party SFP Transceivers	A small percentage of writable third-party SFP transceivers (not purchased from Palo Alto Networks) can stop working or experience other issues after you upgrade the firewall to which they are connected to a PAN-OS 8.0 or PAN-OS 8.1 release. Because it is typically impossible to know if a third-party SFP is writable, Palo Alto Networks® recommends that, if your firewall uses third-party SFPs, you do not upgrade to a PAN-OS 8.0 or PAN-OS 8.1 release until you are able to upgrade to a maintenance release for each that addresses this issue. Additionally, when you are ready to upgrade, make sure that you do not reboot the firewall after you download and install the PAN-OS 8.1 (or PAN-OS 8.0) base image until after you download and install a maintenance release that contains the fix for this issue. This is true when upgrading to PAN-OS 8.1 even if you addressed this issue when you upgraded to a PAN-OS 8.0 maintenance release that includes this fix. For more information about this known issue and maintenance releases related to this issue, refer to the PAN-OS 8.0 or PAN-OS 8.1 release notes, as appropriate.	
Support for Multiple Username Formats	During an upgrade, the User Name in the User Objects section for the Group Mapping profile will be	During a downgrade, the **Primary Username** is used to retrieve the user name and the **Allow matching**

© 2020 Palo Alto Networks, Inc.

Feature	Upgrade Considerations	Downgrade Considerations
	migrated as the **Primary Username** for users. During an upgrade, the firewall clears the user and group mappings. After the upgrade is complete, the firewall must relearn the mappings before it can apply user-based policies. If the firewall uses the user@domain format for the primary username, and is in an HA configuration, and has 500,000 or more users, the User-ID manager resets when it relearns the group mappings after the upgrade.	**users without domain** option and all alternate usernames are removed.
Extensible Authentication Protocol (EAP) Support for RADIUS	Auto has been removed from list of available authentication protocols for a RADIUS Server Profile. During the upgrade, all existing RADIUS Server Profiles using PAP or CHAP will continue to use the selected authentication protocol. If the RADIUS server profile used Auto, the firewall will attempt to change to CHAP or PAP, based on which protocol was in use before the upgrade. If the firewall cannot determine which one was used, CHAP is selected. After you upgrade, Panorama templates use CHAP as the default authentication protocol. If your Panorama templates used the Auto authentication protocol and your RADIUS server requires PAP, update your Panorama template to use PAP to avoid having to log out and log in again after the upgrade.	Any RADIUS Server Profiles configured to use EAP (PEAP-MSCHAPv2, PEAP with GTC, or EAP-TTLS with PAP) will be migrated to Auto (where CHAP is attempted first, then PAP) Any additional options associated with the EAP protocol, such as Outer Identity and Certificate Profiles, will be removed from the configuration.
HSM Client Upgrade and SafeNet HSM Cluster Support	PAN-OS 8.1 upgrades nCipher nShield client version 11.62 to version 12.30. You can upgrade the firewall to use SafeNet client version 5.4.2 or 6.2.2. These HSM client versions provide necessary compatibility with HSM server versions. With an upgraded HSM client version on PAN-OS, you must ensure the corresponding HSM server version is also upgraded. *Downgrading the HSM server versions*	A PAN-OS release is tied to an HSM client version, so if you downgrade from PAN-OS 8.1 to a release prior to PAN-OS 8.0.2 (for example, PAN-OS 8.0.0 or 8.0.1), the downgrade will fail. Downgrading from PAN-OS 8.1 to 8.0.2 or a later release will succeed and the HSM client version will be retained.

© 2020 Palo Alto Networks, Inc.

Feature	Upgrade Considerations	Downgrade Considerations
	might not be an option after they are upgraded. Check your HSM vendor documentation for any upgrade/downgrade considerations.	
FQDN Support for IKE Gateway Peer IP Address		Downgrading will remove the address objects, address groups, and FQDNs from an IKE VPN Peer IP address, so unless the peer IP address is a static IP address, the downgrade will fail.
Dynamic IP Address Support for Destination NAT		Downgrading means that a translated destination NAT address specifying **Dynamic IP (with session distribution)** isn't supported, so the downgrade will fail.
FQDN Refresh Time		If you downgrade a VM-Series firewall that used an FQDN Refresh Time in the range 60-599 seconds, the downgrade will fail because the downgraded release doesn't support an FQDN refresh time that fast; it supports a range of 600-14,399 seconds. Select **Device** > **Setup** > **Services**, change the 8.1 FQDN Refresh Time to a value in the range 600-14,399 seconds, commit, and then downgrade.
DNS Suffix Support		In PAN-OS 8.0.x and earlier releases, you can only add up to 10 DNS suffixes to the GlobalProtect gateway configuration (**Network** > **GlobalProtect** > **Gateways** > **<gateway-config>** > **GlobalProtect Gateway Configuration** > **Agent** > **Network Services** > **DNS Suffix**). You must remove additional DNS suffixes from the configuration prior to downgrading from PAN-OS 8.1.0 to PAN-OS 8.0.x or earlier releases. In Panorama 8.1.0 and earlier releases, you can only push up to 10 DNS suffixes to firewalls running PAN-OS 8.0.x and earlier releases.

© 2020 Palo Alto Networks, Inc.

Feature	Upgrade Considerations	Downgrade Considerations
OPSWAT SDK V4 Support	PAN-OS 8.1 supports OPSWAT SDK V4. When you upgrade to PAN-OS 8.1, the **Antivirus** and **Anti-Spyware** HIP categories merge to form the new **Anti-Malware** category. Vendor and product names are also updated for OPSWAT SDK V4 in all HIP categories. The latest GlobalProtect data file must be installed on your firewall following an upgrade to PAN-OS 8.1. Until the data file is installed, HIP functionality is not available. For PAN-OS 8.1, the GLobalProtect data file is for OPSWAT SDK V4. GlobalProtect data file installation can only be triggered when a data file update schedule is configured under **Device** > **Dynamic Updates** > **GlobalProtect Data File**. We recommend that you set the **Schedule** recurrence to **Hourly**.	PAN-OS 8.0 and earlier releases do not support OPSWAT SDK V4. When you downgrade to an earlier version of PAN-OS, the **Anti-Malware** category is dropped. You must manually reconfigure the **Antivirus** and **Anti-Spyware** categories to display this information in the OPSWAT SDK V3 format. After your system is rebooted, all HIP match logs are persistent, while all existing HIP reports are deleted. The latest GlobalProtect data file must be installed on your firewall following a downgrade to PAN-OS 8.0 or earlier releases. Until the data file is installed, HIP functionality is not available. For PAN-OS 8.0 and earlier releases, the GlobalProtect data file is for OPSWAT SDK V3. GlobalProtect data file installation can only be triggered when a data file update schedule is configured under **Device** > **Dynamic Updates** > **GlobalProtect Data File**. We recommend that you set the **Schedule** recurrence to **Hourly**. When configurations are pushed from Panorama 8.1 to a PAN-OS 8.0 or earlier firewall, the **Anti-Malware** category is dropped from the HIP Match log and HIP Object. You must manually reconfigure the **Antivirus** and **Anti-Spyware** categories to display this information in the OPSWAT SDK V3 format. Vendor and product names may also be different for OPSWAT SDK V3, so you must review and manually reconfigure as necessary.
Content Revert from Panorama		When downgrading a firewall from PAN-OS 8.1 to an earlier PAN-OS release, Panorama will display the content version installed on the firewall while it was running PAN-OS 8.1 regardless of the version currently running on the firewall.
Support for Panorama Virtual Appliances in New Environments	On upgrade to PAN-OS 8.1, you are no longer able to change in to Legacy mode. If your Panorama	

© 2020 Palo Alto Networks, Inc.

Feature	Upgrade Considerations	Downgrade Considerations
	virtual appliance is in Legacy mode on upgrade, the mode will be preserved. However, if you change the Panorama virtual appliance mode, you cannot change back into Legacy mode.	
Management Only Mode		When downgrading a Panorama appliance from PAN-OS 8.1 to an earlier PAN-OS release, the Panorama appliance must be in Panorama mode or Log Collector mode. Downgrade from PAN-OS 8.1 is not supported for a Panorama appliance in Management Only mode.
Configuration Reusability for Templates and Template Stacks	On upgrade to PAN-OS 8.1, templates with devices attached to them will be converted into template stacks that contain the original template and have the device attached to the template stack. The template stack name will be `<Template Name>_mig_stack`. The number of supported templates in a template stack is reduced from 16 to 8.	You are unable to downgrade from PAN-OS 8.1 to an earlier PAN-OS release if variables are used in your template or template stack configuration. Variables must be removed from the template and template stack configuration to downgrade.
VM-Series Firewall for VMware NSX	You must install VMware NSX plugin 2.0.1 before upgrading Panorama to 8.1. You must select a template stack when configuring a service definition for your VM-Series for NSX deployment. Selecting a template causes a commit failure when attempting to commit the changes to Panorama.	
Reporting Engine Enhancements	In PAN-OS 8.1, firewalls write logs from the perspective of who initiated the network connection, resulting in the source (SRC) and destination (DST) values being swapped when compared to PAN-OS 8.0 and earlier releases. When writing queries in the **ACC** tab, **Monitor** tab, and Custom Reports (**Panorama** > **Monitor** > **Manage Custom Reports**), you must swap the SRC and DST values based on the direction of the threat to ensure that the displayed query results are accurate.	

© 2020 Palo Alto Networks, Inc.

Feature	Upgrade Considerations	Downgrade Considerations
	You must swap the SRC and DST values for in your auto-tag configuration based on the direction of the threat to ensure Threat, URL, Data Filtering, and WildFire Submissions logs are correctly tagged.	
	On upgrade to PAN-OS 8.1, all existing Threat, URL Filtering, Data Filtering, and Wildfire Submission log data on Log Collectors preserve original log format where the DST value is the IP address that initiated the connection, and the SRC is the IP address with whom the connection is initiated that exists in PAN-OS 8.0 and earlier releases.	
BGP Minimum Route Advertisement Interval	If you upgrade from PAN-OS 8.0.11 (or a later PAN-OS 8.0 release) to a PAN-OS 8.1 release, the CLI operational command **`set system setting bgp-mrai-timer value`** is deprecated and value configured in the CLI is gone; the BGP peers revert to having a minimum route advertisement interval of 30 seconds. Configure the minimum route advertisement interval for a BGP peer in the user interface instead (**Network** > **Virtual Routers** > **virtual router** > **BGP** > **Peer Group** > **Peer** > **Connection Options**).	If you downgrade from PAN-OS 8.1 to PAN-OS 8.0.11 (or a PAN-OS 8.0 release later than PAN-OS 8.0.11), the minimum route advertisement interval you configured for a BGP peer no longer applies to the peer and the user interface to configure the BGP minimum router advertisement interval does not exist. The default value of 30 seconds applies to all BGP peers. Use the CLI operational command **`setsystem setting bgp-mrai-timer value`** to set an interval for all BGP peers.
URL Filtering Option to Hold Web Requests During URL Category Lookup (PAN-OS 8.1.10 or later)	• If you have this feature enabled in PAN-OS 8.1.10 or later, upgrading will disable it. To re-enable it, see Configure URL Filtering. • The 8.1 and 9.0 versions of this feature have different CLI commands to enable it. For the correct commands, see the appropriate version of Configure URL Filtering.	• If you have this feature enabled, downgrading to an earlier version will disable it. • This feature is not available in PAN-OS versions 9.0 to 9.0.3 and 8.1.9 or earlier. • The 8.1 and 9.0 versions of this feature have different CLI commands to enable it. For the correct commands, see the appropriate version of Configure URL Filtering.
Identity Provider Certificate (PAN-OS 8.1.15 or later)	Ensure that you configure the signing certificate for your SAML Identity Provider as the **Identity Provider Certificate** before you upgrade to PAN-OS 8.1.15 or later so that your	

© 2020 Palo Alto Networks, Inc.

Feature	Upgrade Considerations	Downgrade Considerations
	users can continue to authenticate successfully. Always configure the Identity Provider Certificate when you configure your SAML authentication and, as a best practice, enable certificate validation when available.	

© 2020 Palo Alto Networks, Inc.

Upgrade the Firewall to PAN-OS 8.1

How you upgrade to PAN-OS 8.1 depends on whether you have standalone firewalls or firewalls in a high availability (HA) configuration and, for either scenario, whether you use Panorama to manage your firewalls. Review the PAN-OS 8.1 Release Notes and then follow the procedure specific to your deployment:

- Determine the Upgrade Path to PAN-OS 8.1
- Upgrade Firewalls Using Panorama
- Upgrade a Standalone Firewall to PAN-OS 8.1
- Upgrade an HA Firewall Pair to PAN-OS 8.1

When upgrading firewalls that you manage with Panorama or firewalls that are configured to forward content to a WF-500 appliance, you must first upgrade Panorama and its Log Collectors and then upgrade the WildFire appliance before you upgrade the firewalls.

Additionally, it is not recommended to manage firewalls running a later maintenance release than Panorama as this may result in features not working as expected. For example, it is not recommended to manage firewalls running PAN-OS 8.1.3 or later maintenance releases if Panorama is running PAN-OS 8.1.2.

Determine the Upgrade Path to PAN-OS 8.1

When you upgrade from one PAN-OS feature release version to a later feature release, you cannot skip the installation of any feature release versions in the path to your target release (for example, to get from a 7.1 release version to 8.1, you would also need to install an 8.0 release version). In most cases, the recommended path when moving from one feature release to the next is to download the base image for the next feature release version and then download and install your target maintenance release version. You must have both the base image and the maintenance release image on the firewall for installation to be successful. If the firewall doesn't have enough disk space to hold both the base image and the maintenance release image, it will delete the base image and you will see an error message when you attempt to install the maintenance release. In this case you will need to re-download and install the base image before you can download and install the maintenance release. To minimize downtime for your users, perform upgrades during non-business hours.

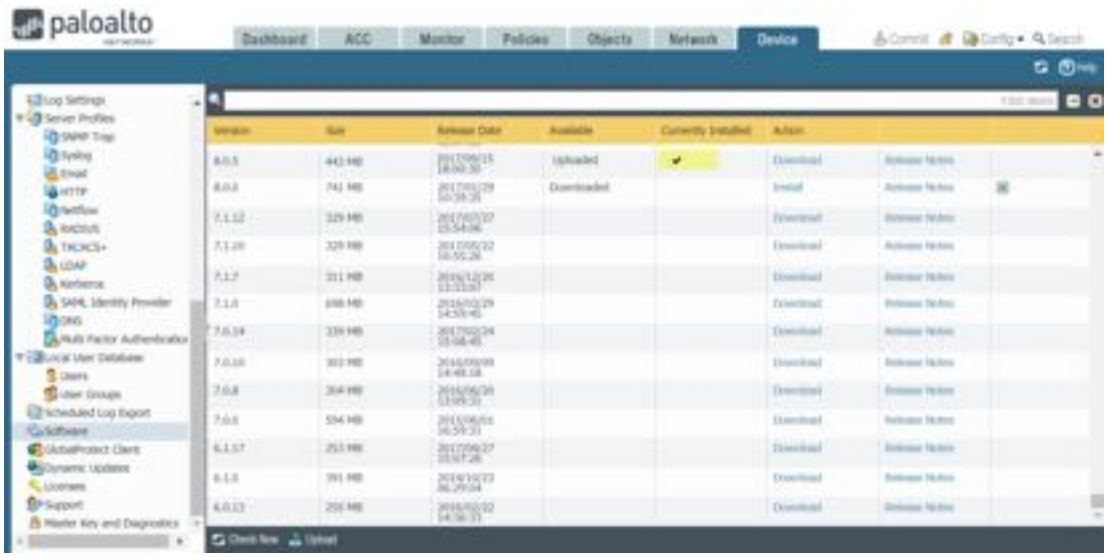
For manual upgrades, you must install the base image for a feature release before you upload and install a maintenance release image.

Determine the upgrade path as follows:

STEP 1 | Identify which version is currently installed.

- From Panorama, select **Panorama** > **Managed Devices** and check the Software Version on the firewalls you plan to upgrade.
- From the firewall, select **Device** > **Software** and check which version has a check mark in the Currently Installed column.

© 2020 Palo Alto Networks, Inc.



STEP 2 | Identify the upgrade path:

Review the known issues and changes to default behavior in the Release Notes *and upgrade/downgrade considerations in the* New Features Guide *for each release through which you pass as part of your upgrade path.*

Installed PAN-OS Version	Recommended Upgrade Path to PAN-OS 8.1
8.0.x	If you are already running a PAN-OS 8.0 release, download and install the latest preferred PAN-OS 8.0 maintenance release and reboot. You can then proceed to Upgrade the Firewall to PAN-OS 8.1.
7.1.x	❑ Download and install the latest preferred PAN-OS 7.1 maintenance release and reboot. ❑ Download PAN-OS-8.0.0. ❑ Download and install the latest preferred PAN-OS 8.0 maintenance release and reboot. Proceed to Upgrade the Firewall to PAN-OS 8.1.
7.0.x	❑ Download PAN-OS 7.1.0. ❑ Download and install the latest preferred 7.1.x maintenance release and reboot. ❑ Download PAN-OS 8.0.0 . ❑ Download and install the latest preferred 8.0.x maintenance release and reboot. Proceed to Upgrade the Firewall to PAN-OS 8.1.

Upgrade Firewalls Using Panorama

Review the PAN-OS 8.1 Release Notes and then use the following procedure to upgrade firewalls that you manage with Panorama. This procedure applies to standalone firewalls and firewalls deployed in a high availability (HA) configuration.

© 2020 Palo Alto Networks, Inc.

If Panorama is unable to connect directly to the updates server, follow the procedure for deploying updates to firewalls when Panoramais not internet-connected so that you can manually download images to Panorama and then distribute the images to firewalls.

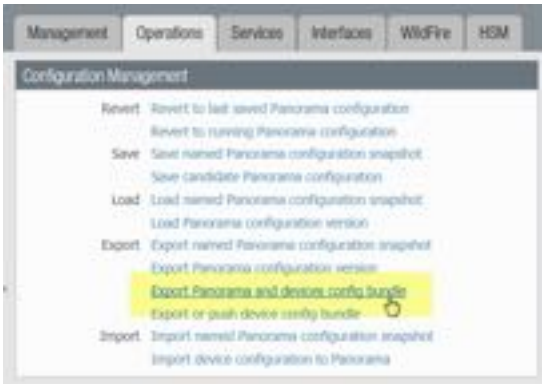
Before you can upgrade firewalls from Panorama, you must:

- Make sure Panorama is running the same or a later PAN-OS version than you are upgrading to. You must upgrade Panorama and its Log Collectors to 8.1 before upgrading the managed firewalls to this version.
- Ensure that firewalls are connected to a reliable power source. A loss of power during an upgrade can make a firewall unusable.

STEP 1 | Save a backup of the current configuration file on each managed firewall you plan to upgrade.

Although the firewall automatically creates a configuration backup, it is a best practice to create and externally store a backup before you upgrade.

1. From the Panorama web interface, select **Panorama** > **Setup** > **Operations** and click **Export Panorama and devices config bundle** to generate and export the latest configuration backup of Panorama and of each managed appliance.

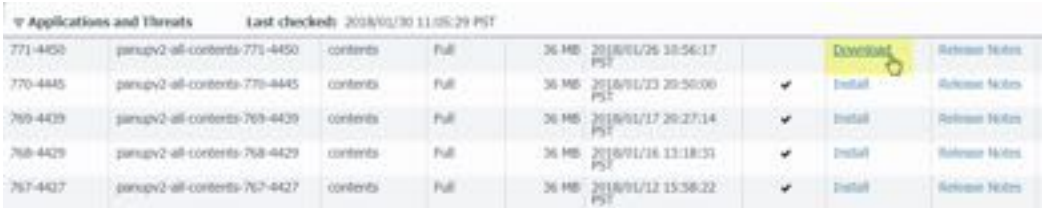


2. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the upgrade.

STEP 2 | Update the content release version on the firewalls you plan to upgrade.

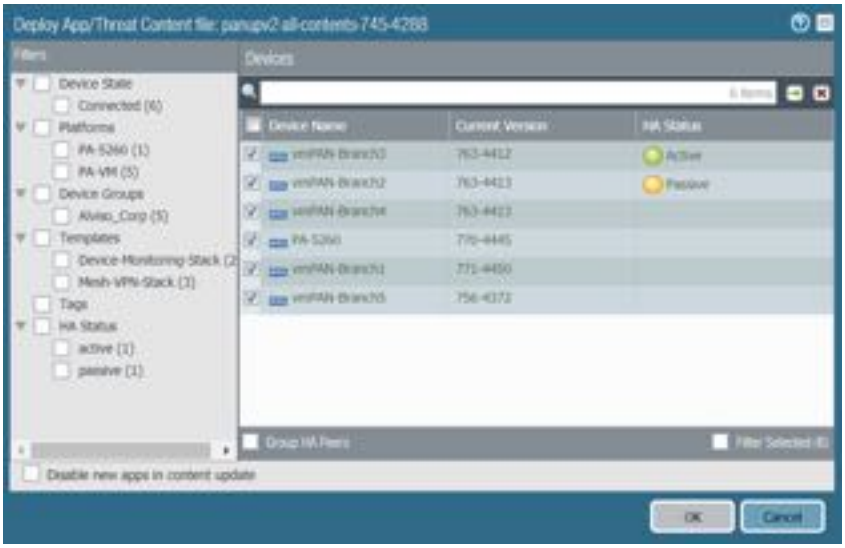
Refer to the Release Notes for the minimum content release version required for PAN-OS 8.1. Make sure to follow the Best Practices for Application and Threat ContentUpdates when deploying content updates to Panorama and managed firewalls.

1. Select **Panorama** > **Device Deployment** > **Dynamic Updates** and **Check Now** for the latest updates. If an update is available, the Action column displays a **Download** link.



2. If not already installed, **Download** the latest content release version.
3. Click **Install**, select the firewalls on which you want to install the update, and click **OK**. If you are upgrading HA firewalls, you must update content on both peers.

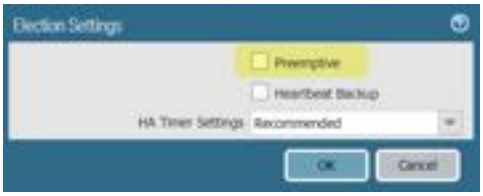
© 2020 Palo Alto Networks, Inc.



By default, you can upload a maximum of two software or content updates of each type to a Panorama appliance and if you download a third update of the same type, Panorama will delete the update for the earliest version of that type. If you need to upload more than two software updates or content updates of a single type, use the ***set max-num-images count <number>*** *CLI command to increase the maximum.*

STEP 3 | (HA firewall upgrades only) If you will be upgrading firewalls that are part of an HA pair, disable preemption. You need only disable this setting on one firewall in each HA pair.

1. Select **Device** > **High Availability** and edit the **Election Settings**.
2. If enabled, disable (clear) the **Preemptive** setting and click **OK**.



3. **Commit** your change. Make sure the commit is successful before you proceed with the upgrade.

STEP 4 | Determine the Upgrade Path to PAN-OS 8.1

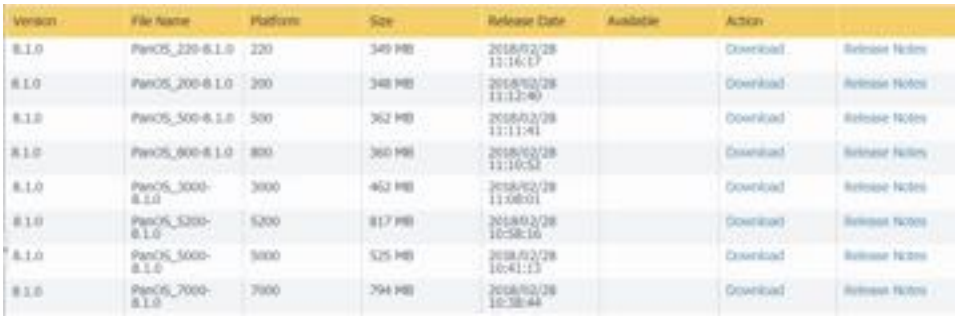
You cannot skip installation of any feature release versions in the path from the currently running PAN-OS version to PAN-OS 8.1.0. Review the known issues and changes to default behavior in the Release Notes and upgrade/downgrade considerations in the New Features Guide for each release through which you pass as part of your upgrade path.

If upgrading more than one firewall, streamline the process by determining upgrade paths for all firewalls before you start downloading images.

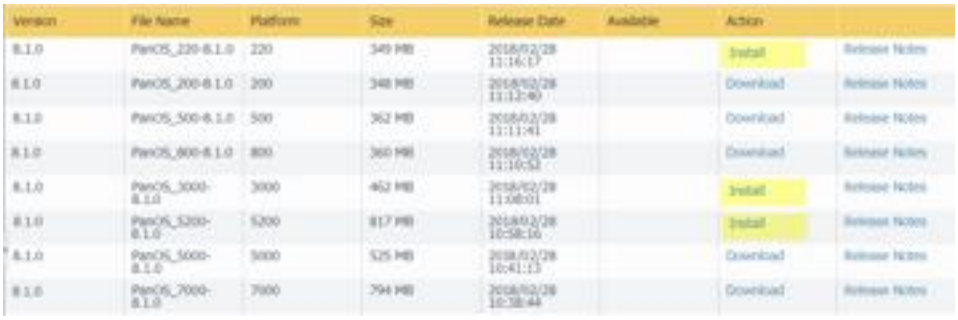
STEP 5 | Download the target PAN-OS 8.1 release image.

1. Select **Panorama** > **Device Deployment** > **Software** and **Check Now** for the latest release versions.
2. **Download** the firewall-specific file (or files) for the release version to which you are upgrading. You must download a separate installation file for each firewall model (or firewall series) that you intend to upgrade.

© 2020 Palo Alto Networks, Inc.

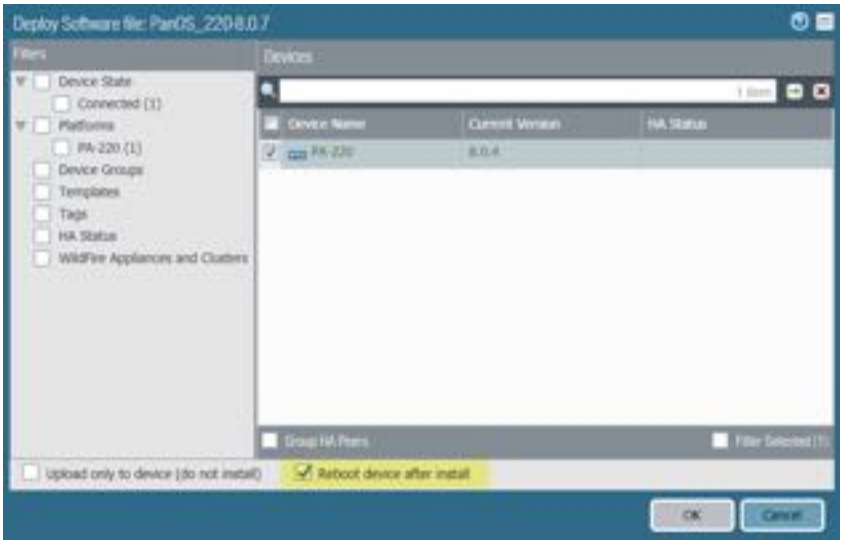


For example, to upgrade your PA-220, PA-3050, and PA-5250 firewalls to PAN-OS 8.1.0, download the `PanOS_220-8.1.0`, `PanOS_3000-8.1.0`, and `PanOS_5200-8.1.0` images. After you successfully download an image, the Action column changes to **Install** for that image.



STEP 6 | Install the PAN-OS 8.1 software update on the firewalls.

1. Click **Install** in the Action column that corresponds to the firewall models you want to upgrade. For example, if you want to upgrade your PA-220 firewalls, click **Install** in the row that corresponds to PanOS_220-8.1.0.
2. In the Deploy Software file dialog, select all firewalls that you want to upgrade. To reduce downtime, select only one peer in each HA pair. For active/passive pairs, select the passive peer; for active/active pairs, select the active-secondary peer.
3. (HA firewall upgrades only) Make sure **Group HA Peers** is not selected.
4. Select **Reboot device after install**.
5. To begin the upgrade, click **OK**.



6. After the installation completes successfully, reboot using one of the following methods:

© 2020 Palo Alto Networks, Inc.

- If you are prompted to reboot, click **Yes**.
- If you are not prompted to reboot, select **Device** > **Setup** > **Operations** and **Reboot Device**.

7. After the firewalls finish rebooting, select **Panorama** > **Managed Devices** and verify the Software Version is 8.1.0 for the firewalls you upgraded. Also verify that the HA status of any passive firewalls you upgraded is still passive.

STEP 7 | (HA firewall upgrades only) Upgrade the second HA peer in each HA pair.

1. (Active/passive upgrades only) Suspend the active device in each active/passive pair you are upgrading.
   1. Switch context to the active firewall.
   2. In the High Availability widget on the **Dashboard**, verify that **Local** firewall state is **Active** and the **Peer** is **Passive**).

   3. Select **Device** > **High Availability** > **Operational Commands** > **Suspend local device**.

   4. Go back to the High Availability widget on the **Dashboard** and verify that **Local** changed to **Passive** and **Peer** changed to **Active**.

2. Go back to the Panorama context and select **Panorama** > **Device Deployment** > **Software**.
3. Click **Install** in the Action column that corresponds to the firewall models of the HA pairs you are upgrading.
4. In the Deploy Software file dialog, select all firewalls that you want to upgrade. This time, select only the peers of the HA firewalls you just upgraded.
5. Make sure **Group HA Peers** is not selected.
6. Select **Reboot device after install**.
7. To begin the upgrade, click **OK**.
8. After the installation completes successfully, reboot using one of the following methods:
   - If you are prompted to reboot, click **Yes**.
   - If you are not prompted to reboot, select **Device** > **Setup** > **Operations** and **Reboot Device**.
9. (Active/passive upgrades only)From the CLI of the peer you just upgraded, run the following command to make the firewall functional again:

   ```
   request high-availability state functional
   ```

STEP 8 | Verify the software and content release version running on each managed firewall.

1. On Panorama, select **Panorama** > **Managed Devices**.
2. Locate the firewalls and review the content and software versions in the table.

   For HA firewalls, you can also verify that the HA Status of each peer is as expected.

© 2020 Palo Alto Networks, Inc.



STEP 9 | (HA firewall upgrades only)If you disabled preemption on one of your HA firewalls before you upgraded, then edit the **Election Settings** (**Device** > **High Availability**) and re-enable the **Preemptive** setting for that firewall and then **Commit** the change.

Upgrade a Standalone Firewall to PAN-OS 8.1

Review the PAN-OS 8.1 Release Notes and then use the following procedure to upgrade a firewall that is not in an HA configuration to PAN-OS 8.1.

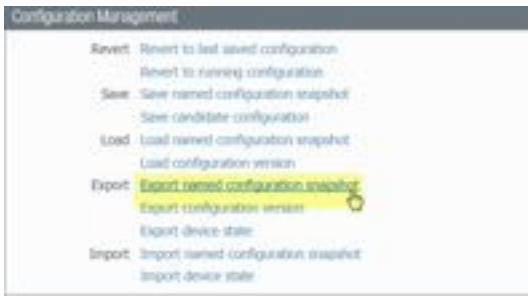
If your firewalls are configured to forward samples to a WF-500 appliance for analysis, you must upgrade the WildFire appliance to PAN-OS 8.1 before upgrading the forwarding firewalls.

To avoid impacting traffic, plan to upgrade within the outage window. Ensure the firewall is connected to a reliable power source. A loss of power during an upgrade can make the firewall unusable.

STEP 1 | Save a backup of the current configuration file.

Although the firewall automatically creates a configuration backup, it is a best practice to create and externally store a backup before you upgrade.

1. Select **Device** > **Setup** > **Operations** and click **Export named configuration snapshot**.



2. Select the XML file that contains your running configuration (for example, **running-config.xml**) and click **OK** to export the configuration file.



3. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the upgrade.

© 2020 Palo Alto Networks, Inc.

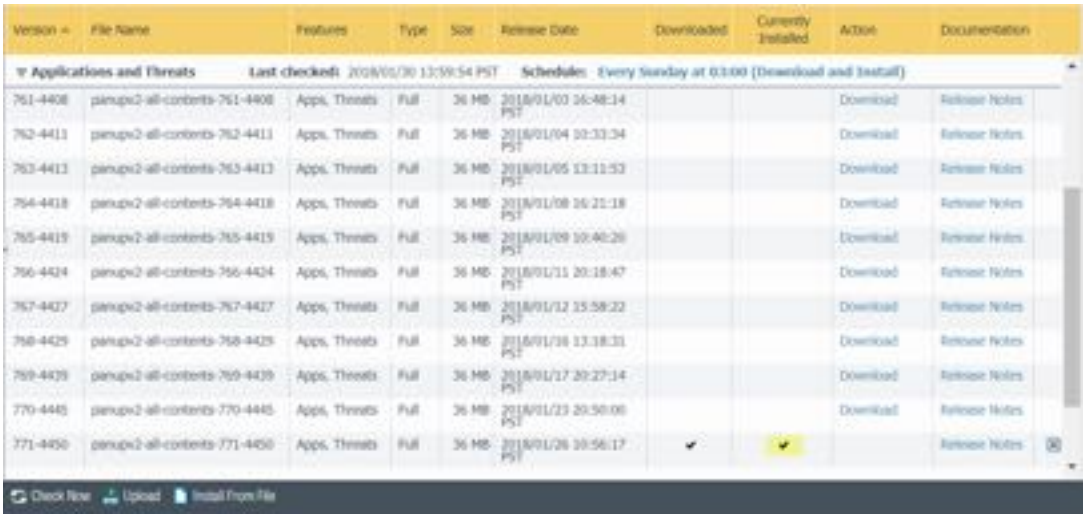
STEP 2 | If you have enabled User-ID, after you upgrade, the firewall clears the current IP address-to-username and group mappings so that they can be repopulated with the attributes from the User-ID sources. To estimate the time required for your environment to repopulate the mappings, run the following CLI commands on the firewall.

- For IP address-to-username mappings:
  - **`show user user-id-agent state all`**
  - **`show user server-monitor state all`**
- For group mappings: **`show user group-mapping statistics`**

STEP 3 | Ensure that the firewall is running the latest content release version.

Refer to the Release Notes for the minimum content release version you must install for a PAN-OS 8.1 release. Make sure to follow the Best Practices for Application and Threat Content Updates.

1. Select **Device** > **Dynamic Updates** and check which **Applications** or **Applications and Threats** to determine which update is Currently Installed.



2. If the firewall is not running the minimum required content release version or a later version required for PAN-OS 8.1, **Check Now** to retrieve a list of available updates.
3. Locate and **Download** the desired content release version.

   After you successfully download a content update file, the link in the Action column changes from **Download** to **Install** for that content release version.
4. **Install** the update.

STEP 4 | Determine the Upgrade Path to PAN-OS 8.1

You cannot skip installation of any feature release versions in the path from the currently running PAN-OS version to PAN-OS 8.1.0.

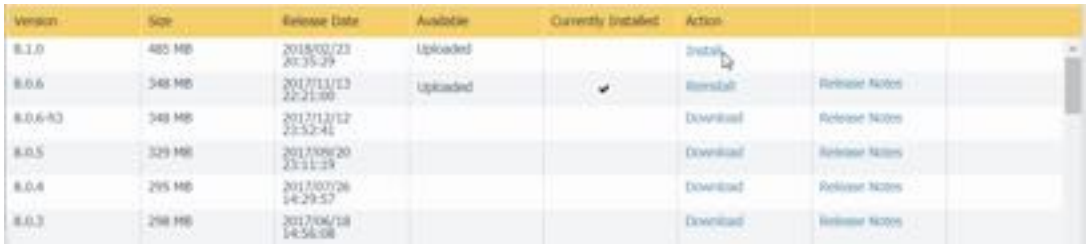
Review the known issues and changes to default behavior in the Release Notes and upgrade/downgrade considerations in the New Features Guide for each release through which you pass as part of your upgrade path.

STEP 5 | Upgrade to PAN-OS 8.1.

If your firewall does not have internet access from the management port, you can download the software image from the Palo Alto Networks Support Portal and then manually Upload it to your firewall.

© 2020 Palo Alto Networks, Inc.

1. Select **Device** > **Software** and click **Check Now** to display the latest PAN-OS updates.
2. Locate and **Download** PAN-OS 8.1.0.
3. After you download the image (or, for a manual upgrade, after you upload the image), **Install** the image.



4. After the installation completes successfully, reboot using one of the following methods:
   - If you are prompted to reboot, click **Yes**.
   - If you are not prompted to reboot, select **Device** > **Setup** > **Operations** and click **Reboot Device**.

   At this point, the firewall clears the User-ID mappings, then connects to the User-ID sources to repopulate the mappings.
5. If you have enabled User-ID, use the following CLI commands to verify that the firewall has repopulated the IP address-to-username and group mappings before allowing traffic.
   - **`show user ip-user-mapping all`**
   - **`show user group list`**

STEP 6 | Verify that the firewall is passing traffic.

Select **Monitor** > **Session Browser** and verify that you are seeing new sessions.



Upgrade an HA Firewall Pair to PAN-OS 8.1

Review the PAN-OS 8.1 Release Notes and then use the following procedure to upgrade a pair of firewalls in a high availability (HA) configuration. This procedure applies to both active/passive and active/active configurations.

To avoid downtime when upgrading firewalls that are in a high availability (HA) configuration, update one HA peer at a time: For active/active firewalls, it doesn't matter which peer you upgrade first (though for simplicity, this procedure shows you how to upgrade the active-secondary peer first). For active/passive firewalls, you must upgrade the passive peer first, suspend the active peer (fail over), update the active peer, and then return that peer to a functional state (fail back). To prevent failover during the upgrade of the HA

© 2020 Palo Alto Networks, Inc.

peers, you must make sure preemption is disabled before proceeding with the upgrade. You only need to disable preemption on one peer in the pair.

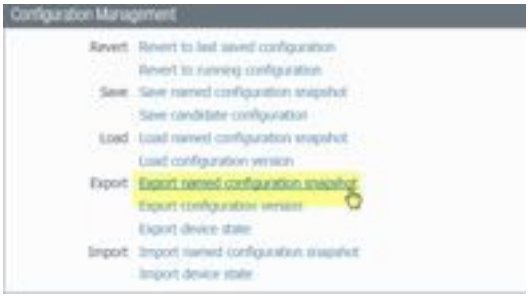
To avoid impacting traffic, plan to upgrade within the outage window. Ensure the firewalls are connected to a reliable power source. A loss of power during an upgrade can make firewalls unusable.

STEP 1 | Save a backup of the current configuration file.

Although the firewall automatically creates a backup of the configuration, it is a best practice to create and externally store a backup before you upgrade.

Perform these steps on each firewall in the pair:

1. Select **Device** > **Setup** > **Operations** and click **Export named configuration snapshot**.



2. Select the XML file that contains your running configuration (for example, **running-config.xml**) and click **OK** to export the configuration file.



3. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the upgrade.

STEP 2 | If you have enabled User-ID, after you upgrade, the firewall clears the current IP address-to-username and group mappings so that they can be repopulated with the attributes from the User-ID sources. To estimate the time required for your environment to repopulate the mappings, run the following CLI commands on the firewall.

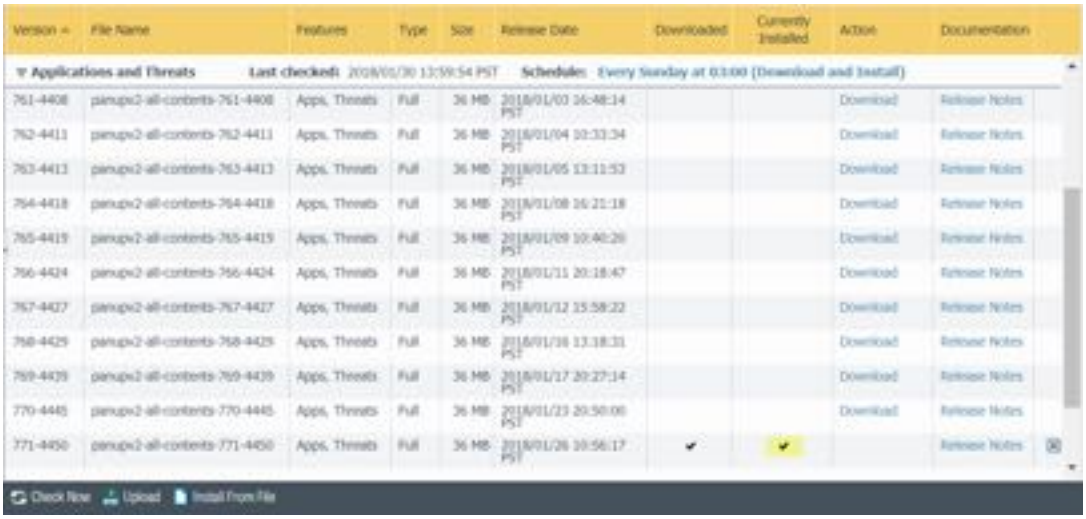
- For IP address-to-username mappings:
  - **`show user user-id-agent state all`**
  - **`show user server-monitor state all`**
- For group mappings: **`show user group-mapping statistics`**

STEP 3 | Ensure that each firewall in the HA pair is running the latest content release version.

Refer to the Release Notes for the minimum content release version you must install for a PAN-OS 8.1 release. Make sure to follow the Best Practices for Application and Threat Content Updates.

1. Select **Device** > **Dynamic Updates** and check which **Applications** or **Applications and Threats** to determine which update is Currently Installed.

© 2020 Palo Alto Networks, Inc.

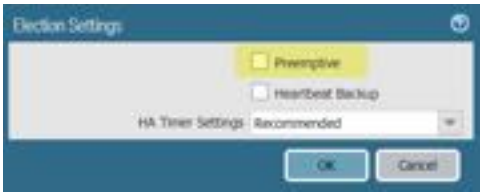


2. If the firewalls are not running the minimum required content release version or a later version required for PAN-OS 8.1, **Check Now** to retrieve a list of available updates.
3. Locate and **Download** the desired content release version.

   After you successfully download a content update file, the link in the Action column changes from **Download** to **Install** for that content release version.
4. **Install** the update. You must install the update on both peers.

STEP 4 | Disable preemption on the first peer in each pair. You only need to disable this setting on one firewall in the HA pair but ensure that the commit is successful before you proceed with the upgrade.

1. Select **Device** > **High Availability** and edit the **Election Settings**.
2. If enabled, disable (clear) the **Preemptive** setting and click **OK**.



3. **Commit** the change.

STEP 5 | Determine the Upgrade Path to PAN-OS 8.1

You cannot skip installation of any feature release versions in the path from the currently running PAN-OS version to PAN-OS 8.1.0.

Review the known issues and changes to default behavior in the Release Notes and upgrade/downgrade considerations in the New Features Guide for each release through which you pass as part of your upgrade path.

STEP 6 | Install PAN-OS 8.1 on the first peer.

To minimize downtime in an active/passive configuration, upgrade the passive peer first. For an active/active configuration, upgrade the secondary peer first. As a best practice, if you are using an active/active configuration, we recommend upgrading both peers during the same maintenance window.

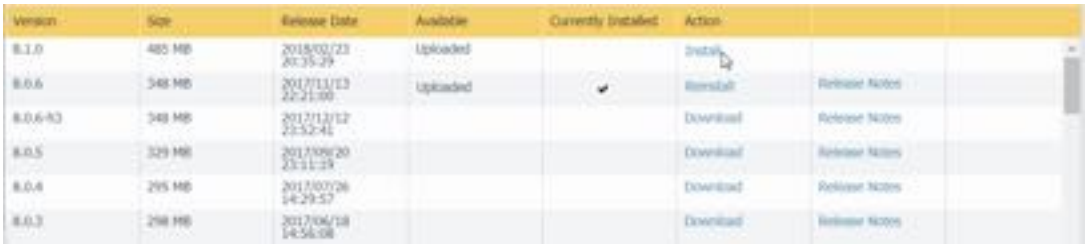
© 2020 Palo Alto Networks, Inc.

Because the User-ID mappings are managed differently in 8.1.0, after the passive peer is upgraded and available, the mappings are not synced until both peers are running the same release version.

If you want to test that HA is functioning properly before the upgrade, consider upgrading the active peer in an active/passive configuration first to ensure that failover occurs without incident.

1. On the first peer, select **Device** > **Software** and click **Check Now** for the latest updates.
2. Locate and **Download** PAN-OS 8.1.0.

   If your firewall does not have internet access from the management port, you can download the software image from the PaloAltoNetworksSupport Portal *and then manually Upload it to your firewall.*
3. After you download the image (or, for a manual upgrade, after you upload the image), **Install** the image.



4. After the installation completes successfully, reboot using one of the following methods:
   - If you are prompted to reboot, click **Yes**.
   - If you are not prompted to reboot, select **Device** > **Setup** > **Operations** and **Reboot Device**.
5. After the device finishes rebooting, view the High Availability widget on the **Dashboard** and verify that the device you just upgraded is still the passive or active-secondary peer in the HA configuration.



STEP 7 | If you have enabled User-ID, the firewall connects to the User-ID sources to repopulate the mappings. To confirm the firewall has successfully repopulated the mappings, use the following CLI commands:

- **`show user ip-user-mapping all`**
- **`show user group list`**

After the firewall repopulates the mappings, it can enforce user-based policy and user-based visibility is available.

STEP 8 | Install PAN-OS 8.1 on the second peer.

1. (Active/passive configurations only) Suspend the active peer so that HA fails over to the peer you just upgraded.
   1. On the active peer, select **Device** > **High Availability** > **Operational Commands** and click **Suspend local device**.

© 2020 Palo Alto Networks, Inc.



2. View the High Availability widget on the **Dashboard** and verify that the state changes to **Passive**.
3. On the other peer, verify that it is active and is passing traffic (**Monitor** > **Session Browser**).

2. On the second peer, select **Device** > **Software** and click **Check Now** for the latest updates.
3. Locate and **Download** PAN-OS 8.1.0.
4. After you download the image, **Install** it.
5. After the installation completes successfully, reboot using one of the following methods:
   - If you are prompted to reboot, click **Yes**.
   - If you are not prompted to reboot, select **Device** > **Setup** > **Operations** and **Reboot Device**.
6. (Active/passive configurations only) From the CLI of the peer you just upgraded, run the following command to make the firewall functional again:

   ```
   request high-availability state functional
   ```

STEP 9 | Verify that the other peer has successfully received the mappings using the CLI commands from Step 7.

STEP 10 | Verify that both peers are passing traffic as expected.

In an active/passive configuration, only the active peer should be passing traffic; both peers should be passing traffic in an active/active configuration.

Run the following CLI commands to confirm that the upgrade succeeded:

- (Active peers only) To verify that active peers are passing traffic, run the **`show session all`** command.
- To verify session synchronization, run the **`show high-availability interface ha2`** command and make sure that the Hardware Interface counters on the CPU table are increasing as follows:
  - In an active/passive configuration, only the active peer shows packets transmitted; the passive peer will show only packets received.

    If you enabled HA2 keep-alive, the hardware interface counters on the passive peer will show both transmit and receive packets. This occurs because HA2 keep-alive is bi-directional, which means that both peers transmit HA2 keep-alive packets.
  - In an active/active configuration, you will see packets received and packets transmitted on both peers.

STEP 11 | If you disabled preemption prior to the upgrade, re-enable it now.

1. Select **Device** > **High Availability** and edit the **Election Settings**.
2. Select **Preemptive** and click **OK**.
3. **Commit** the change.

© 2020 Palo Alto Networks, Inc.

Downgrade from PAN-OS 8.1

The way you downgrade a firewall from PAN-OS 8.1 depends on whether you are downgrading to a previous feature release (where the first or second digit in the PAN-OS version changes, for example, from 8.0.2 to 7.1.7 or from 7.1.6 to 7.0.9) or downgrading to a maintenance release version within the same feature release (where the third digit in the release version changes, for example, from 8.0.2 to 8.0.0). When you downgrade from one feature release to an earlier feature release, you can migrate the configuration from the later release to accommodate new features. To migrate the PAN-OS 8.1 configuration to an earlier PAN-OS release, first restore the configuration for the feature release to which you are downgrading. You do not need to restore the configuration when you downgrade from one maintenance release to another within the same feature release.

- Downgrade a Firewall to a Previous Maintenance Release
- Downgrade a Firewall to a Previous Feature Release
- Downgrade a Windows Agent from PAN-OS 8.1



Always downgrade into a configuration that matches the software version. Unmatched software versions and configurations can result in failed downgrades or force the system into maintenance mode. This only applies to a downgrade from one feature release to another (for example 8.1.0 to 8.0.7), not to downgrades to maintenance releases within the same feature release version (for example, 8.0.7 to 8.0.4).

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.

Downgrade a Firewall to a Previous Feature Release

Use the following workflow to restore the configuration that was running before you upgraded to a different feature release. Any changes made since the upgrade are lost. Therefore, it is important to back up your current configuration so you can restore those changes when you return to the newer feature release.

Use the following procedure to downgrade to a previous feature release.

STEP 1 | Save a backup of the current configuration file.



Although the firewall automatically creates a backup of the configuration, it is a best practice to create a backup before you upgrade and store it externally.

1. **Export named configuration snapshot** (**Device** > **Setup** > **Operations**).
2. Select the XML file that contains your running configuration (for example, **running-config.xml**) and click **OK** to export the configuration file.
3. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the downgrade.

STEP 2 | Install the previous feature release image.



Autosave versions are created when you upgrade to a new release.

1. **Check Now** (**Device** > **Software**) for available images.
2. Locate the image to which you want to downgrade. If the image is not already downloaded, then **Download** it.

© 2020 Palo Alto Networks, Inc.

3. After the download completes, **Install** the image.
4. **Select a Config File for Downgrading**, which the firewall will load after you reboot the device. In most cases, you should select the configuration that was saved automatically when you upgraded from the release to which you are now downgrading. For example, if you are running PAN-OS 8.1 and are downgrading to PAN-OS 8.0.7, select `autosave-8.0.7`.
5. After the installation completes successfully, reboot using one of the following methods:
   - If you are prompted to reboot, click **Yes**.
   - If you are not prompted to reboot, go to Device Operations (**Device** > **Setup** > **Operations**) and **Reboot Device**.

Downgrade a Firewall to a Previous Maintenance Release

Because maintenance releases do not introduce new features, you can downgrade to a previous maintenance release in the same feature release without having to restore the previous configuration. A maintenance release is a release in which the third digit in the release version changes, for example a downgrade from 7.1.7 to 7.1.2 is considered a maintenance release downgrade because only the third digit in the release version is different.

Use the following procedure to downgrade to a previous maintenance release within the same feature release.

STEP 1 | Save a backup of the current configuration file.



Although the firewall automatically creates a backup of the configuration, it is a best practice to create a backup before you downgrade and store it externally.

1. **Export named configuration snapshot** (**Device** > **Setup** > **Operations**).
2. Select the XML file that contains your running configuration (for example, **running-config.xml**) and click **OK** to export the configuration file.
3. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the downgrade.

STEP 2 | Install the previous maintenance release image.



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 it to your firewall.*

1. **Check Now** (**Device** > **Software**) for available images.
2. Locate the version to which you want to downgrade. If the image is not already downloaded, then **Download** it.
3. After the download completes, **Install** the image.
4. After the installation completes successfully, reboot using one of the following methods:
   - If you are prompted to reboot, click **Yes**.
   - If you are not prompted to reboot, go to Device Operations (**Device** > **Setup** > **Operations**) and **Reboot Device**.

Downgrade a Windows Agent from PAN-OS 8.1

After you uninstall the PAN-OS 8.1 Windows-based User-ID agent, perform the following steps before you install an earlier agent release.

© 2020 Palo Alto Networks, Inc.

STEP 1 | Open the Windows Start menu and select **Administrative Tools**.

STEP 2 | Select **Computer Management** > **Services and Applications** > **Services** and double-click **User-ID Agent**.

STEP 3 | Select **Log On**, select **This account**, and specify the username for the User-ID agent account.

STEP 4 | Enter the **Password** and **Confirm Password**.

STEP 5 | Click **OK** to save your changes.

© 2020 Palo Alto Networks, Inc.

App-ID Features

- SaaS Application Hosting Characteristics
- Simplified App-ID
- HTTP Header Insertion and Modification
- Service-Based Session Timeouts

© 2020 Palo Alto Networks, Inc.

SaaS Application Hosting Characteristics

You can now identify, assess, and reduce risks related to SaaS applications with unfavorable hosting characteristics by leveraging the enhanced ACC filters in App-ID™. To help you determine which applications you should allow in your environment, five new application hosting characteristics are now available: data breaches, poor terms of service, no certifications, poor financial viability, and IP-based access restrictions. You can use these application characteristics to gain visibility, control and the information you need to make informed decisions about the apps on your network, and define actions to eliminate risks or prevent future violations.

The enhanced ACC filters allow you to view the detailed risk profiles and usage statistics for the SaaS applications on your network based on the following hosting characteristics:

SaaS Application Characteristic	Description
Data Breaches	It is important for you to know if the SaaS application has ever had an intentional or unintentional release of secure information to an untrusted environment without proper information security precautions. Data loss and exfiltration is a possible risk when you use SaaS applications that do not disclose breaches or that have been breached within the past three years.
Poor Terms of Service	Applications with poor terms of service such as termination of service at any time for any reason, or making your sensitive data public for non-payment can make your enterprise vulnerable to loss of intellectual property, data, or nonavailability of service.
No Certifications	Certifications can help you assess applications to verify if the application has gone through certain basic security precautions that have been validated and officially recorded by a certification party. Most certifications are not a one-time evaluation and provide you the assurance that the security of the application is periodically evaluated. Applications are checked for compliance with industry certifications such as SOC1, SOC2, SSAE16, PCI, HIPAA, FINRAA, or FEDRAMP.
Poor Financial Viability	If the application vendor goes out of business, the liquidation activities can have a serious impact your enterprise. For example, during liquidation the customer data is considered an asset and sold, potentially exposing your IP and losing the data to the liquidator. This characteristic checks for applications with the potential to be out of business within the next 18 to 24 months and reports the applications with poor financial viability.
IP-based Restrictions	An application without IP-based restrictions lack the ability to restrict end-user access from outside of the corporate network, and can expose your enterprise data to end-users outside of your known enterprise IP range (such as a corporate campus).

© 2020 Palo Alto Networks, Inc.

SaaS Application Characteristic	Description
	The data for the IP-based Restrictions characteristic is available to view in the Network Activity graphs and for generating a custom report.

For ease of use, the applications with unfavorable hosting characteristics are displayed automatically in the predefined Application Usage Report.

The application usage report displays in a new window, so your browser must allow pop-ups. If the preview window does not open, refer to your browser documentation for the steps to enable pop-ups.

The Applications with Risky Characteristics displays on page 2 of the report.



For a more tailored view, you can also use the characteristics to build a custom report.

© 2020 Palo Alto Networks, Inc.

Simplified App-ID

App-ID identifies the applications traversing your network—regardless of port, protocol, encryption, or any evasive tactics—so that you can safely enable desired applications and block unwanted applications. Palo Alto Networks releases new App-IDs on a monthly basis that your security policy can begin to enforce without any additional configuration. While this enables the firewall to dynamically control application traffic with ever-increasing precision, changes in how the firewall enforces applications can potentially impact availability for the mission-critical applications on which your organization relies.

Now, you no longer need to weigh availability for mission-critical applications against equipping the firewall with the latest application knowledge; with the following options, you can do both, and it'll be easier for you to move to and maintain an application-based security policy:

- New App-ID Threshold— Fine tune content update installation thresholds so that threat updates install immediately as they're released, and new App-IDs are installed only after you've had an opportunity to make any necessary security policy updates.
- New App-ID Characteristic—Automatically allow new App-IDs in critical categories (like software-development), and get visibility into new App-ID activity on your network.
- Extended Policy Impact Review for Content Releases—Now you can see how currently installed content releases impact security policy enforcement, and policy review is extended to include modified App-IDs. (Previously, you could only perform a policy impact review for new App-IDs before you installed a content release).
- Coverage Change Details for Modified App-IDs—Get details on how coverage for modified applications is now expanded or more precise.

New App-ID Threshold

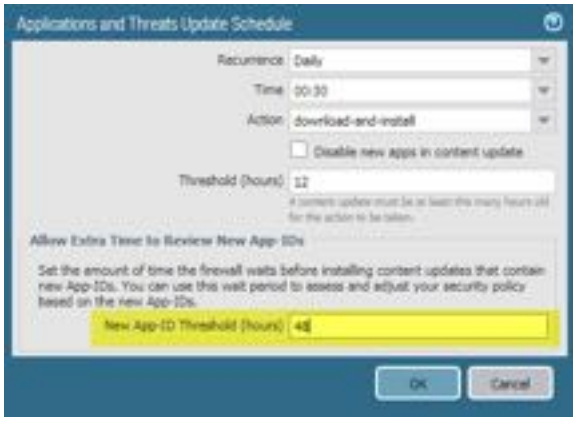
New and updated application and threat signatures are released together in a single package as Applications and Threats content updates. Previously, this meant that you had to install threat updates and new and modified applications at the same time, and choose one of two approaches to deploying content updates:

- A mission-critical approach, where you delay content update installation until you can assess it's impact to application availability.
- A security-first approach, where you install application and threat updates as they're made available; you prioritize the latest threat protection over a possible impact to application availability.

Now, the New App-ID Threshold allows you to install content updates that include new App-IDs on a separate schedule from those that don't. Especially if your business has a mission-critical approach, this gives you extra time to review how new App-IDs impact security policy enforcement, and make any necessary policy updates.

To configure a New App-ID Threshold, modify the Applications and Threats Update Schedule (select **Device** > **Dynamic Updates** > **Schedule**).

© 2020 Palo Alto Networks, Inc.



The New App-ID Threshold must always be a value greater than the overall content update threshold. The example above shows what an Application and Threats Update Schedule might look like if you decide on a mission-critical approach to deploying content updates. Note that the overall content Threshold is set to 12 hours, while the new App-ID threshold is set to 48 hours.

New App-ID Characteristic

The New App-ID characteristic is now found throughout the firewall web interface; in several different contexts, it enables you to filter for only the App-IDs that were introduced in the most recently installed content releases. This means that you can enforce brand new App-IDs without adjusting your security policy, and monitor new App-ID activity.

- Ensure Critical New App-IDs are Allowed—Use the New App-ID characteristic in policy to allow critical applications based on category (for example, always allow authentication or software development apps). Ensure availability for these critical applications, while continuing to get the latest App-ID updates.
- Get Increased Visibility Into New App-ID Activity—Monitor and generate reports for new App-IDs. What you learn can help you make the right decisions about how you to update your security policy to enforce the most recently-categorized App-IDs.

Regardless of where you use it, the New App-ID characteristic always matches to only the new App-IDs in the most recently installed content releases. When a new content release is installed, the new App-ID characteristic automatically begins to enforce only the new App-IDs in that content release version.

Ensure Critical New App-IDs are Allowed

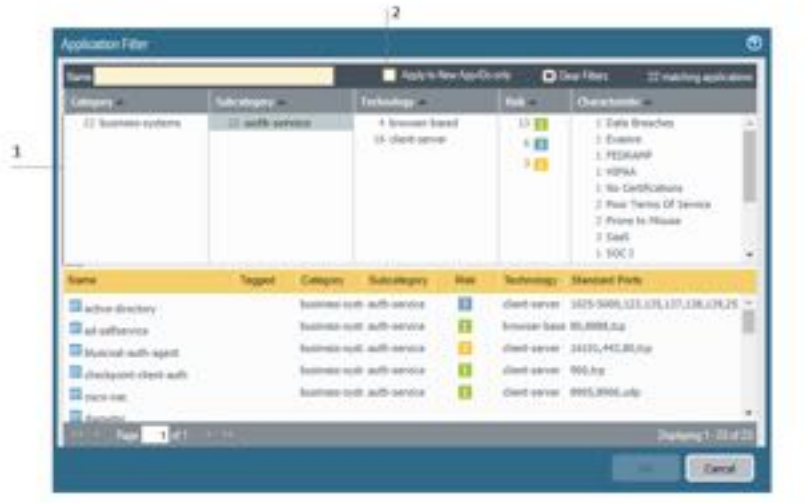
New App-IDs can cause a change in policy enforcement for traffic that is newly-identified as belonging to a certain application. You can now add the predefined **New App-ID** application filter to a security policy rule in order to use only new App-IDs as match criteria for the rule. You can choose to enforce all new App-IDs, or target the security policy rule to enforce certain types of new App-IDs (for example, enforce only authentication or software development applications). Set the security policy rule to **Allow** to ensure that even if an App-ID release introduces expanded or more precise coverage for critical applications, the firewall continues to allow them. New App-IDs are released on the third Tuesday of every month, and the application filter provides coverage for only those applications in the latest release. This gives you a month's time (or, if the firewall is not installing content updates on a schedule, until the next time you manually install content) to assess how newly-categorized applications might impact security policy enforcement and make any necessary adjustments.

STEP 1 | Select **Objects** > **Application Filters** and **Add** a new application filter.

© 2020 Palo Alto Networks, Inc.

STEP 2 | Define the types of new applications for which you want to ensure constant availability based on subcategory or characteristic. For example, select the category “auth-service” to ensure that any newly-installed applications that are known to perform or support authentication are allowed.

STEP 3 | Only after narrowing the types of new applications that you want to allow immediately upon installation, select **Apply to New App-IDs only**.



STEP 4 | Select **Policies** > **Security** and add or edit a security policy rule that is configured to allow matching traffic.

STEP 5 | Select **Application** and add the new **Application Filter** to the policy rule as match criteria.

STEP 6 | Click **OK** and **Commit** to save your changes.

STEP 7 | To continue to adjust your security policy to account for any changes to enforcement that new App-IDs introduce:

- Get Increased Visibility Into New App-ID Activity—Monitor and get reports on new App-ID activity.
- Extended Policy Impact Review for Content Releases—See how the newly-installed App-IDs impact your existing security policy rules.

Get Increased Visibility Into New App-ID Activity

The New App-ID characteristic enables you to monitor new applications on your network, so that you can better assess the security policy updates you might want to make. Use the New App-ID characteristic on the ACC to get visibility into the new applications on your network, and to generate reports that detail newly-categorized application activity.

- Generate a report with details specifically regarding new applications (applications introduced only in the latest content release).

© 2020 Palo Alto Networks, Inc.



- Use the ACC to monitor new application activity: select **ACC**and under **Global Filters**, select **Application** > **Application Characteristics** > **New App-ID**.



Extended Policy Impact Review for Content Releases

Policy impact review is now available for installed content releases and includes details on how modified App-IDs affect security policy enforcement. Previously, policy impact review details, where you can view all security policy rules where application enforcement might have changed due to a content update, were only available for downloaded content updates and detailed only how new App-IDs affected security policy rules.

Take the following steps to see a list of the applications modified in a content release and assess how those changes impact security policy enforcement. Importantly, updates to applications with network-wide impact (for example, LDAP or IKE) are prominently flagged as a **Policy Review Recommendation**.

STEP 1 | Select **Device** > **Dynamic Updates** and **Download** the latest Applications and Threats content release.

STEP 2 | For any content release (downloaded or currently installed) select either **Review Policies** or **Review Apps** in the Action column of the content update row.



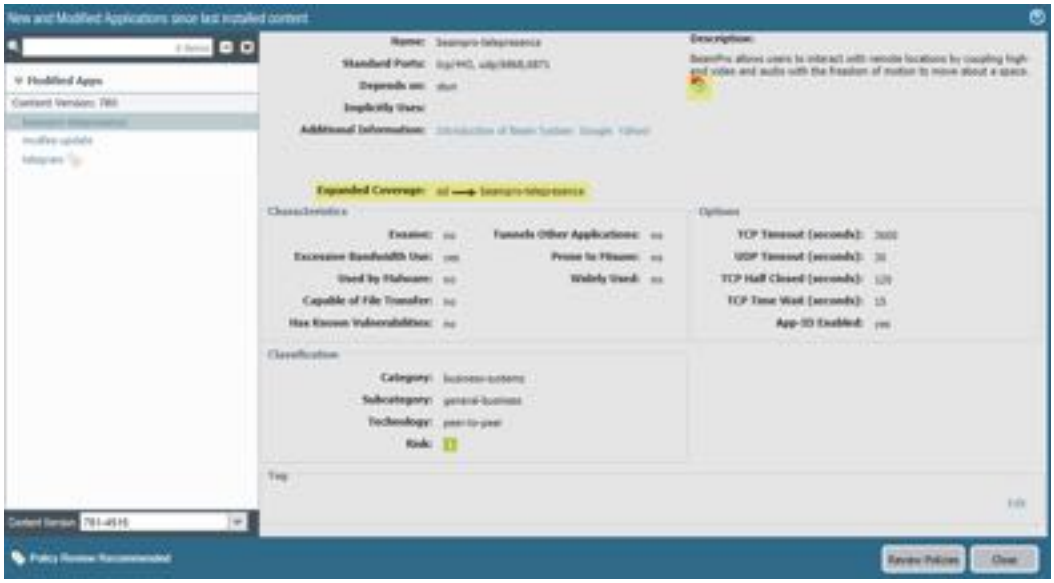
© 2020 Palo Alto Networks, Inc.

STEP 3 | **Review Apps** to see a list of all applications that are modified in a content release, and see details for each application. In addition to modified applications being listed, updates to applications with network-wide impact (for example, LDAP or IKE) are prominently flagged as a **Policy Review Recommendation**. This signals to you to evaluate potential changes in enforcement for those critical applications.

STEP 4 | **Review Policies** to see the security policy rules that might enforce traffic differently now that the application is modified.

Coverage Change Details for Modified App-IDs

Application details now include information on: **Expanded Coverage**, **Remove False Positive**, and application metadata changes. The **Expanded Coverage** and **Remove False Positive** fields both indicate how the application's coverage has changed (it's either more comprehensive or has been narrowed) and a clock icon indicates a metadata change, where certain application details are updated.



© 2020 Palo Alto Networks, Inc.

HTTP Header Insertion and Modification

Unsanctioned usage of SaaS applications can be a way for your users to transmit sensitive information outside of your network, usually by accessing a consumer version of an application. However, if you need to allow access to the enterprise version of these applications for specific individuals or organizations, then you can't block the SaaS application entirely. With the HTTP header insertion and modification feature, you can now manage HTTP header information to disallow SaaS consumer accounts while allowing a specific enterprise account.

Because many SaaS applications allow or disallow access to applications based on information contained in specific HTTP headers, you can use predefined header insertion rules to manage access to popular SaaS applications. Predefined rules are specific to particular SaaS applications. Currently there are four available predefined applications: Google, Dropbox, YouTube, and Office 365. These applications (and any additional applications that might be predefined in the future) are populated and maintained by Palo Alto Networks® using content updates. If a SaaS vendor changes its implementation so that additional header or domain information is needed, Palo Alto Networks will use a content update to modify the relevant predefined rule for you.

If you want to perform HTTP header insertion for an application that is not yet predefined in a content update, you can create a custom rule. Custom rules allow you to manage custom HTTP headers but you can also use them to manage standard HTTP headers depending on your requirements.

Custom rules that you create will never be managed or modified by Palo Alto Networks. It is your responsibility to maintain the custom rules that you create so that they accurately interact with the SaaS application.



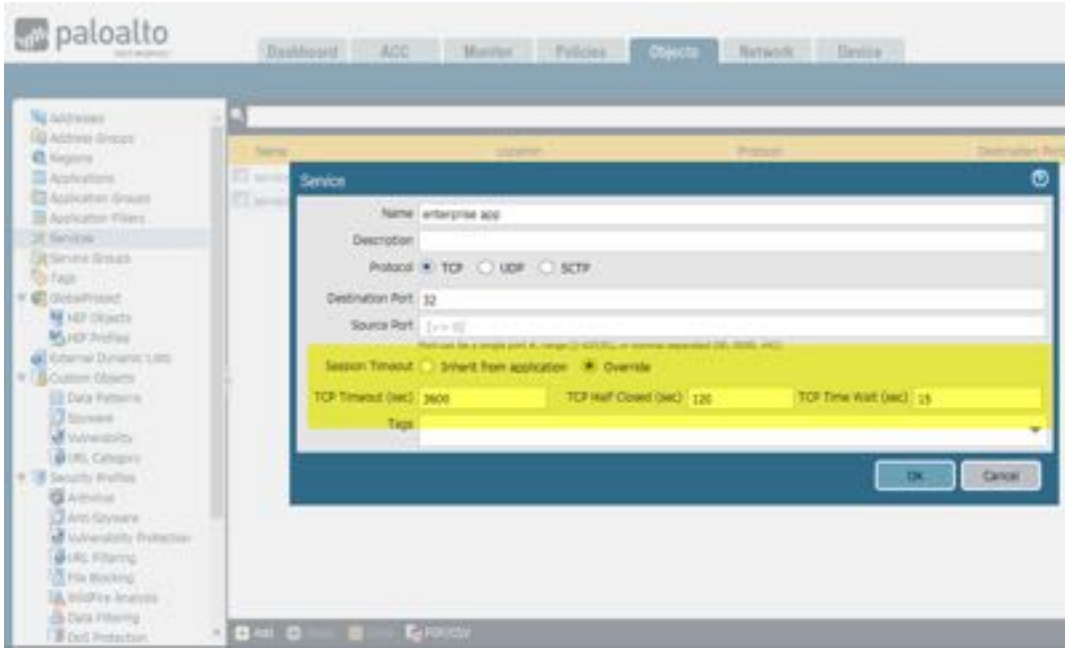
You don't need a URL Filtering license to use this feature.

© 2020 Palo Alto Networks, Inc.

Service-Based Session Timeouts

You can now more easily maintain custom timeouts for applications as you move from a port-based policy to an application-based policy. Previously, you might have had to override App-ID (losing application visibility) or create a custom App-ID (expending time and research) in order to maintain custom timeouts for applications.

To get started, configure custom timeout settings as part of a service object:



Then add the service object in a policy rule to apply the custom timeouts to the application(s) the rule enforces. See the full workflow to Maintain Custom Timeouts for Legacy Applications—you can follow these same steps to apply custom timeouts to user groups.

© 2020 Palo Alto Networks, Inc.

Virtualization Features

- VM-50 Lite
- Integration with Azure Security Center
- Support for Azure Application Insights
- Bootstrapping Enhancements for VM-Series on Azure
- VM-Series Firewall on Google Cloud Platform

© 2020 Palo Alto Networks, Inc.

VM-50 Lite

The VM-50 Lite provides an alternative for environments where hardware resources are constrained; specifically environments where the standard VM-50's 4.5 GB memory footprint is too large. The VM-50 Lite requires 4GB of memory instead of the 4.5GB required by the standard VM-50.

The VM-50 Lite is a mode of operation, not a new model. Enabling lite mode uses the same license as the standard VM-50 but changes to a VM-50 Lite when you lower the allocated memory. And while the VM-50 Lite has lower capacity, it has the same performance of the standard VM-50.

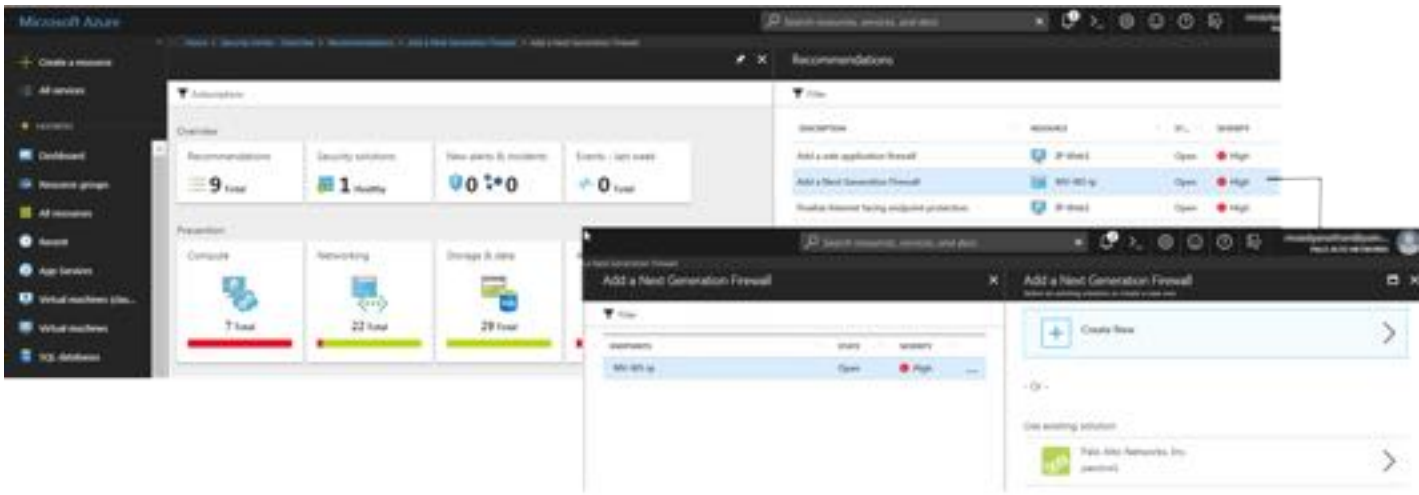


The VM-50 Lite does not support the Application Command Center (ACC).

© 2020 Palo Alto Networks, Inc.

Integration with Azure Security Center

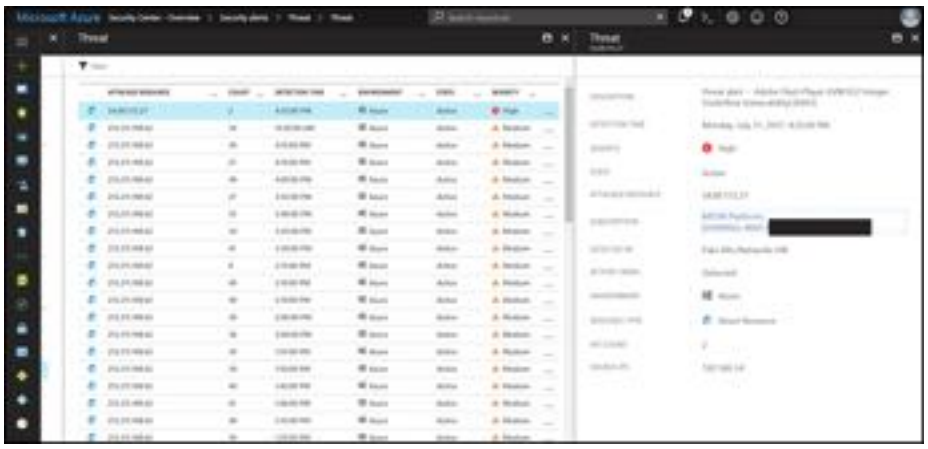
The VM-Series firewall integration with Azure Security Center provides a single pane of glass for high-priority security alerts so you can start triaging an incident directly from the Azure Security Center dashboard. To start using this integration, you must enable Azure Security Center on your Azure subscription.



When you deploy a VM-Series firewall on Azure directly from Azure Security Center, the firewall is automatically configured with two example Security policy rules to safely inspect and allow inbound web-browsing traffic and outbound traffic, and it includes a log forwarding rule to send security-related logs to Azure. With this log forwarding profile, Threat and WildFire Submissions logs of low, medium, high, or critical severity generated on the firewall are displayed as security alerts on the Azure Security Center dashboard.

Currently Azure restricts you from deploying a multi NIC appliance in an existing resource group. Therefore, you cannot deploy the VM-Series firewall in a resource group where you have deployed the workloads you want to secure. To work around this limitation and make practical use of the default configuration for the VM-Series firewall, you can stage a security risk and deploy a workload with a public IP address that is exposed to the internet. Doing so will trigger the Azure Security Center recommendation for a next-generation firewall, and you can use this recommendation to deploy the firewall in an empty resource group. The deployment workflow is the same as the Azure Marketplace for VM-Series firewall. After you deploy the VM-Series firewall, delete the internet exposed workload, you can deploy your applications or workloads in the resource group when needed within the resource group where you've already deployed the firewall.

© 2020 Palo Alto Networks, Inc.



Azure Security Center can also automatically discover an existing or new VM-Series firewall instance that you launch with PAN-OS 8.1 from the Azure marketplace or have a custom deployment using the Azure CLI, PowerShell or ARM template. To enable the discovery of the VM-Series firewall as a Security Solution on the Azure Security Center dashboard, you must have the Standard tier of Azure Security Center enabled on your subscription. The workflow to forward security-related logs from the VM-Series firewall to Azure Security Center is more involved and requires you to use an intermediate Linux virtual machine and configure Syslog forwarding to send the required logs from the VM-Series firewall. If you have already configured your firewall, you can manually attach the Azure Security Center- default Log Forwarding profile directly on the firewall or use Panorama templates and device groups to enable managed firewalls to forward logs to Azure Security Center.

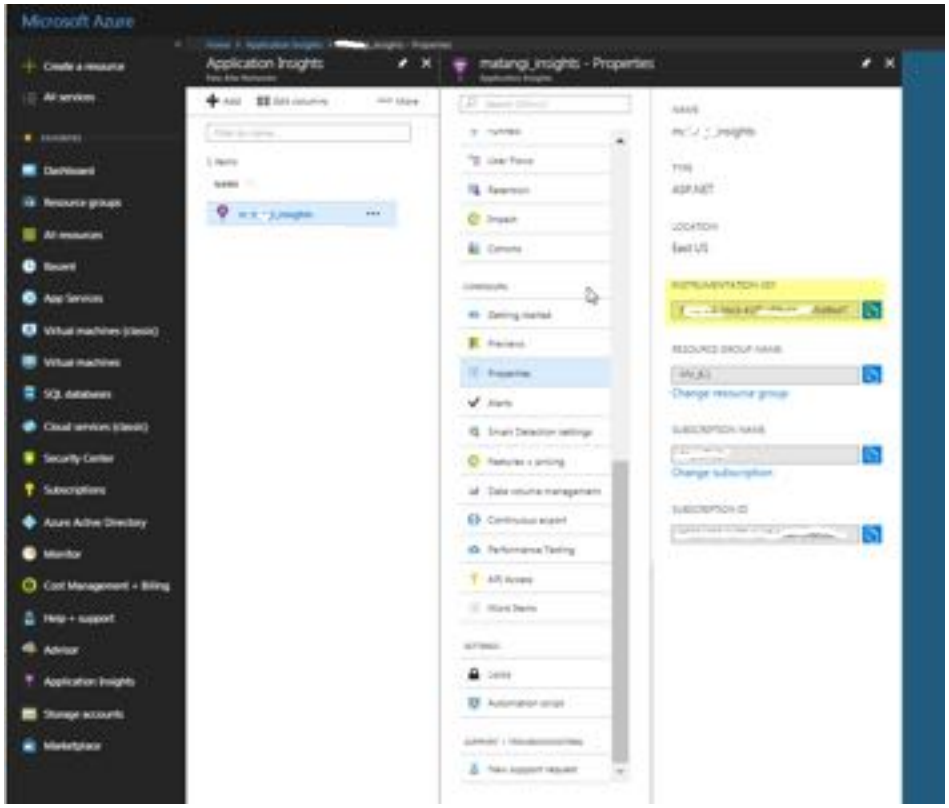
© 2020 Palo Alto Networks, Inc.

Support for Azure Application Insights

In PAN-OS 8.1, the VM-Series firewall on Azure includes support for natively publishing PAN-OS metrics to an Azure Application Insights instance where you can monitor firewall performance and health status, view trends on performance and usage patterns, and set up alerts to detect anomalies proactively.

STEP 1 | Set up an Application Insights instance of type **General**. Refer to Azure documentation for details on how to create an Application Insights resource.

STEP 2 | Log in to the VM-Series firewall on Azure and configure it to publish PAN-OS metrics to Azure Application Insights.



STEP 3 | Select **Device** > **Setup** > **Operations**.

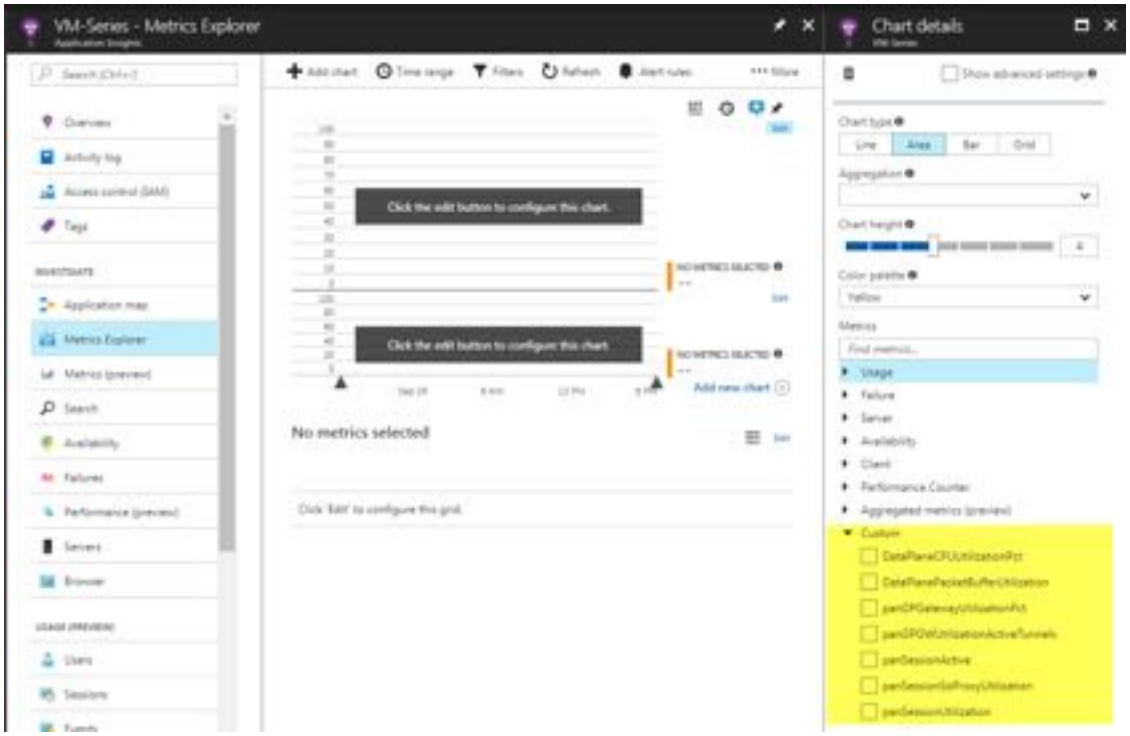
STEP 4 | Select **Enable Application Insights,** and copy and paste the **Instrumentation Key** from the Azure portal. The firewall requires this key to authenticate to the Application Insights instance and publish metrics to it.

STEP 5 | **Commit** the changes.

STEP 6 | Verify that the firewall is publishing metrics.

Log in to the Azure portal and on the Application Insights instance that you entered above, view the custom metrics on the Metrics Explorer.

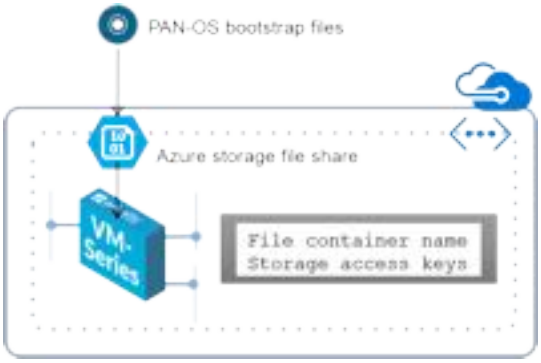
© 2020 Palo Alto Networks, Inc.



© 2020 Palo Alto Networks, Inc.

Bootstrapping Enhancements for VM-Series on Azure

The solution template in the Azure Marketplace has been updated to enable bootstrapping, and you can now use Azure file storage to upload the configuration files and licenses to easily bootstrap the VM-Series firewall at launch. In order to bootstrap successfully, you must provide the name of the storage account that holds the bootstrap folders and the keys to authenticate for access the storage account.



STEP 1 | Log in to the Azure portal.

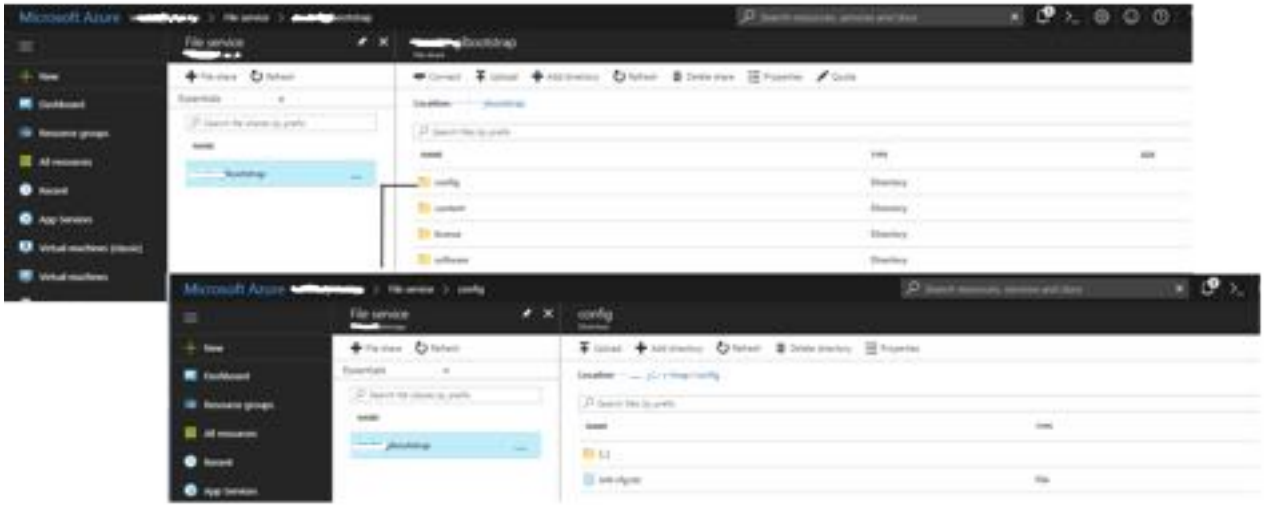
For example, **`https://portal.azure.com`** or **`https://portal.azure.us`**.

STEP 2 | Set up the storage account with a File service. The File service allows the VM-Series firewalls to read files stored there and get provisioning information at launch.

If you are using the solution template in the Azure marketplace to deploy the VM-Series, make sure to keep this storage account in a different resource group that the one in which you plan to deploy the firewall. This is because the solution template can be deployed in either a new or empty resource group.

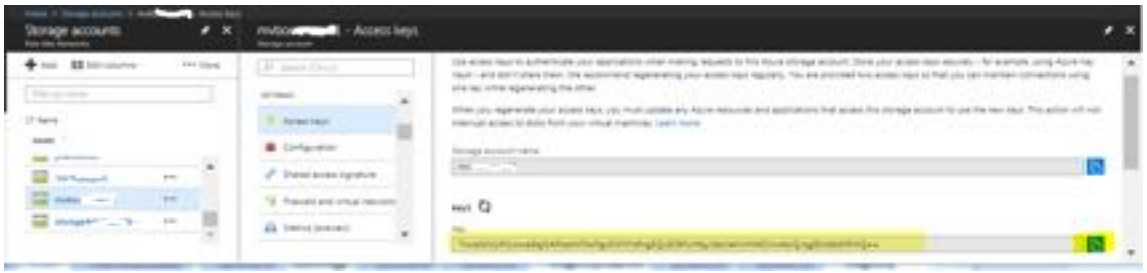
Within the File service, you must create a **File share**. In the File share you must **Add** the directory structure for the bootstrap package so that you can upload the files required for bootstrapping. Optionally, you can specify a **Share-directory**. If you have a common File share that serves as a repository for bootstrap configuration for different set ups, using a share-directory gives you the flexibility to create a folder hierarchy and access a specific set of sub-folders within the common File share.

© 2020 Palo Alto Networks, Inc.



STEP 3 | Add the information to bootstrap the firewall, when you Deploy the VM-Series firewall.

1. Select Enable Bootstrap **yes**.
2. Enter the **Storage Account Name** that holds the bootstrap package.
3. Enter the **Storage Account Access Key**. This firewall needs this access key to authenticate to the storage account and access the files stored in it.

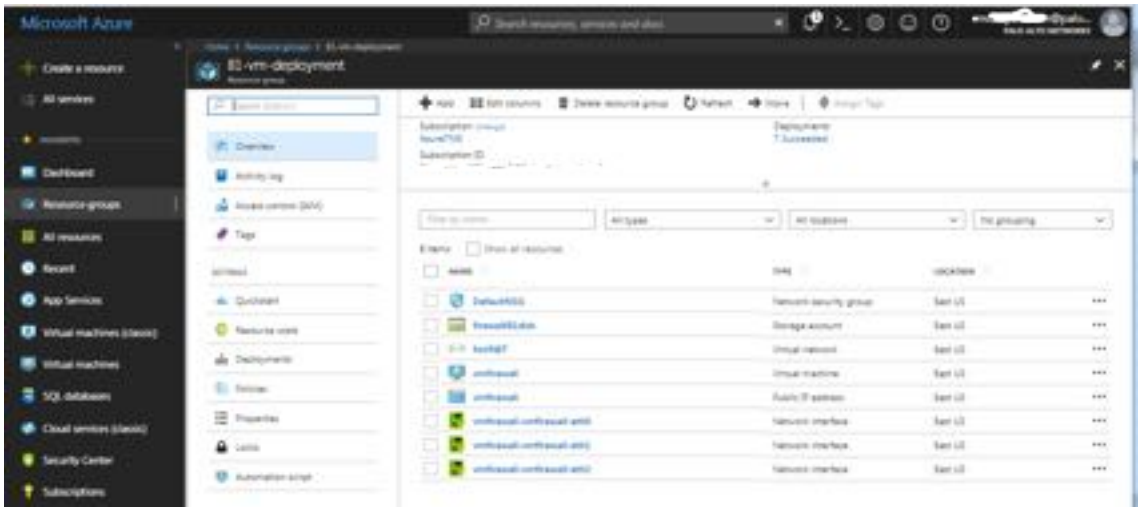


4. Add the **File share name** to which you have uploaded the files required for bootstrapping the firewall. The storage account must be set up with the correct folder structure for the bootstrapping.

STEP 4 | Verify that you have successfully bootstrapped the VM-Series firewall.

1. Select **Dashboard** > **Resource Groups**, select the resource group.
2. Select **Settings** > **Deployments** > **Deployment History** for detailed status.

© 2020 Palo Alto Networks, Inc.



STEP 5 | To publish PAN-OS metrics to Azure Application Insights, see Support for Azure Application Insights.

© 2020 Palo Alto Networks, Inc.

VM-Series Firewall on Google Cloud Platform

You can now deploy the VM-Series firewall on a Google Compute Engine instance within a Google Cloud Platform project to secure your applications and workloads.

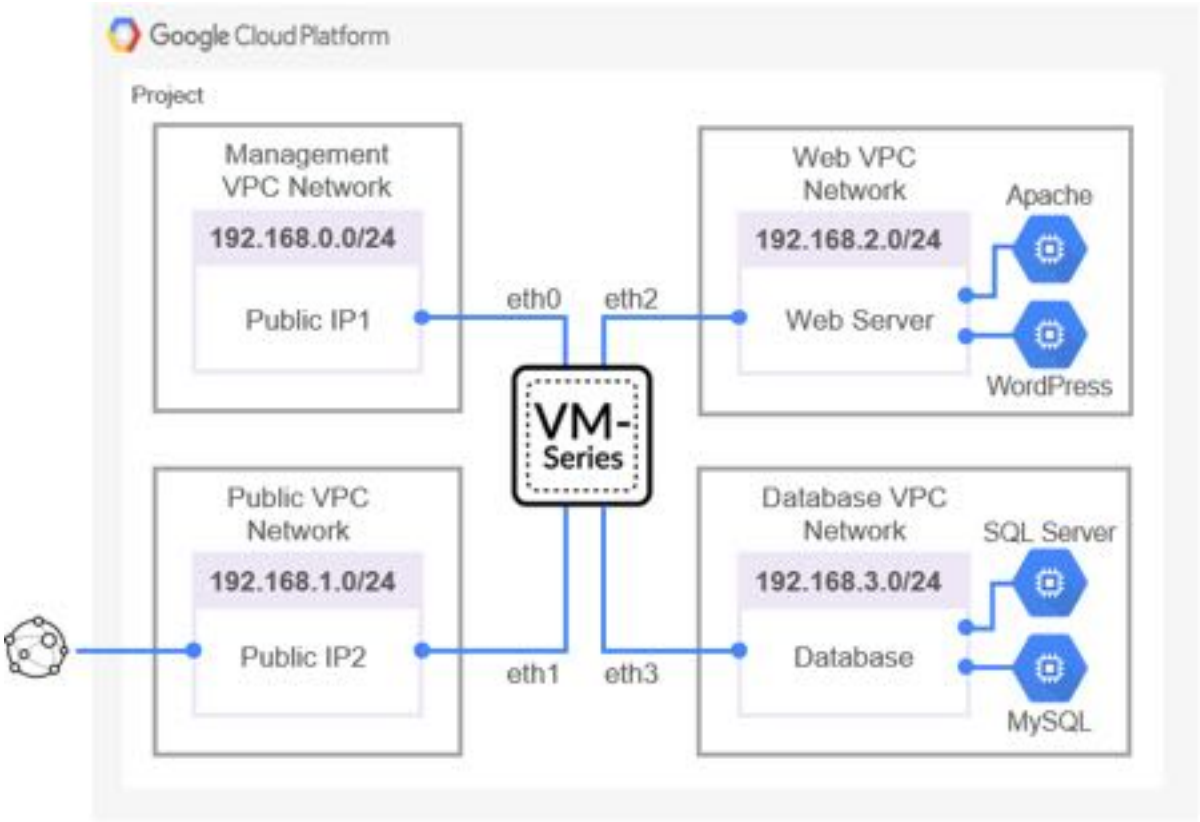
The VM-Series firewall on Google Cloud Platform can publish custom PAN-OS metrics to Google Stackdriver. With Stackdriver Monitoring, you can monitor the firewall, and set up alerts based on firewall health and performance.

You can also enable any firewall that runs PAN-OS 8.1 (virtual or physical) to monitor application workloads deployed on Google Compute Engine instances. With an awareness of virtual machine adds, moves, or deletes within a Google VPC, you can create security policy rules that automatically adapt to changes in your application environment.

- Deploy the VM-Series Firewall from Google Cloud Platform Marketplace
- Enable Google Stackdriver Monitoring
- Enable VM-Series Firewall to Track Changes on Google Cloud Platform VMs

Deploy the VM-Series Firewall from Google Cloud Platform Marketplace

The VM-series firewall is part of your Google project, using the VPC networks to communicate with other compute engine instances. In addition to serving as an internet gateway, the VM-series firewall can secure east-west traffic between VPCs to ensure data protection compliance and application access.



© 2020 Palo Alto Networks, Inc.

Google Marketplace provides templates based on license types. The templates deploy an instance of the VM-Series firewall with a management interface and two dataplane interfaces.

Before you deploy the VM-Series firewall, you must choose a project in your organization, and create a minimum of three networks and subnetworks that the firewall requires at launch.

STEP 1 | Locate the VM-Series firewall listing in Google Marketplace.

1. Log in to the Google Cloud Console.
2. From the Products and Services menu, choose **Marketplace**.
3. Search for "VM-Series".
4. Select one of the VM-Series licensing options.

STEP 2 | Click **Launch on Compute Engine**.

STEP 3 | Deploy the VM-Series Firewall from Google Cloud Platform Marketplace.

Enable Google Stackdriver Monitoring

You can enable any firewall that runs PAN-OS 8.1 (virtual or physical) to monitor application workloads deployed on Google Compute Engine instances. With an awareness of virtual machine adds, moves, or deletes within a Google VPC, you can create security policy rules that automatically adapt to changes in your application environment.

When Stackdriver is enabled, the firewall can retrieve metadata on eight predefined attributes—hostname of the VM, machine type, status (running or not), source (OS type), VPC Network, subnetwork, zone, and Project ID. In addition to these attributes, you can retrieve up to 24 user-defined attributes such as labels, tags, and other key-value pairs defined using metadata or startup scripts on the Google Compute Engine instances.

STEP 1 | Select **Device** > **VM Information Sources**, and **Add** a new source to monitor.

STEP 2 | Enter a unique **Name** for the source.

STEP 3 | Select the **Service Authentication Type**.

- **VM-Series running in GCE**—Use this option if a VM-Series firewall deployed on GCE is monitoring the virtual machines on GCE. You do not need to provide account credentials if the service account that you used to provision the firewall has the permissions required to authenticate to the Google Cloud Project you want to monitor.
- **Service Account**—Use this option on any hardware-based firewall or VM-Series firewall that is not running on GCE. You must the provide the **Service Account Credential** as a JSON file so that the firewall can authenticate to the GCP infrastructure and retrieve the attributes.

STEP 4 | Enter the **Project ID** and the **Zone** in which the resources are deployed.

STEP 5 | Click **OK** and **Commit** your changes.

STEP 6 | Verify the connection **Status** is successful and that the firewall is able to connect to the GCE project you want to monitor.

STEP 7 | Use the attributes as match criteria in dynamic address groups.

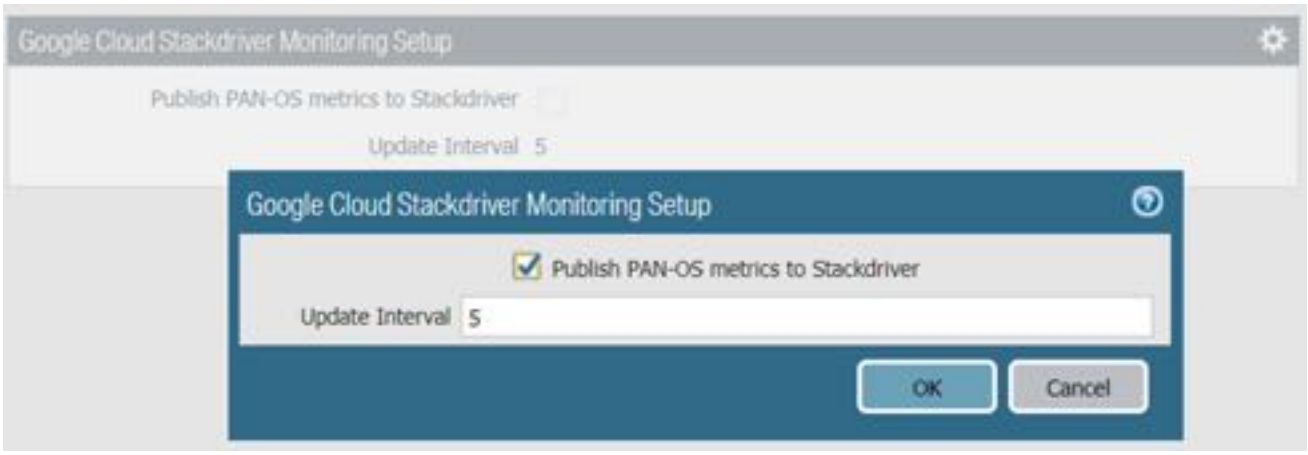
© 2020 Palo Alto Networks, Inc.

Enable VM-Series Firewall to Track Changes on Google Cloud Platform VMs

You can also enable any firewall that runs PAN-OS 8.1(virtual or physical) to monitor application workloads deployed on Google Compute Engine instances. For a description of the PAN-OS metrics that you can publish to Google Stackdriver, see Custom PAN-OS Metrics Published for Monitoring.

STEP 1 | Push PAN-OS metrics from a VM-Series firewall on a Google Compute Engine instance to Stackdriver.

1. Log in to the web interface on the VM-Series firewall.
2. Select **Device** > **Operations**. On the Google Cloud Stackdriver Monitoring Setup panel, click Edit .
    1. Check **Publish PAN-OS metrics to Stackdriver**.



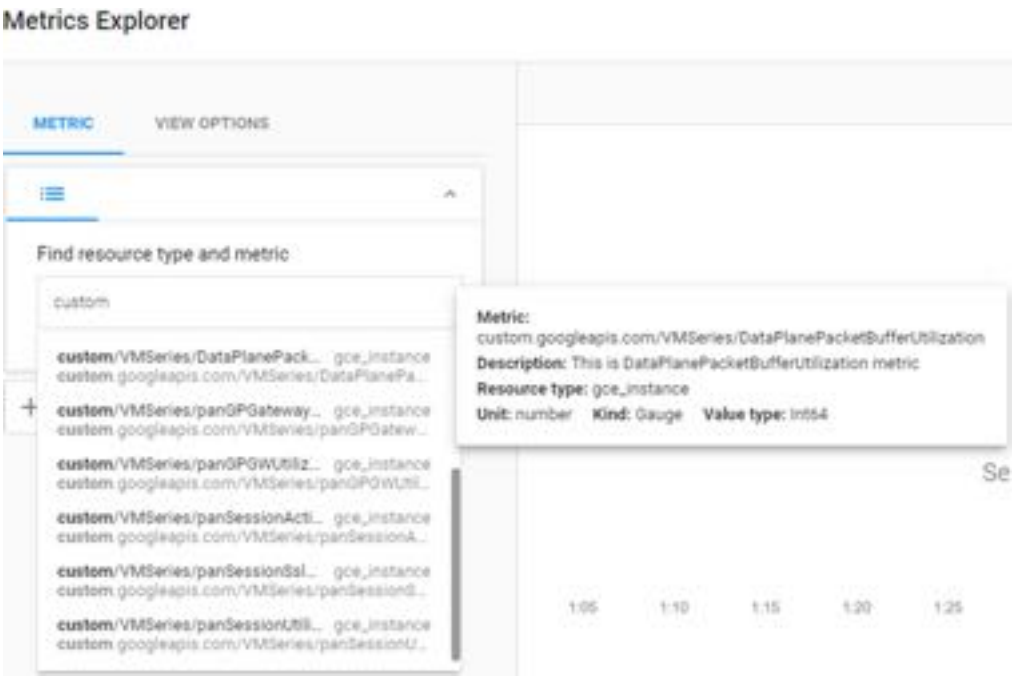
    2. Set **Update Interval** to a value between 1- 60 minutes. This is the frequency at which the firewall publishes the metrics to Stackdriver. The default is 5 minutes.
    3. Click **OK**.
3. **Commit** the changes.

    Wait until the firewall starts to publish metrics to Stackdriver before you configure alarms for PAN-OS metrics.

STEP 2 | Verify that you can see the metrics on Stackdriver.

1. In the Google Cloud Console, select **Products and Services** > **Monitoring**.
2. In Stackdriver, choose **Resources** > **Metrics Explorer**.
3. Under “Find resource type and metric”, click in the search field and type **`custom`** to filter the PAN-OS metrics.

© 2020 Palo Alto Networks, Inc.



STEP 3 | Configure alerts and actions for PAN-OS metrics on Stackdriver. See Monitoring Quickstart for Google Compute Engine, and Stackdriver Introduction to Alerting.

Decryption Features

- Decryption Broker
- Automatic SAN Support for SSL Decryption
- HSM Client Upgrade and SafeNet HSM Cluster Support
- ECDSA Certificate Support for SSL Decryption with HSMs
- Decryption Port Mirroring Support Extension

© 2020 Palo Alto Networks, Inc.

Decryption Broker

Offload SSL decryption to the Palo Alto Networks firewall and decrypt traffic only once. A firewall enabled as a decryption broker forwards clear text traffic to security chains (sets of inline, third-party appliances) for additional enforcement. This allows you to consolidate security functions on the firewall, optimize network performance, and reduce the number of devices in your security infrastructure.

A decryption broker firewall uses a pair of designated forwarding interfaces to connect to the security chain. Together, the firewall and the security chain function as private analysis network—the clear text traffic flowing through this network is totally segmented from dataplane traffic. The decryption broker firewall first inspects the decrypted (now clear text) SSL traffic, and then sends it to the security chain. If you've configured multiple security chains, the firewall can perform session distribution to avoid oversubscribing any one chain. Then, last device in a security chain sends the clear text traffic back to the firewall. The firewall re-encrypts the traffic and sends it to its destination.

How you deploy decryption broker might vary depending on what type of security chain you plan to use. Two types of security chain deployments are supported: Layer 3 security chains (devices have assigned IP addresses and are configured with static routes to direct traffic) and transparent bridge security chains (devices do not have IP addresses or local routing tables and are serially connected).

Decryption broker is supported for PA-7000 Series, PA-3200 Series, PA-5200 Series, and VM-Series devices, and is supported only for outbound SSL traffic (from internal users to the internet) that is being decrypted using SSL Forward Proxy decryption. To learn more about decryption broker, and for detailed and complete steps to enable this feature, see Decryption Broker. Enabling decryption broker includes:

- Deciding what security chain deployment to use—a Layer 3 security chain or a Transparent Bridge security chain—and follow the guidelines to configure that security chain.
- Activating the free Decryption Broker license.
- Confirming that SSL Forward Proxy decryption is enabled.
- Enabling the firewall to act as a decryption broker with a Layer 3 security chain or a Transparent Bridge security chain. This includes designating a pair of Layer 3 interfaces to connect the firewall to the security chain, and optionally configuring the firewall to forward to multiple security chains.

© 2020 Palo Alto Networks, Inc.

Automatic SAN Support for SSL Decryption

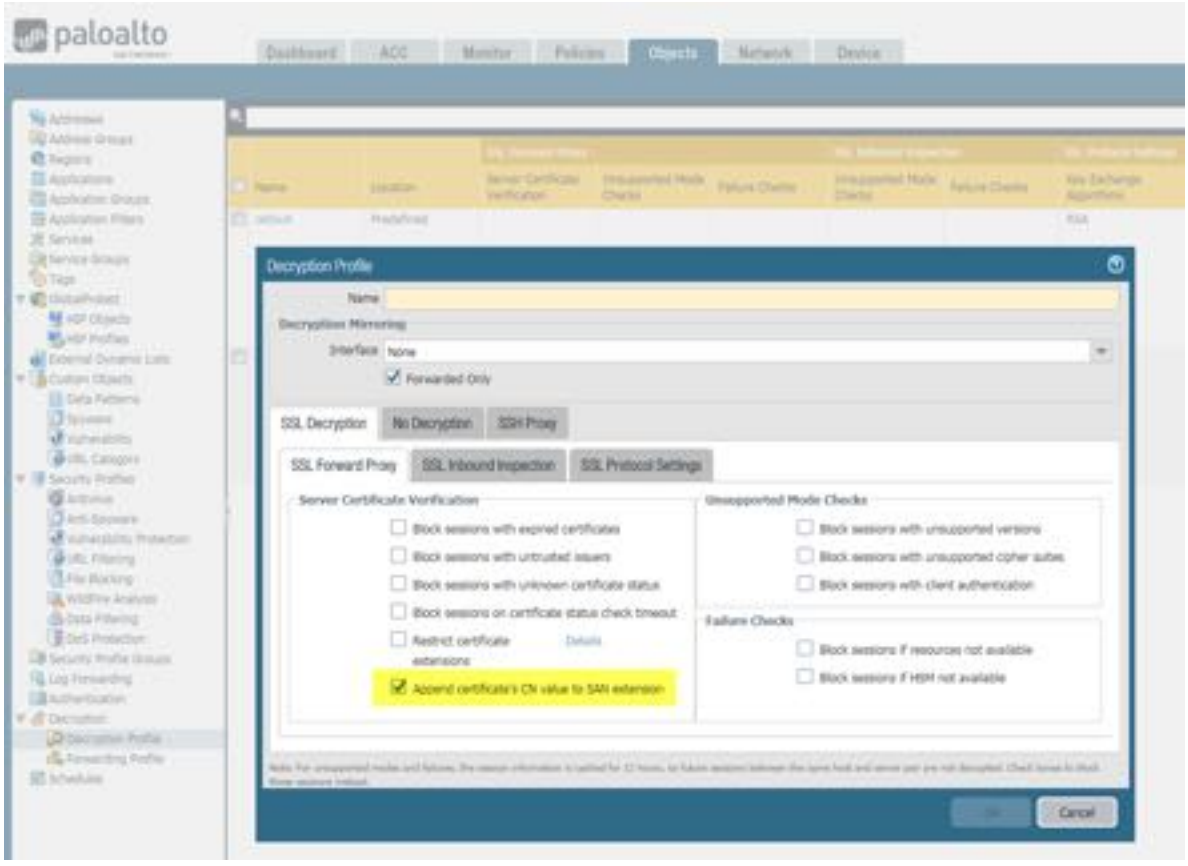
Some browsers require server certificates to use a Subject Alternative Name (SAN) to specify the domains the certificate protects, and no longer support certificate matching based on a server certificate Common Name (CN). SANs enable a single server certificate to protect multiple names; CNs are less well-defined than SANs and can protect only a single domain or all first-level subdomains on a domain. However, if a server certificates contains only a CN, browsers that require a SAN will not allow end users to connect to the requested web resource.

Now, the firewall can add a SAN to the impersonation certificate it generates to establish itself as a trusted third-party during SSL decryption. When a server certificate contains only a CN, a firewall performing SSL decryption copies the server certificate CN to the impersonation certificate SAN. The firewall presents the impersonation certificate with the SAN to the client, and the browser is able to support the connection. End users can continue to access the resources they need, and the firewall can decrypt the sessions.

STEP 1 | To decrypt and inspect SSL/TLS traffic from internal users to the web, configure SSL Forward Proxy Decryption.

STEP 2 | Enable the firewall to add a SAN to impersonation certificates.

1. Select **Objects** > **Decryption Profile** and **Add** or modify a profile.
2. Select **SSL Decryption** > **SSL Forward Proxy** and then enable the firewall to **Append certificate's CN value to SAN extension**



STEP 3 | Attach the updated Decryption profile to a Decryption policy to enforce the new setting on matching traffic (**Policies** > **Decryption** > **Options** > **Decryption Profile**).

© 2020 Palo Alto Networks, Inc.

HSM Client Upgrade and SafeNet HSM Cluster Support

When you use your firewall as a hardware security module (HSM) client to manage your digital keys, that firewall HSM client can run SafeNet client versions 5.4.2 and 6.2.2 and nCipher nShield client version 12.30. These newer HSM client versions provide necessary compatibility with newer HSM server versions. Refer to the HSM vendor documentation for the client-server compatibility matrix and for any upgrade/downgrade considerations.

It is possible that downgrading an HSM server won't be an option after you upgrade it. See HSM Client Upgrade and SafeNet HSM Cluster Support in Upgrade/Downgrade Considerations.

nCipher nShield client—Your firewall is automatically upgraded from client version 11.62 to 12.30 in PAN-OS® 8.1. The firewall HSM client can support up to two independent nCipher HSM servers.

SafeNet HSM client—Use the following information and task to configure your firewall as needed.

View the current running HSM client version by selecting **Device** > **Setup** > **HSM** and then **Select HSM Client Version** (in the Hardware Security Operations window). When you upgrade from a PAN-OS 8.0 release to PAN-OS 8.1, the firewall will use a specific SafeNet client version as follows:

- If your PAN-OS 8.0 release uses SafeNet client version 5.2.1, upgrading to PAN-OS 8.1 results in the firewall using SafeNet client version 5.4.2.
- If your PAN-OS 8.0.2 or later PAN-OS 8.0 release uses SafeNet client version 5.4.2 or 6.2.2, upgrading to PAN-OS 8.1 results in the firewall using the same SafeNet client version it was using before the upgrade. If the firewall was using SafeNet client version 5.4.2 and you want to install SafeNet client version 6.2.2, you can perform the task below after you upgrade to PAN-OS 8.1

Additionally, the number of SafeNet HSM servers supported in a high availability (HA) configuration is enhanced from an HA pair of HSMs (two) to a cluster of up to 16 HSMs. However, the HSM servers in the cluster must all run the same SafeNet version and must authenticate separately. Use a SafeNet HSM cluster only when you need to replicate the keys across the cluster. Alternatively, you can add up to 16 SafeNet HSM servers that function independently.

HSM client integration is supported on Panorama and all firewall models except for PA-800 Series, PA-500, PA-220, PA-220-R, and PA-200 firewalls.

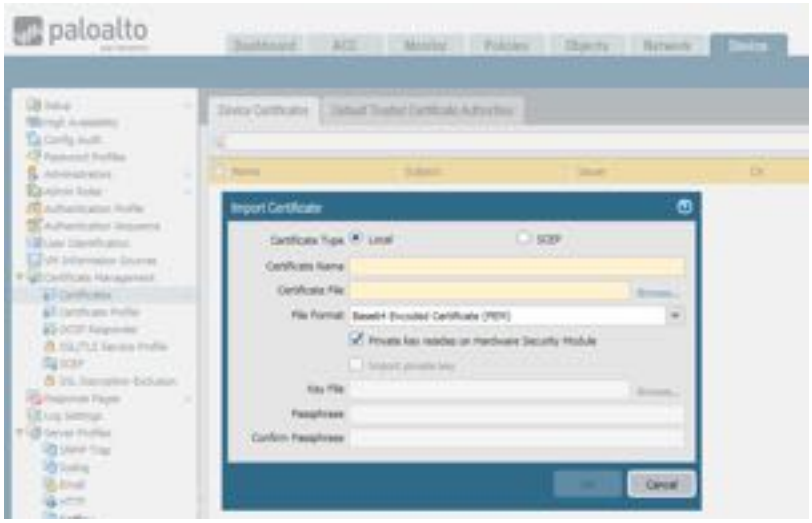
STEP 1 | Install the SafeNet Client RPM Packet Manager as described when you Set Up Connectivity with an HSM.

STEP 2 | Set up connectivity with SafeNet HSM servers as described when you Set Up Connectivity with a SafeNet Network HSM. You can establish HA with a cluster of up to 16 HSM servers.

© 2020 Palo Alto Networks, Inc.

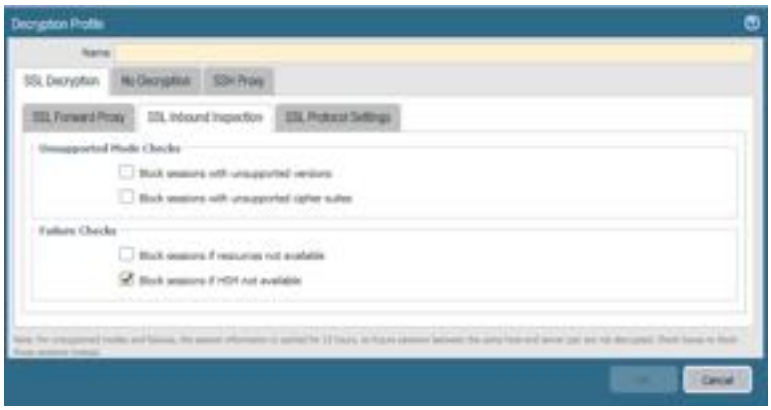
ECDSA Certificate Support for SSL Decryption with HSMs

If you use Elliptic Curve Digital Signature Algorithm (ECDSA) certificates for SSL decryption, you can now securely store your elliptic curve private keys on a third-party network HSM. This means the firewall can import or generate an ECDSA certificate when you enable the **Private key resides on Hardware Security Module** option so that the firewall can get the ECDSA key from the HSM to decrypt traffic between a client and server. Prior to PAN-OS 8.1, you could not store ECDSA certificates on an HSM.



HSM support for ECDSA certificates applies to SSL decryption in both forward proxy and inbound inspection modes.

At the start of a TLS handshake, the firewall checks the HSM connection. If the HSM connection fails at the start of a session, the firewall blocks or allows the session depending on whether you enabled the **Block session if HSM not available** option. However, when the HSM connection fails in the middle of a session, the firewall will block the session.



© 2020 Palo Alto Networks, Inc.

Decryption Port Mirroring Support Extension

Support for decryption port mirroring is now available on all hardware-based firewalls and VM-Series firewall models. Decryption port mirroring allows you to copy decrypted traffic from a firewall and then send it to a traffic collection tool, such as NetWitness or Solera. Decryption mirroring requires a Decryption Port Mirror license. This license is free of change and you can activate it through the customer support portal.

This feature is not supported on the VM-Series firewall on VMware NSX and public cloud hypervisors (Amazon Web Services, Microsoft Azure, and Google Cloud Platform).

© 2020 Palo Alto Networks, Inc.

Panorama Features

- Device Monitoring on Panorama
- Configuration Reusability for Templates and Template Stacks
- Panorama Virtual Appliance and Virtual Dedicated Log Collector on AWS
- Panorama Virtual Appliance and Virtual Dedicated Log Collector on AWS GovCloud
- Panorama Virtual Appliance and Virtual Dedicated Log Collector on Azure
- Panorama Virtual Appliance and Virtual Dedicated Log Collector on Google Cloud Platform
- Panorama Virtual Appliance and Virtual Dedicated Log Collector on KVM
- Panorama Virtual Appliance and Virtual Dedicated Log Collector on Hyper-V
- Virtual Dedicated Log Collector on ESXi
- Streamlined Panorama Deployment for Application and Threat Content Updates
- Content Update Revert from Panorama
- Direct Query of PA-7000 Series Firewalls from Panorama
- Panorama Interconnect

© 2020 Palo Alto Networks, Inc.

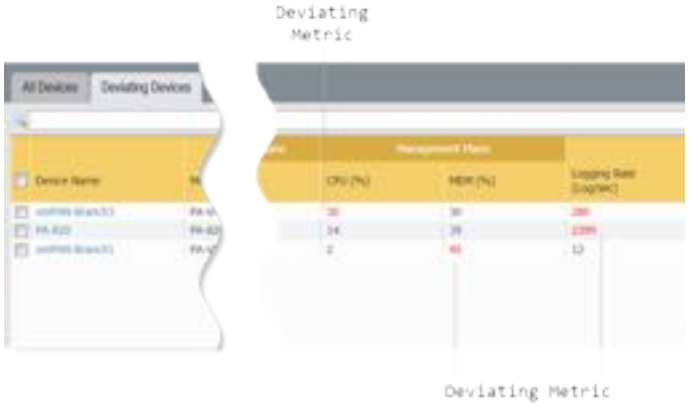
Device Monitoring on Panorama

Now you can display the health information (**Panorama** > **Managed Devices** > **Health**) or the summary information (**Panorama** > **Managed Devices** > **Summary**) for your managed firewalls. You must upgrade both Panorama™ and firewalls to PAN-OS 8.1 before you can monitor the summary session, logging, resource, and environmental performance for managed firewalls in the summary. You can view the time-trended graphs and tables for the device sessions, environmentals, interfaces, logging, resources, and high availability performance for selected devices. Panorama calculates the baseline performance for each metric using seven-day averages and standard deviation to determine a normal operating range for the specific firewall. You can identify firewalls experiencing performance related issues by viewing **Panorama** > **Managed Devices** > **Health** > **Deviating Devices** to isolate issues before they can seriously impact your business. When Panorama identifies a metric that is outside the normal operating range, the metric is marked and the firewall populates the Deviating Devices tab. For more information, see Monitor Device Health.



Device Monitoring on Panorama	
1	Managed Devices is now split in to two sections: • **Summary**—Displays summary information about firewalls managed by the Panorama. The same summary information is shown here as in previous releases. • **Health** (New)—Displays the health status of all healthy firewalls and all firewalls outside normal operating range (Deviating Devices). Selecting a firewall takes you to the Detailed Device View where you can view the time-trended graphs and tables of monitored metrics.
2	View **All Devices** for a list of all managed firewalls regardless of device health status.
3	View the **Deviating Devices** for a list of all firewalls that are reporting metrics of the calculated baseline. Metrics that deviate outside the calculated baselines appear in red text:

© 2020 Palo Alto Networks, Inc.

Device Monitoring on Panorama	
	

© 2020 Palo Alto Networks, Inc.

Configuration Reusability for Templates and Template Stacks

Templates and template stacks are improved to make it easier to manage firewalls and appliances from Panorama™ using a single template or template stack. To more easily reuse templates and template stacks, you can now create template variables in place of firewall-specific and appliance-specific IP literals in your configurations. You can define template variables at either the template level or the template stack level and you can use them to replace IP addresses, IP ranges, FQDNs, group IDs, and interfaces in IKE, VPN and HA configurations. You can override variables associated with and managed by the template or template stack on a per-firewall or appliance basis. Overriding a template or template stack variable allows you to minimize the number of templates and template stacks you need to manage while still allowing you to keep any device-specific configurations as needed. To reduce the number template and template stack configurations you need to manage, you now add firewalls and appliances to a template stack rather than to a template; the firewall or appliance prioritizes the template stack configuration. Additionally, template stacks can now reference any named object in a template that belongs to the template stack, which means you use template stacks to manage the base firewall and appliance configurations while template configurations enable you to create specific configurations. For more information, see Configure a Template or Template Stack Variables and Manage Templates and Template Stacks.



To create a template variable:

STEP 1 | Create a template and template stack using a variable name for an object. Variable names must start with the dollar sign (“$”) symbol. For example, you could use **$Panorama** as a variable for the Panorama IP address that you want to configure on multiple managed firewalls and appliances.

1. Add a Template.
2. Configure a Template Stack.

STEP 2 | In **Panorama** > **Templates**, click **Manage** (Variables column).

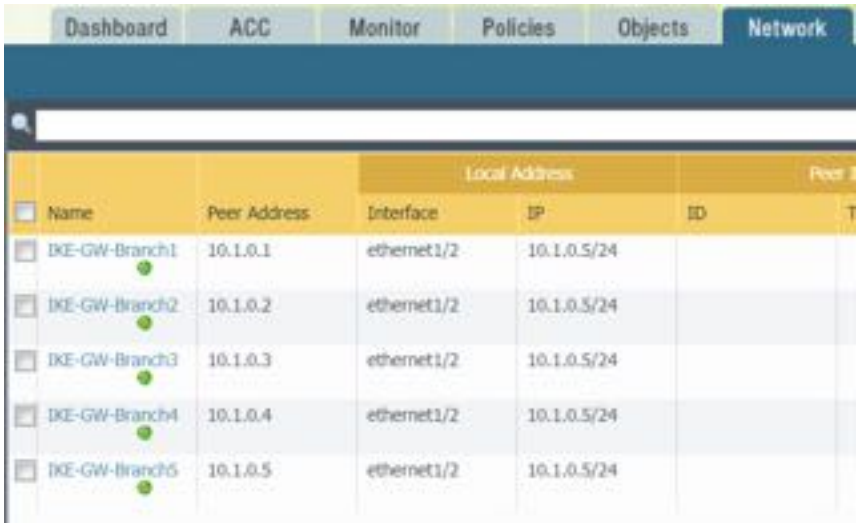
STEP 3 | **Add** a new variable.

© 2020 Palo Alto Networks, Inc.

Variables can also be created inline where variables are supported.

STEP 4 | **Commit and Push** the configuration changes to the managed firewalls and appliances.

STEP 5 | Verify that the variables were pushed to the managed firewalls and appliances. Values defined by a variable display a variable symbol ( ) and you can hover over the symbol to view the template or template stack to which the variable definition belongs. A variable that is overridden displays an overridden variable symbol ( ).



© 2020 Palo Alto Networks, Inc.

Panorama Virtual Appliance and Virtual Dedicated Log Collector on AWS

Panorama™ continues to expand the number of supported virtual environments to help reduce your physical footprint and giving you more flexibility when deploying your Palo Alto Networks® management and log collection services alongside other applications you deployed on Amazon® Web Services (AWS®). The Panorama virtual appliance on AWS supports all deployment modes (Panorama, Log Collector, and Management) and each mode shares the same processes and functionality as its M-Series counterpart. Refer to Panorama Models for more information on Panorama modes.

For more information, see Install Panorama on AWS.

STEP 1 | Log in to the Amazon Web Service Console and select the EC2 dashboard.

STEP 2 | Set up the VPC for your network needs.

STEP 3 | Deploy Panorama on Amazon Web Services.

1. On the EC2 Dashboard, **Launch Instance** and then select **My AMIs** and **Select** the Panorama 8.1.0 AMI.
2. Launch the Panorama virtual appliance as an EC2 instance.
    1. Choose the **EC2 instance type** for allocating the resources required for the Panorama virtual appliance, and click **Next: Configure Instance Details**. See the Setup Prerequisites for the Panorama VirtualAppliance for resource requirements.
    2. Select the VPC.
    3. Select the public subnet for the Panorama virtual appliance management interface.
    4. **Automatically assign a public IP address**.
    5. Click **Next: Add Storage** and **Add New Volume** to add log storage.
    6. (Optional) Add one or more tags as metadata to help you identify and group the Panorama virtual appliance. For example, add a **Name** tag with a **Value** that helps you identify which firewalls the Panorama virtual appliance manages.
    7. Create a new **Security Group** or select an existing one with—at minimum—HTTPS and SSH enabled
    8. **Review and Launch** and then verify that your selections are accurate before you **Launch.**
    9. Select an existing key pair or create a new one and acknowledge the disclaimer.

        As a best practice, create a new key for each instance of Panorama or Dedicated Log Collector on AWS.

    10. If you create a new key, download and save the private key to a safe location; the file extension is `.pem`. You cannot regenerate this key if it is lost.

        It takes 20 to 30 minutes to launch the Panorama virtual appliance using the minimum requirements.

STEP 4 | Enable management access to the Panorama virtual appliance.

1. Shut down the Panorama virtual appliance.
2. Create virtual network interfaces as needed and attach the interfaces to the Panorama virtual appliance. The virtual network interfaces are called Elastic Network Interfaces (ENIs) on AWS. You use these interfaces to manage devices from the virtual appliance.

© 2020 Palo Alto Networks, Inc.

The Panorama virtual appliance on AWS supports a single management interface; you must use this interface for log collection and device management.

3. Create or assign an Elastic IP (EIP) address to the management interface.
4. Power on the Panorama virtual appliance.

STEP 5 | Configure a new administrative password for the Panorama virtual appliance.

You must configure a unique administrative password before you can access the web interface of the Panorama virtual appliance. The private key you used to launch the Panorama virtual appliance is required to access the CLI.

STEP 6 | Activate the licenses on the Panorama virtual appliance.

STEP 7 | Finish configuring the Panorama virtual appliance for your deployment needs.

- (Management Only mode) Set Up a Panorama Virtual Appliance in Management Only Mode.
- (Log Collector mode) Skip to Step 6 to set up the Panorama virtual appliance as a Log Collector.

  When configuring the Management interface in Step 9, enter the Public IP Address of the Dedicated Log Collector. You cannot specify the IP Address, Netmask, or Gateway.
- (Panorama and Management Only mode) Configure a Managed Collector. Manage a Dedicated Log Collector from the Panorama virtual appliance for log collection of managed firewalls.

© 2020 Palo Alto Networks, Inc.

Panorama Virtual Appliance and Virtual Dedicated Log Collector on AWS GovCloud

Panorama continues to expand supported virtual environments to help reduce your physical footprint by giving you more flexibility in deploying your Palo Alto Networks management and log collection services along side other application you have deployed on Amazon Web Services. The Panorama virtual appliance on AWS supports all deployment modes (Panorama, Log Collector, and Management) and share the same processes and functionality as their M-Series counterparts. Refer to Panorama Models for more information on the Panorama modes.

For more information, see Install Panorama on AWS GovCloud.

STEP 1 | Log in to the AWS GovCloud Web Service Console and select the EC2 dashboard.

STEP 2 | Set up the VPC for your network needs.

STEP 3 | Deploy Panorama on Amazon Web Services.

1. On the EC2 Dashboard, click **Launch Instance** and then select **My AMIs** and **Select** the Panorama 8.1.0 AMI.
2. Launch the Panorama virtual appliance as an EC2 instance.
    1. Choose the **EC2 instance type** for allocating the resources required for the Panorama virtual appliance, and click **Next: Configure Instance Details**. See the Setup Prerequisites for the Panorama Virtual Appliance for resource requirements.
    2. Select the VPC.
    3. Select the public subnet for the Panorama virtual appliance management interface.
    4. Select **Automatically assign a public IP address**.
    5. Click **Next: Add Storage** and **Add New Volume** to add log storage.
    6. (Optional) Add one or more tags as metadata to help you identify and group the Panorama virtual appliance. For example, add a **Name** tag with a **Value** that helps you identify which firewalls the Panorama virtual appliance manages.
    7. Create a new **Security Group** or select an existing one with HTTPS and SSH enabled at a minimum.
    8. Select **Review and Launch** and verify that your selections are accurate before you select **Launch.**
    9. Select an existing key pair or create a new one and acknowledge the disclaimer.

        As a best practice, create a new key for each instance of Panorama or Dedicated Log Collector on AWS.
    10. If you created a new key, download and save the private key to a safe location; the file extension is `.pem`. You cannot regenerate this key if lost.

        It takes 20-30 minutes to launch the Panorama virtual appliance using the minimum requirements.

STEP 4 | Enable management access to the Panorama virtual appliance.

1. Shut down the Panorama virtual appliance.
2. Create virtual network interface(s) and attach the interface(s) to the Panorama virtual appliance. The virtual network interfaces are called Elastic Network Interfaces (ENIs) on AWS. These interfaces are used for managing devices from the virtual appliance.

    The Panorama virtual appliance on AWS supports a single management interface. You must use this interface for log collection and device management.

© 2020 Palo Alto Networks, Inc.

3. Create or assign an Elastic IP (EIP) address to the management interface.
4. Power on the Panorama virtual appliance.

STEP 5 | Configure a new administrative password for the Panorama virtual appliance.

You must configure a unique administrative password before you can access the web interface of the Panorama virtual appliance. To access the CLI, the private key used to launch the Panorama virtual appliance is required.

STEP 6 | Activate the licenses on the Panorama virtual appliance.

STEP 7 | Complete configuring the Panorama virtual appliance for your deployment needs.

- (Management Only mode) Set Up a Panorama Virtual Appliance in Management Only Mode.
- (Log Collector mode) Begin at Step 6 to set up the Panorama virtual appliance as a Log Collector.

  When configuring the Management interface in Step 9.3, enter the Public IP Address of the Dedicated Log Collector. You do not have the option to enter the IP Address, Netmask, or Gateway.
- (Panorama and Management Only mode) Configure a Managed Collector. Manage a Dedicated Log Collector from the Panorama virtual appliance for log collection of managed firewalls.

© 2020 Palo Alto Networks, Inc.

Panorama Virtual Appliance and Virtual Dedicated Log Collector on Azure

You can now deploy Panorama™ and a Dedicated Log Collector on Microsoft® Azure®. The Panorama virtual appliance on Azure supports all deployment modes (Panorama, Log Collector, and Management Only) and each mode shares the same processes and functionality as its M-Series counterpart. Refer to Panorama Models for more information on Panorama modes.

For more information, see Install Panorama on Azure.

STEP 1 | Log in to the Microsoft Azure portal.

STEP 2 | Deploy the Panorama virtual appliance.

1. In the Azure Dashboard, select **Virtual machines**, **Add** a VM, search for Panorama (BYOL), and **Create** the Panorama virtual appliance.
2. Configure the basic settings for the Panorama virtual appliance.
3. Configure the Panorama virtual appliance size. Refer to the Setup Prerequisites for the Panorama VirtualAppliance for sizing requirements. The Panorama virtual appliance defaults to Management Mode if you do not configure it to meet the minimum requirements for Panorama or Log Collector mode.
4. Configure storage and networking.

   The Panorama virtual appliance on Azure supports a single management interface. You must use this interface for log collection and device management.
5. Review the summary, accept the terms of use and privacy policy, and then **Create** (deploy) the Panorama virtual appliance.
6. Verify that you have successfully deployed the Panorama virtual appliance.

   It takes 20 to 30 minutes to launch the Panorama virtual appliance and can take longer depending on the resources you configured for the virtual machine. ICMP protocol is not permitted by Microsoft Azure so you cannot ping the Panorama virtual appliance to determine whether it is successfully deployed.

STEP 3 | Log in to the web interface of the Panorama virtual appliance.

STEP 4 | Activate the licenses on the Panorama virtual appliance.

STEP 5 | Finish configuring the Panorama virtual appliance for your deployment needs.

- (Management Only mode) Set Up a Panorama Virtual Appliance in Management Only Mode.
- (Log Collector mode) Skip to Step 6 to set up the Panorama virtual appliance as a Log Collector.

  When configuring the Management interface in Step 9, enter the Public IP Address of the Dedicated Log Collector. You cannot specify IP Address, Netmask, or Gateway.
- (Panorama and Management Only mode) Configure a Managed Collector. Manage a Dedicated Log Collector from the Panorama virtual appliance for log collection of managed firewalls.

© 2020 Palo Alto Networks, Inc.

Panorama Virtual Appliance and Virtual Dedicated Log Collector on Google Cloud Platform

You can now deploy Panorama™ and a Dedicated Log Collector on Google® Cloud Platform (GCP™). The Panorama virtual appliance on GCP supports all deployment modes (Panorama, Log Collector, and Management Only) and each mode shares the same processes and functionality as its M-Series counterpart. Refer to Panorama Models for more information on Panorama modes. Panorama virtual appliance and virtual Dedicated Log Collector on GCP is available only on PAN-OS 8.1.1 and later releases.

For more information, see Install Panorama on Google Cloud Platform.

STEP 1 | Download the Panorama virtual appliance image.

1. Log in to the Palo Alto Networks Support Portal.
2. Select **Updates** > **Software Updates** and filter by **Panorama Base Images**.
3. Download the latest version of the Panorama on GCP `tar.gz` image.

STEP 2 | Upload the Panorama virtual appliance image to the Google Cloud Platform.

1. Log in to the Google Cloud Console.
2. From the **Products and Services** menu, select **Storage**.
3. Click **Create Bucket**, configure the new storage bucket and click **Create.**
4. Select the storage bucket you created in the previous step, click **Upload files**, and select the Panorama virtual appliance image you downloaded.
5. From the **Products and Services** menu, select **Compute Engine** > **Images**.
6. Click **Create Image** and create the Panorama virtual appliance image:
    1. **Name** the Panorama virtual appliance image.
    2. In the **Source** field, select **Cloud Storage file** from the drop-down menu.
    3. Click **Browse** and navigate to the storage bucket where you uploaded the Panorama virtual appliance image, and **Select** the uploaded image.
    4. **Create** the Panorama virtual appliance image.

STEP 3 | Configure the Panorama virtual appliance.

1. From the **Products and Services** menu and select the **Compute Engine**.
2. Click **Create Instance** to begin deploying the Panorama virtual appliance.
3. Add a descriptive **Name** to easily identify the Panorama virtual appliance.
4. Select the **Zone** where you want to deploy the Panorama virtual appliance.
5. Allocate the **Machine Type** CPU cores and memory. Refer to the Setup Prerequisites for the Panorama Virtual Appliance for minimum resource requirements.
6. For the **Boot Disk**, click **Change** > **Custom image**, select the Panorama image file you uploaded in Step 2, and click **Select**.
7. Under **Identity and API access**, **Allow full access to all Cloud APIs**.
8. Under **Firewall**, **Allow HTTPS traffic**.

STEP 4 | Enable management access to the Panorama virtual appliance.

1. Expand **Management, disks, networking SSH keys**.
2. Enable access to the serial port so you can manage the Panorama virtual appliance.

© 2020 Palo Alto Networks, Inc.

3. Configure the management interface on the Panorama virtual appliance.

   To learn more about how to reserve IP addresses, refer to the information about how to Reserve a Static Internal IP Address and how to Reserve a Static External IP Address.

   The Panorama virtual appliance on GCP supports a single management interface; ou must use this interface for log collection and device management.
4. Configure the SSH key. You need an SSH key to access the Panorama virtual appliance CLI to configure the administrative user password after initial deployment.

STEP 5 | (Panorama and Log Collector mode) Add additional storage for log collection. Repeat this step as needed to add additional virtual logging disks.

If you intend to use the Panorama virtual appliance in Panorama mode or as a Dedicated Log Collector, add the virtual logging disks during the initial deployment. By default, the Panorama virtual appliance is in Panorama mode for the initial deployment when you meet the Panorama mode resource requirements and you have added at least one virtual logging disk. Otherwise, the Panorama virtual appliance defaults to Management Only mode. Change the Panorama virtual appliance to Management Only mode if you need to only manage devices and Dedicated Log Collectors and you don't need to collect logs locally.

The Panorama virtual appliance on GCP supports only 2TB logging disks and, in total, supports up to 24TB of log storage. You cannot add a logging disk smaller than 2TB or a logging disk with a size that is not evenly divisible by the 2TB logging disk requirement because the Panorama virtual appliance partitions logging disks larger than 2TB into 2TB partitions.

STEP 6 | **Create** (deploy) the Panorama virtual appliance. The Panorama virtual appliances takes approximately 10 minutes to boot up after initial deployment.

STEP 7 | Configure a new administrative password for the Panorama virtual appliance.

Use a private key to access the CLI and configure a unique administrative password so that you can access the web interface of the Panorama virtual appliance.

STEP 8 | Activate the licenses on the Panorama virtual appliance.

STEP 9 | Finish configuring the Panorama virtual appliance for your deployment needs.

- (Management Only mode) Set Up a Panorama Virtual Appliance in Management Only Mode.
- (Log Collector mode) Skip to Step 6 to set up the Panorama virtual appliance as a Log Collector.

  When configuring the Management interface in Step 9, enter the Public IP Address of the Dedicated Log Collector. You cannot specify IP Address, Netmask, or Gateway.
- (Panorama and Management Only mode) Configure a Managed Collector. Manage a Dedicated Log Collector from the Panorama virtual appliance for log collection of managed firewalls.

© 2020 Palo Alto Networks, Inc.

Panorama Virtual Appliance and Virtual Dedicated Log Collector on KVM

You can now deploy Panorama™ and a Dedicated Log Collector on KVM. The Panorama virtual appliance on KVM supports all deployment modes (Panorama, Log Collector, and Management Only) and each mode shares the same processes and functionality as its M-Series counterpart. Refer to Panorama Models for more information on Panorama modes. Panorama virtual appliance and virtual Dedicated Log Collector on KVM is available only on PAN-OS 8.1.2 and later releases.

For more information, see Install Panorama on KVM.

STEP 1 | Launch the Panorama virtual appliance.

1. Download the Panorama virtual appliance for KVM image.
2. Create a new virtual machine image and add the Panorama virtual appliance image for KVM to the Virtual Machine Manager.
3. Configure the **Memory** and **CPU** settings.
4. **Name** the Panorama virtual appliance, enable configuration customization, and select the management interface bridge.

   The Panorama virtual appliance on KVM supports a single management interface. You must use this interface for log collection and device management.
5. Configure the virtual system disk settings.
6. Configure the virtual machine console display to use the VNC server to interact with the virtual machine.
7. (Panorama and Log Collector mode) Add additional storage for log collection. Repeat this step as needed to add additional virtual logging disks.

   If you intend to use the Panorama virtual appliance in Panorama mode or as a Dedicated Log Collector, add the virtual logging disks during the initial deployment. By default, the Panorama virtual appliance is in Panorama mode for the initial deployment when you meet the Panorama mode resource requirements and you have added at least one virtual logging disk. Otherwise, the Panorama virtual appliance defaults to Management Only mode. Change the Panorama virtual appliance to Management Only mode if you need to only manage devices and Dedicated Log Collectors and you don't need to collect logs locally.

   The Panorama virtual appliance on KVM supports only 2TB logging disks and, in total, supports up to 24TB of log storage. You cannot add a logging disk smaller than 2TB or a logging disk with a size that is not divisible by the 2TB logging disk requirement because the Panorama virtual appliance partitions logging disks larger than 2TB into 2TB partitions.
8. **Begin Installation**. The Panorama virtual appliances takes approximately 10 minutes to boot up.

STEP 2 | Configure the network access settings for the management interface.

STEP 3 | Activate the licenses on the Panorama virtual appliance.

STEP 4 | Finish configuring the Panorama virtual appliance for your deployment needs.

- (Management Only mode) Set Up a Panorama Virtual Appliance in Management Only Mode.
- (Log Collector mode) Skip to Step 6 to set up the Panorama virtual appliance as a Log Collector.



When configuring the Management interface in Step 9, enter the Public IP Address of the Dedicated Log Collector. You cannot specify IP Address, Netmask, or Gateway.

© 2020 Palo Alto Networks, Inc.

- (Panorama and Management Only mode) Configure a Managed Collector. Manage a Dedicated Log Collector from the Panorama virtual appliance for log collection of managed firewalls.

© 2020 Palo Alto Networks, Inc.

Panorama Virtual Appliance and Virtual Dedicated Log Collector on Hyper-V

Panorama continues to expand supported virtual environments to help reduce your physical footprint by giving you more flexibility in deploying your Palo Alto Networks management and log collection services along side other application you have deployed on Hyper-V The Panorama virtual appliance on Hyper-V supports all deployment modes (Panorama, Log Collector, and Management) and share the same processes and functionality as their M-Series counterparts. Refer to Panorama Models for more information on the Panorama modes.

For more information, see Install Panorama on Hyper-V.

STEP 1 | Download the VHDX file from the Palo Alto Networks Support Portal.

STEP 2 | Set up any vSwitch(es) that you will need. For more information, review the Virtual Switch Types for more information.

STEP 3 | Install the Panorama virtual appliance.

1. On the Hyper-V Manager, select the host and select **Action** > **New** > **Virtual Machine.** Configure the following settings in the New Virtual Machine Wizard:
    1. Choose a **Name** and **Location** for the Panorama virtual appliance. The Panorama virtual appliance stores the VHDX file at the specified location.
    2. Choose **Generation 1**. This is the default option and the only version supported.
    3. For **Startup Memory**, assign the memory based on the intended system mode. See the Setup Prerequisites for the Panorama Virtual Appliance for the memory requirements for each mode.

        Do not enable dynamic memory; the Panorama virtual appliance requires static memory allocation.
    4. Configure **Networking**. Select an external vSwitch to connect the management interface on the firewall.
    5. To connect the **Virtual Hard Disk**, select **Use an existing virtual hard disk** and browse to the VHDX file you downloaded earlier.
    6. Review the summary and click **Finish**.

STEP 4 | Allocate the Panorama virtual appliance CPU cores.

Review the Setup Prerequisites for the Panorama Virtual Appliance for minimum resource requirements.

If you plan to use the Panorama virtual appliance as a Dedicated Log Collector, ensure that you configure the appliance with the required resources during initial deployment. The Panorama virtual appliance does not remain in Log Collector mode if you resize the virtual machine after you deploy it, and this results in a loss of log data.

1. In the **Hardware** list, select **Processor**.
2. Edit the currently allocated **Number of virtual processors**.

STEP 5 | Connect at least one network adapter for the dataplane interface on the firewall. Repeat this to create additional network interfaces on the Panorama virtual appliance.

1. Select **Settings** > **Hardware** > **Add Hardware** and select the **Hardware type** for your network adapter.

© 2020 Palo Alto Networks, Inc.

Legacy Network Adapter and SR-IOV are not supported. If selected, the VM-Series firewall will boot into maintenance mode.

2. Click **OK**.

STEP 6 | (Panorama and Log Collector mode) Add additional storage for log collection. Repeat this step as needed to add additional virtual logging disks. If you want to deploy the Panorama virtual appliance in Management Only mode, continue to Step 7.

If you intend to use the Panorama virtual appliance in Panorama mode or as a Dedicated Log Collector, add the virtual logging disks during the initial deployment. By default, the Panorama virtual appliance is in Panorama mode for the initial deployment when you meet the Panorama mode resource requirements and have added at least one virtual logging disk.

The Panorama virtual appliance on Hyper-V only supports 2TB logging disks, and in total supports up to 24TB of log storage. You are unable to add a logging disk smaller than 2TB, or a logging disk with a size not divisible by the 2TB logging disk requirement. The Panorama virtual appliance partitions logging disks larger than 2TB into 2TB partitions.

1. On the Hyper-V Manager, select the host and select **Action** > **New** > **Hard Disk.**
2. If you see the Before You Begin prompt, click **Next** to begin adding the virtual logging disk.
3. For the Disk Format, select **VHDX**. Click **Next**. to continue
4. For the Disk Type, select **Fixed Size** or **Dynamically Expanding** based on your needs. Click **Next** to continue.
5. Specify the **Name** and **Location** for the virtual logging disk file. Click **Next** to continue.
6. To configure the disk, select **Create a new virtual hard disk** and enter the disk size. Click **Next** to continue.
7. Review the Summary and **Finish** adding the virtual hard logging disk.

STEP 7 | Power on the Panorama virtual appliance.

STEP 8 | Configure the IP address of the management interface.

STEP 9 | Configure a new administrative password for the Panorama virtual appliance.

You must configure a unique administrative password before you can access the web interface of the Panorama virtual appliance. To access the CLI, the private key used to launch the Panorama virtual appliance is required.

STEP 10 | Activate the licenses on the Panorama virtual appliance.

STEP 11 | Complete configuring the Panorama virtual appliance for your deployment needs.

- (Management Only mode) Set Up a Panorama Virtual Appliance in Management Only Mode.
- (Log Collector mode) Begin at Step 6 to set up the Panorama virtual appliance as a Log Collector.

  When configuring the Management interface in Step 9.3, enter the Public IP Address of the Dedicated Log Collector. You do not have the option to enter the IP Address, Netmask, or Gateway.
- (Panorama and Management Only mode) Configure a Managed Collector. Manage a Dedicated Log Collector from the Panorama virtual appliance for log collection of managed firewalls.

© 2020 Palo Alto Networks, Inc.

Virtual Dedicated Log Collector on ESXi

Panorama continues to expand supported virtual environments to help reduce your physical footprint by giving you more flexibility in deploying your Palo Alto Networks management and log collection services along side other application you have deployed on the ESXi hypervisor. The Panorama virtual appliance deployed as a Dedicated Log Collector on ESXi supports all deployment modes share the same processes and functionality as their M-Series counterparts. For more information, see Set Up the M-Series Appliance as a Dedicated Log Collector.

STEP 1 | Install Panorama on an ESXi server.

STEP 2 | Switch from Panorama mode to Log Collector mode on each Panorama virtual appliance to set up a Dedicated Log Collector.

STEP 3 | Record the serial number of the Log Collector. You will need the serial number to add the Log Collector as a managed collector.

STEP 4 | Add the Log Collector as a managed collector on the Panorama management server.

STEP 5 | Enable the logging disks.

© 2020 Palo Alto Networks, Inc.

Streamlined Panorama Deployment for Application and Threat Content Updates

When using Panorama to deploy content updates to managed firewalls, you can now more easily configure dynamic updates schedules for multiple firewalls at once based on platform, device group, and tag (instead of applying schedules based on individual firewall serial numbers, as was previously supported). If yours is a mission-critical network, this is especially useful to stagger content updates across your network and to enforce content update installation threshold. A content update installation threshold only allows firewalls to install Applications and Threat content updates that have been successfully functioning in customer environments for a given amount of time, and a separate installation threshold for content updates with new App-IDs gives you extra time to assess how new application signatures impact your security policy. Previously, you could only use Panorama to configure content updates thresholds for managed firewalls that were connecting to the Palo Alto Networks Update Server directly, not those firewalls that retrieve content updates from Panorama.

For guidance on how to best deploy Application and Threat content updates based on your organization's network security and application availability requirements, review the Best Practices for Application and Threat Content Updates. Then, to easily apply a Dynamic Updates Schedule to several devices at once, and to set a Content Update Threshold for those devices, start by editing or adding a Dynamic Updates schedule for managed firewalls:

STEP 1 | Select **Panorama** > **Device Deployment** > **Dynamic Updates** > **Schedules**.

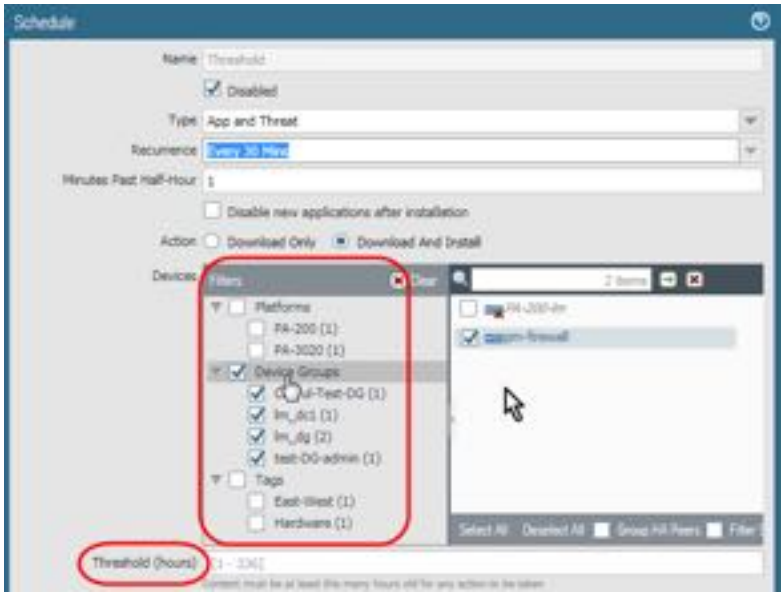
STEP 2 | Set the schedule **Type** to **App and Threat**.

STEP 3 | Set the schedule **Action** to **Download and Install**.

STEP 4 | Use **Filters** to set a Dynamic Updates schedule for many devices at once based on Platforms, Device Groups, and Tags.

STEP 5 | Set a **Threshold** requirement for new Applications and Threat content releases—managed firewalls only retrieve and install content releases from Panorama that have been available and functioning in customer environments for at least the amount of time that you define.

© 2020 Palo Alto Networks, Inc.



© 2020 Palo Alto Networks, Inc.

Content Update Revert from Panorama

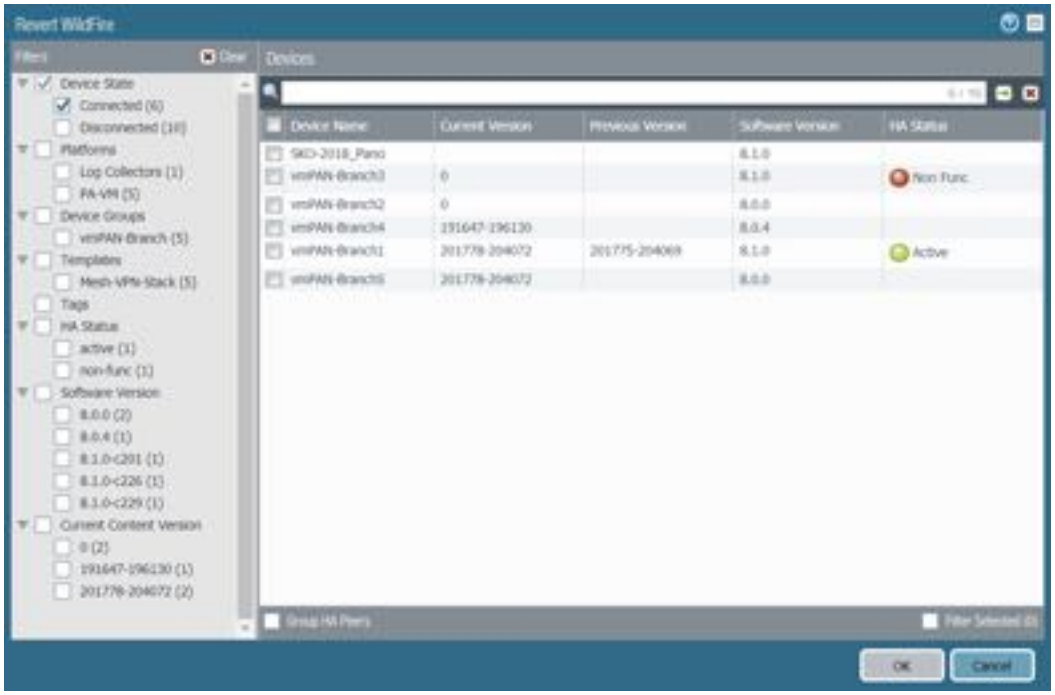
You can quickly revert the Applications, Applications and Threats, Antivirus, WildFire®, and WildFire content versions of one or more firewalls, Log Collectors, or WildFire appliances directly from Panorama™. Using Panorama to revert content versions installed on managed firewalls and appliances provides a centralized workflow that helps mitigate risks associated with new or modified applications or new threat signatures that are introduced in a content update. The ability to revert content release versions on one or more managed firewalls or appliances from Panorama allows you to quickly revert a content update (such as one that adversely affected your network operations and security) from a single location rather than having to access and revert content updates for one managed firewall or appliance at a time. Additionally, for each firewall and appliance you revert, Panorama generates a system log. For more information, see the Best Practices for Content and Threat Content Updates before you deploy content updates to your managed firewalls and appliances.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Select **Panorama** > **Device Deployment** > **Dynamic Updates** and **Revert Content**.

STEP 3 | Select the content update type you need to revert.

STEP 4 | Select one or more firewalls or appliances on which you need to revert content and click **OK**. The content version you revert to must be an earlier version than the version currently installed on the firewall or appliance.



© 2020 Palo Alto Networks, Inc.

Direct Query of PA-7000 Series Firewalls from Panorama

Because PA-7000 Series firewall can now forward logs to Panorama, Panorama no longer treats the PA-7000 Series firewalls it manages as Log Collectors. If you have not configured the PA-7000 Series firewalls to forward logs to Panorama, all logs a managed PA-7000 Series firewall generates are only viewable from the local firewall and not from Panorama. If you do not yet have a log forwarding infrastructure that is capable of handling the logging rate and volume from the PA-7000 Series firewalls, you can now enable Panorama to directly query PA-7000 Series firewalls when monitoring logs.



For Panorama to directly query PA-7000 Series firewalls, the firewalls must be running PAN-OS 8.0.8 or later.

With this new functionality, Panorama now provides two options for monitoring logs and running reports for managed PA-7000 Series firewalls:

- (New) Enable Panorama to directly query managed PA-7000 Series firewalls when monitoring logs.

  To enable Panorama to directly query the PA-7000 Series firewalls without requiring the firewalls to forward logs, you must enter the following command from the Panorama CLI:

  ```
  admin@panorama> debug reportd send-request-to-7k yes
  ```

  After running the command, you will be able to view logs for managed PA-7000 Series firewalls on the Panorama **Monitor** tab. Additionally, as with all managed devices, you can also generate reports that include PA-7000 Series log data by selecting **Remote Device Data** as the **Data Source**.

- Configure the managed PA-7000 Series firewalls to forward logs to Panorama.

  Before enabling your PA-7000 Series firewalls to forward logs to Panorama, make sure you have a logging infrastructure that will handle the logging rate and volume. Refer to the table in Panorama Models to determine if you have the right logging capacity. Additionally, if you have enabled Panorama to directly query PA-7000 Series firewalls, you must disable this before you enable log forwarding by entering the following command from the Panorama CLI:

  ```
  > debug reportd send-request-to-7k no
  ```

  After you have enabled your PA-7000 Series firewalls to forward logs to Panorama, the PA-7000 Series log data will be aggregated within all Panorama views: Application Command Center (ACC), the App-Scope, the log viewer (**Monitor** tab), and the standard, customizable reporting options on Panorama.

© 2020 Palo Alto Networks, Inc.

Panorama Interconnect

The Panorama™ Interconnect plugin enables you to establish a Panorama Controller that manages Panorama Nodes and centralizes configuration management for consistent firewall configurations and policies across large-scale firewall deployments. This also helps reduce the time you spend synchronizing configurations on multiple Panorama management servers and improves the response time for security vulnerabilities in the event of a misconfiguration.

For more information, see Manage Large-Scale Firewall Deployments.

STEP 1 | Set Up Panorama management servers that shall function as the Controller and Nodes.

STEP 2 | Install the Panorama Interconnect plugin. You must install the plugin on the Panorama Controller and all Panorama Nodes.

1. Log in to the Panorama Web Interface.
2. Select **Panorama** > **Plugins** and search for **`Interconnect`**.
3. **Download** and **Install** the Panorama Interconnect plugin.
4. Click **Commit** > **Commit to Panorama** to finish installing the Panorama Interconnect plugin.

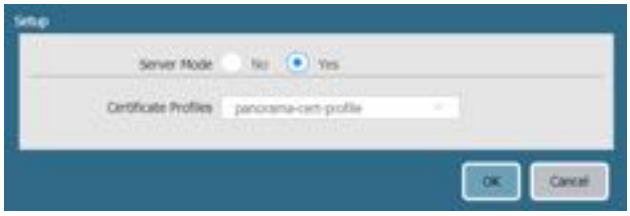
STEP 3 | Enable Authentication Between the Panorama Controller and the Panorama Nodes.

1. Obtain the certificate authority (CA) certificate for the Panorama Controller.
2. Generate the Panorama Node certificate.
3. Create a certificate profile for authenticating the Panorama Nodes.

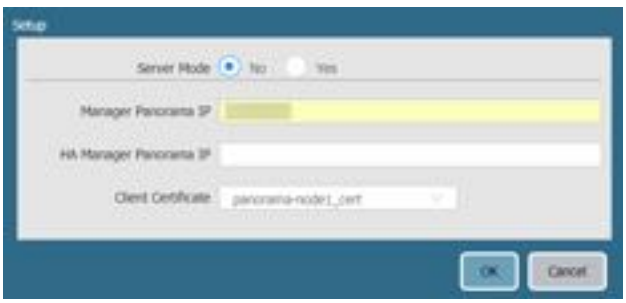
STEP 4 | Set Up the Panorama Interconnect Plugin.

1. Set up the Panorama Controller. Repeat this step on the high availability (HA) peer if the Panorama Controller is in an HA configuration.

   After you configure the Panorama management server as the Panorama Controller, you cannot reconfigure that Controller as a Panorama Node. Verify that you are configuring the correct Panorama management server as the Panorama Controller before you continue to the next step.



2. Set up the Panorama Node. Repeat this step for all Panorama Nodes.



3. Add the Panorama Nodes to the Panorama Controller.

© 2020 Palo Alto Networks, Inc.



STEP 5 | Add a Device Group on the Panorama Controller.

STEP 6 | Configure a Template Stack on the Panorama Controller.

STEP 7 | Push the Common Panorama Configuration to Panorama Nodes.

STEP 8 | Synchronize the Panorama Node with the Panorama Controller.

STEP 9 | Add one or more firewalls to be managed by a Panorama Node.

1. Add a Firewall to a Panorama Node.
2. Import Multiple Firewalls to a Panorama Node.

STEP 10 | Push the Panorama Node Configuration to Managed Devices.

Content Inspection Features

- SCTP Security
- Rapid Deployment of the Latest Threat Prevention
- Tools to Avoid or Mitigate Content Update Issues
- Option to Hold Web Requests During URL Category Lookup

© 2020 Palo Alto Networks, Inc.

SCTP Security

Stream Control Transmission Protocol (SCTP) is a reliable message-based transport protocol (number 132) that sends multiple streams of signaling, voice, and other data simultaneously. Mobile networks use SCTP to transport signaling traffic on various interfaces, such as S1-MME, S6a, and X2.

You use the multilayered approach of your firewall to secure your SCTP traffic. You can filter SCTP traffic based on payload protocol IDs (PPIDs). You can apply granular-level filtering on Diameter traffic over SCTP and SS7 traffic over SCTP. You can validate SCTP packets to ensure they comply with RFC4960. You can also protect against flooding of SCTP initiation (INIT) packets. In the case of mobile networks, these SCTP security measures help to prevent attackers from causing network congestion and outages that disrupt data and voice services of mobile subscribers and IoT devices connected to these networks. Additionally, you can view SCTP logs, ACC information, and reports to verify configurations and gain visibility into the SCTP events and traffic between two endpoints.

Only PA-5200 Series and VM-Series firewall support SCTP security in PAN-OS 8.1 releases.

STEP 1 | Enable SCTP security.

1. Select **Device** > **Setup** > **Management** and edit the General Settings to enable **SCTP Security**.
2. Click **OK**.

STEP 2 | Create an SCTP Protection profile and specify the checks and filters you want to apply to SCTP traffic.

1. **Add** a profile by **Name** (**Objects** > **Security Profiles** > **SCTP Protection**).
2. Select **SCTP Inspection** to configure the action the firewall takes on unknown chunks, non-compliant chunks, and chunks of invalid length. Generating a log to alert you and blocking packets that have invalid chunks help you secure your SCTP traffic.

STEP 3 | Select the Log Settings for the profile—options to generate SCTP logs for allowed chunks, association start or end, and state failure events.

STEP 4 | Select **Filtering Options** for the profile so you can filter protocols running on top of SCTP.

1. Add SCTP filters to allow, block, or generate an alert for PPIDs.
2. Add Diameter filters to allow, block, or generate an alert for Diameter Application IDs, Command Codes, and Attribute/Value Pairs.
3. Add SS7 filters to allow, block, or generate an alert for SS7 chunks based on SCCP Calling Party SSN, SCCP Calling Party GT, and Operation Codes.

STEP 5 | Apply the SCTP Protection profile to a security policy rule.

1. Select **Policies** > **Security** and select a Security policy rule.
2. Select **Actions** and in the Profile Setting section, select the **SCTP Protection** profile you created. Configure the rest of the Security policy rule and save it.

STEP 6 | Allocate SCTP log storage on the firewall if you want to capture SCTP logs.

Select **Device** > **Setup** > **Management**, edit the Logging and Reporting Settings, and select **Log Storage**. Enter quota percentages for SCTP, SCTP Summary, and the SCTP hourly, daily, and weekly summaries.

STEP 7 | View information about your SCTP traffic.

1. Select **Monitor** > **Logs** > **SCTP** to view the SCTP logs and detailed logs.

© 2020 Palo Alto Networks, Inc.

2. Select **Monitor** > **Logs** > **Traffic** and select the Detailed Log View ( ) for a log where the **Application** is **sctp** to view a detailed traffic log for an SCTP association.
3. Select **ACC** > **Mobile Network Activity** to view SCTP events and association activity.
4. View predefined reports about SCTP events and errors by selecting **Device** > **Setup** > **Management**. Edit the Logging and Reporting Setting section and, for Predefined Reports, select any of the SCTP reports.
5. Create a custom report on SCTP events by selecting **Monitor** > **Manage Custom Reports** and adding a custom report that uses the SCTP database.

© 2020 Palo Alto Networks, Inc.

Rapid Deployment of the Latest Threat Prevention

The following two features enable you to quickly deploy the latest threat protections while ensuring availability for mission-critical applications:

- New App-ID Threshold— Fine tune content update installation thresholds so that threat updates install immediately as they're released, and new App-IDs are installed only after you've had an opportunity to make any necessary security policy updates.
- Streamlined Panorama Deployment for Application and Threat Content Updates—Use Panorama to set an installation threshold for new App-IDs, and stagger the roll-out of new content to locations with less business risk first.

© 2020 Palo Alto Networks, Inc.

Tools to Avoid or Mitigate Content Update Issues

Palo Alto Networks Application and Threat Content Updates undergo rigorous performance and quality assurance; however, because there are so many possible variables in a customer environment, there are rare occasions where a content release might impact a network in a way that we did not foresee. The following features are intended to help both you and us avoid and mitigate an issue with a content release, so that there is as little impact to your network as possible.

Review the Best Practices for Application and Threat Content Updates for guidance on how to best deploy content updates based on your organization's network security and application availability needs.

- The firewall can now validate that a downloaded content update is still Palo Alto Networks-recommended at the time of installation.

  This check, which the firewall performs by default, is helpful in cases where content updates are downloaded from the Palo Alto Networks update server (either manually or on a schedule) ahead of installation. Because there are rare instances where Palo Alto Networks removes a content update from availability, this option prevents the firewall from installing a content update that Palo Alto Networks has deprecated, even if the firewall has already downloaded it.

- The threat intelligence telemetry data that the firewall sends to Palo Alto Networks now includes information that Palo Alto Networks can use to identify and troubleshoot issues with content updates.

  This new telemetry data helps us to quickly recognize a content update that is impacting firewall performance or security policy enforcement in unexpected ways, across the Palo Alto Networks customer base. We can quickly determine the firewall platforms or types of firewall deployments that are affected in order to help you to mitigate impact to your own network, or avoid the issue altogether.

  Make sure that you've enabled the firewall to collect and share telemetry data with Palo Alto Networks:

  1. Select **Device** > **Setup** > **Telemetry**.
  2. Edit the **Telemetry** settings and **Select All**.
  3. Click **OK** and **Commit** to save your changes.

- Palo Alto Networks can now directly alert you to a critical content release issue; we'll give you the information you need to understand if and how the issue affects you, along with steps to move forward.

  Palo Alto Networks can now issue alerts about content update issues directly to the firewall web interface or—if you have log forwarding enabled—to the external service you use for monitoring.

  In the firewall web interface, critical alerts about content issues are displayed similarly to the Message of the Day. When Palo Alto Networks issues a critical alert about a content update, the alert is displayed by default when you log into the firewall web interface. If you're already logged into the firewall web interface, you will notice an exclamation appear over the message icon on the menu bar located at the bottom of the web interface—click on the message icon to view the alert.

  Critical content update alerts are also logged as system log entries with the Type `general`and the event ID `palo-alto-networks-message`. Use the following filter to view these log entries: `( subtype eq dynamic-updates ) and ( eventid eq palo-alto-networks-message)`.

© 2020 Palo Alto Networks, Inc.

Set up log forwarding to send these entries to any external services that you use for monitoring network and firewall activity. This allows you to make sure that the appropriate personnel is notified when Palo Alto Networks issues an alert, so that they can take action as needed.

- After being notified about an issue with a content update, you can now use Panorama to revert managed firewalls to the last content update version, instead of manually reverting the content version for individual firewalls. To learn more, see Content Update Reversion from Panorama.

© 2020 Palo Alto Networks, Inc.

Option to Hold Web Requests During URL Category Lookup

Beginning with PAN-OS 8.1.10 (and PAN-OS 9.0.4), you have the flexibility to hold or allow initial web requests while the firewall checks the URL category for the destination websites.

When a user visits a website, a firewall with URL Filtering enabled checks its local cache of URL categories in an attempt to categorize the site. If the firewall doesn't find the category for that URL in the cache, that firewall performs a lookup in PAN-DB (the Palo Alto Networks URL database). By default, the firewall allows the web requests to pass through during this cloud lookup and enforces policy as needed after the server responds.

However, you can now choose to hold web requests (the firewall blocks the request) until the firewall either finds the URL category or the lookup times out. If the lookup times out, the firewall considers the URL category `not-resolved`.

Configure URL Filtering to hold web requests by following these steps:

STEP 1 | Access the CLI.

STEP 2 | Enter the command to enable the feature:

```
admin@PA-3060> set system setting ctd hold-client-request yes
```

Authentication Features

- Extensible Authentication Protocol (EAP) Support for RADIUS
- Authentication Using Custom Certificates for WildFire and PAN-DB

© 2020 Palo Alto Networks, Inc.

Extensible Authentication Protocol (EAP) Support for RADIUS

To securely transport administrator or end user credentials between RADIUS servers and the firewall, you can now use the following Extensible Authentication Protocols (EAP): PEAP-MSCHAPv2, PEAP with GTC, or EAP-TTLS with PAP.

The supported EAP methods create encrypted tunnels between the firewall and the RADIUS server to securely transmit usernames, passwords, and other credential information. The tunnel created by EAP consists of an inner tunnel and an outer tunnel. After the RADIUS server's certificate is validated, the firewall creates the outer tunnel using SSL. After the encrypted TLS outer tunnel has been established, the firewall creates the inner tunnel to transmit the user's credentials to the server.

To further protect user information from eavesdropping, you can mask the username by anonymizing the user's identity in the outer tunnel. For authentication using GlobalProtect, you can optionally allow GlobalProtect users with expired passwords to successfully log in by permitting them to change their password.

You can use EAP with RADIUS to:

- Secure administrative access to the web interface on the firewall and Panorama
- Secure administrative access using CLI on the firewall and Panorama
- Authenticate end users through Captive Portal and GlobalProtect, including Clientless VPN and GlobalProtect Gateway

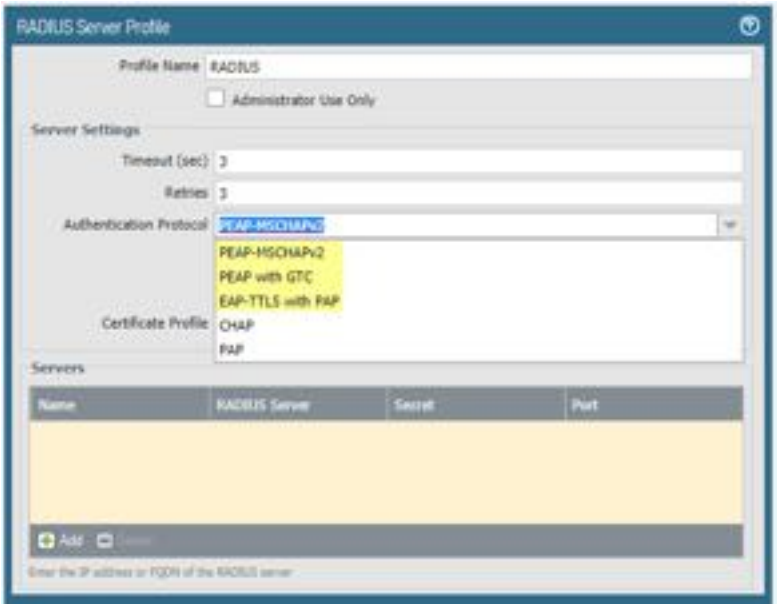
STEP 1 | Add a Certificate Profile to allow the RADIUS server to authenticate with the firewall.

1. Install the Server Certificate and Key on the RADIUS server.
2. Add the hostname or IP address of the firewall as the RADIUS client.
3. Upload the Root-CA and Intermediate CA certificates used to sign the RADIUS server certificate on the firewall.

STEP 2 | Create a RADIUS server profile that uses an EAP **Authentication Protocol**:

- **PEAP-MSCHAPv2** (Default)—Protected EAP (PEAP) with Microsoft Challenge-Handshake Authentication Protocol (MSCHAPv2) provides improved security over PAP or CHAP by transmitting both the username and password in an encrypted tunnel. In addition, if you are using GlobalProtect, you can allow GlobalProtect users to change expired passwords.
- **PEAP with GTC**—Select Protected EAP (PEAPv0) with Generic Token Card (GTC) to use one-time tokens in an encrypted tunnel.
- **EAP-TTLS with PAP**—Select EAP with Tunneled Transport Layer Security and PAP to transport plaintext credentials for PAP in an encrypted tunnel.

© 2020 Palo Alto Networks, Inc.



The Auto option is no longer supported. For more information on changes to default behavior for this feature, see Changes to Default Behavior and Upgrade/Downgrade Considerations in the Release Notes.

STEP 3 | (Optional) Select whether GlobalProtect users can change expired passwords:

- **Allow users to change passwords after expiry** (PEAP-MSCHAPv2 with GlobalProtect only)—Select this option to allow GlobalProtect users to change expired passwords.

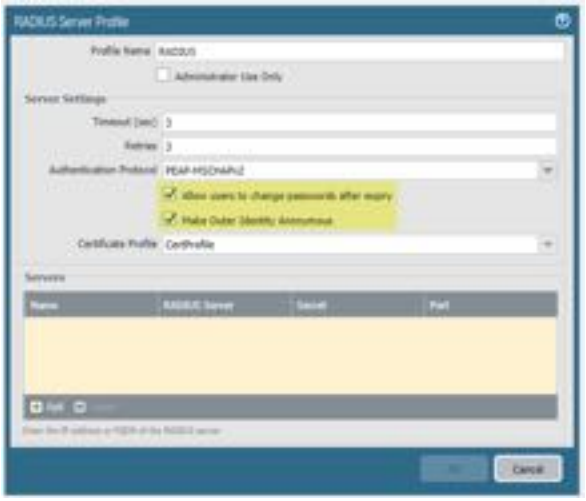
  This feature is only supported with GlobalProtect client 4.1 or later.

STEP 4 | Select whether the user's identity is anonymous in the outer tunnel that the firewall creates to authenticate with the server:

- **Make Outer Identity Anonymous**—This option is enabled by default to anonymize the user's identity in the outer tunnel, which is created after authenticating with the server.

  You must configure the RADIUS server to allow access for anonymous users. Some RADIUS server configurations may not support anonymous outer IDs, and you may need to clear the option. When cleared, usernames are transmitted using cleartext.

© 2020 Palo Alto Networks, Inc.



STEP 5 | Select the **Certificate Profile** that the server uses to authenticate the firewall.

STEP 6 | **Add** each RADIUS server.

STEP 7 | Select **Device** > **Authentication Profile**, **Add** a profile, and assign the RADIUS server profile to the authentication profile.

STEP 8 | To enable administrator access, select **Device** > **Setup** > **Management**, and select the Authentication Profile. Configure an Admin Role profile if the administrator uses a custom role and configure an access domain if the firewall has more than one virtual system.

STEP 9 | If you are using GlobalProtect, see Enable Authentication Using an Authentication Profile.

STEP 10 | If you are using Captive Portal, see Map IP Addresses to Usernames Using Captive Portal.

STEP 11 | **Commit** your changes to activate them on the firewall.

To confirm the authentication is successful, use the **`test authentication authentication-profile <auth-profile-name> username <username> password`** command. You can also use the **Description** column in **Monitor** > **Authentication** to view the EAP outer and inner identities. To troubleshoot, use **`grep`** to search the logs for the number in the **Authentication ID** column. For more information, refer to the PAN-OS Admin Guide.

© 2020 Palo Alto Networks, Inc.

Authentication Using Custom Certificates for WildFire, PAN-DB, and Between Log Collectors

You can now configure mutual authentication between WildFire® or PAN-DB appliances and other Palo Alto Networks appliances using custom certificates. This allows you to establish a unique chain of trust between WildFire or PAN-DB and connected Palo Alto Networks appliances instead of relying on predefined certificates for inter-device communication. You can generate these certificates locally on Panorama or the firewall, obtain them from a trusted third-party certificate authority (CA), or obtain certificates from your own enterprise CA.

Additionally, you can configure mutual authentication between Log Collectors in a Collector Group. Local Log Collectors use the same client and server certificates as Panorama.

You can configure the client certificate and certificate profile on each client device or push the configuration from Panorama to each device as part of a template.



This feature is an extension of Authentication Using Custom Certificates *introduced in PAN-OS 8.0.*

- Custom certificates for a standalone WildFire appliance or a PAN-DB appliance—You can deploy custom certificates between a WildFire or PAN-DN appliance that receives samples or URL information from a firewall. In this case, the firewall acts as the client and the WildFire appliance or PAN-DB appliance acts as the server. Use the CLI to deploy custom certificates directly on the WildFire appliance or PAN-DB appliance.
- Custom certificate for a WildFire appliance managed by Panorama—This allows you to configure custom certificate communication between WildFire appliance that receives samples from firewalls. In this deployment, the WildFire appliance is the server and the firewalls are the clients. However, you can complete the server and client communication configuration through the Panorama web interface instead of the WildFire appliance CLI.
- Custom certificates for a WildFire appliance as a client—You can deploy custom certificates for the communication channel Panorama uses to push configuration information to a WildFire appliance. In this deployment, Panorama is the server and the WildFire appliance is the client.
- Single custom certificate for a WildFire Cluster—Instead of assigning unique certificates to each WildFire appliance in a cluster, you can assign a single, shared client certificate to the entire WildFire cluster, which, in turn, allows you to push a single certificate to all WildFire appliances in the cluster instead of configuring separate certificates for each cluster member. In this scenario, Panorama is the server and the WildFire cluster is the client.
- Custom certificates for communication between Log Collectors—This allows you to configure custom server and client certificates for inter-Log Collector communication. You must configure secure server communication and secure client communication on each Log Collector in a Collector Group because the server and client roles are chosen dynamically.

The following procedure provides a high-level overview of the steps involved in deploying custom certificates on your WildFire appliance, PAN-DB appliance, or Log Collectors.

STEP 1 | Generate or obtain your server and client certificates.

© 2020 Palo Alto Networks, Inc.

Based in the needs of your organization, choose one of the supported methods for generating or obtaining your custom certificates.

STEP 2 | Configure the server certificate profile and SSL/TLS service profile for the server.

STEP 3 | Configure Secure Server Communication on the server.

STEP 4 | Configure the client certificate profile for the client device. The method for configuring this profile depends on your deployment.

STEP 5 | Configure Secure Client Communication on the client devices.

STEP 6 | Enforce the use of custom certificates.

© 2020 Palo Alto Networks, Inc.

GlobalProtect Features

- Optimized Split Tunneling for GlobalProtect
- Kerberos Authentication Support for macOS
- SAML SSO for GlobalProtect on Chromebooks
- OPSWAT SDK V4 Support
- GlobalProtect Tunnel Preservation on User Logout
- Automatic Launching of Web Browser in Captive Portal Environment

For more information about the features introduced in GlobalProtect app 4.1, including GlobalProtect App for Linux, see the GlobalProtect App 4.1 New Features Guide.

© 2020 Palo Alto Networks, Inc.

Optimized Split Tunneling for GlobalProtect

Software Support: Starting with GlobalProtect™ App 4.1 and with PAN-OS® 8.1 and later releases

OS Support: Windows 7 Service Pack 2 and later releases and macOS 10.10 and later releases

In addition to route-based split tunneling, the GlobalProtect app for Windows and macOS endpoints now supports split tunneling based on destination domain, client process, and HTTP/HTTPS video streaming application.

This enhancement requires a GlobalProtect subscription.

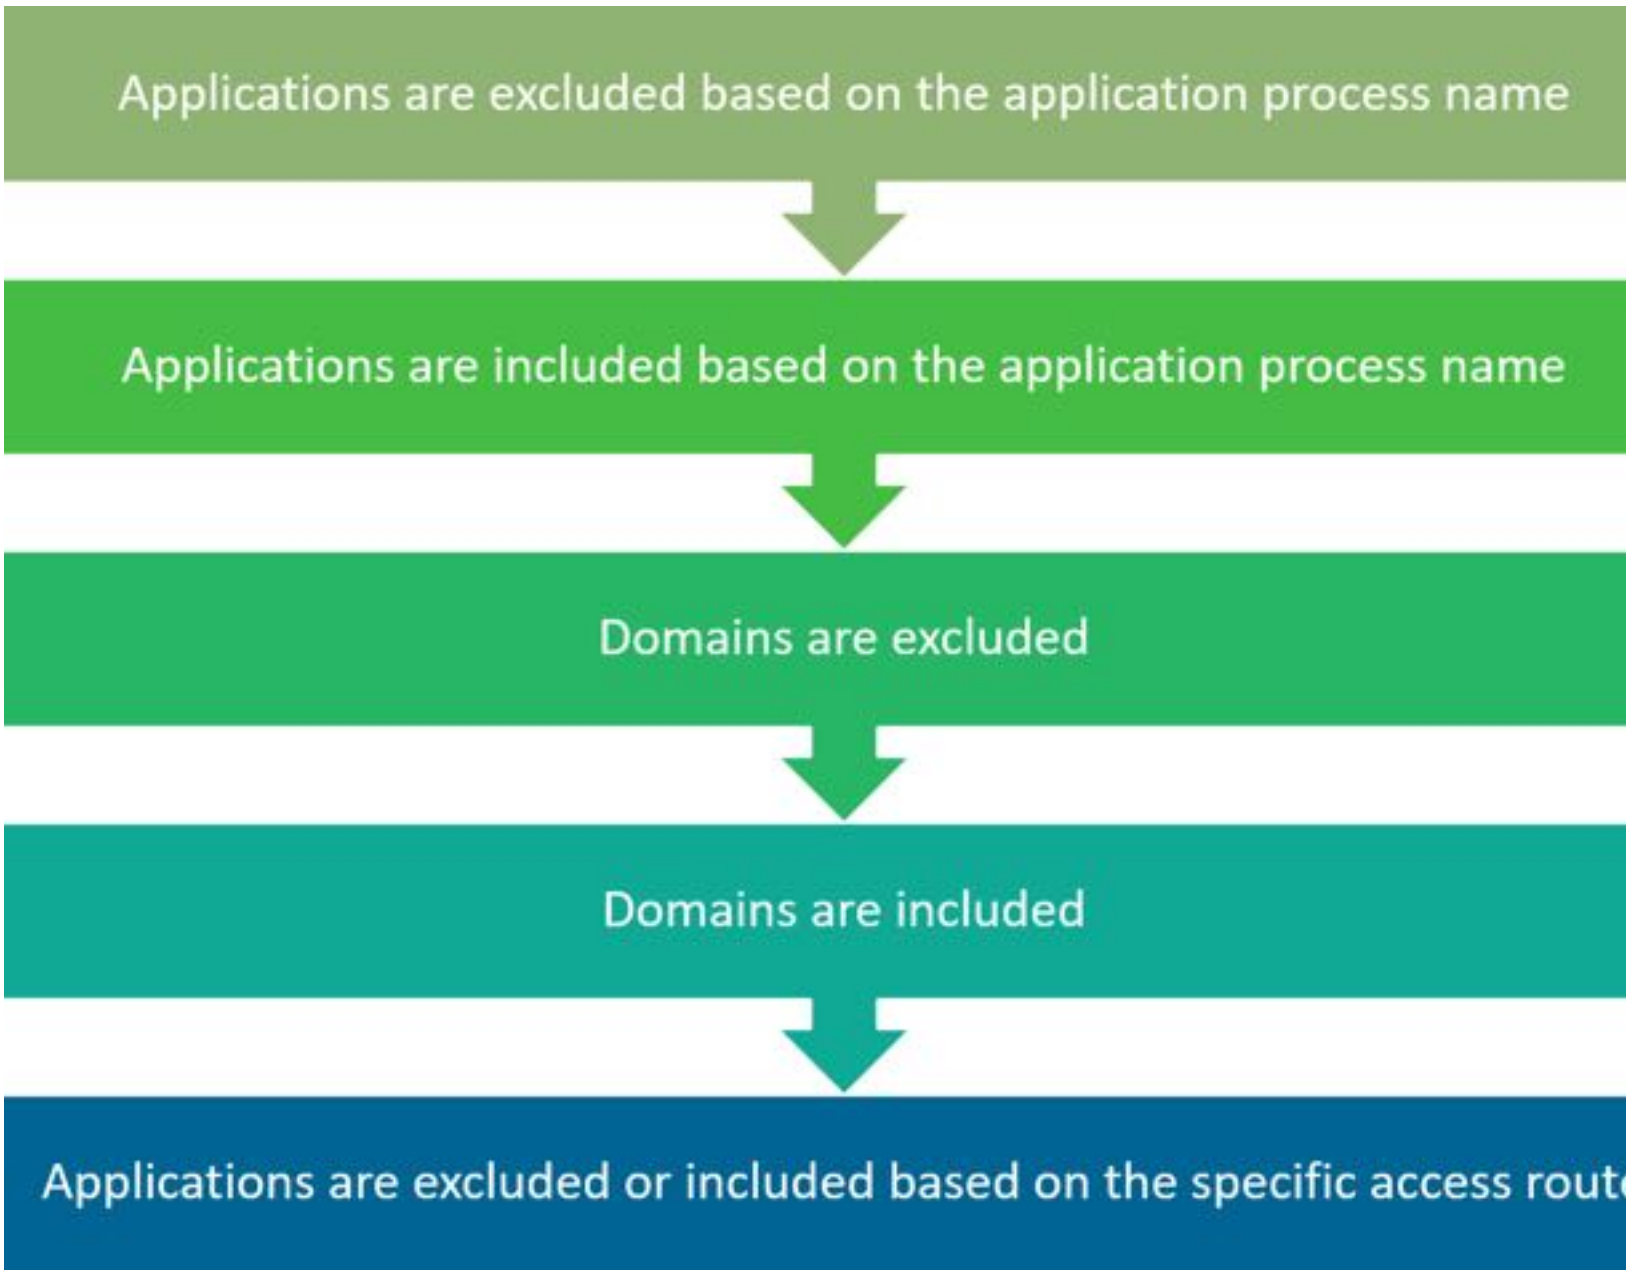
This enhancement enables you to:

- Tunnel enterprise SaaS and public cloud applications for comprehensive SaaS application visibility and control to avoid risks associated with Shadow IT in environments where it is not feasible to tunnel all traffic.
- Send latency-sensitive traffic, such as VoIP, outside the VPN tunnel, while all other traffic goes through the VPN for inspection and policy enforcement by the GlobalProtect gateway.
- Exclude HTTP/HTTPS video streaming traffic from the VPN tunnel. Video streaming applications, such as YouTube and Netflix, consume large amounts of bandwidth. By excluding lower risk video streaming traffic from the VPN tunnel, you can decrease bandwidth consumption on the gateway.

The firewall App-ID functionality identifies the video stream before allowing traffic to be split tunneled.

The following list describes the order in which the split tunnel rules are applied:

© 2020 Palo Alto Networks, Inc.



When you configure a split tunnel to include traffic based on the application process name or destination domain and port (optional), all traffic for that specific application or domain is sent through the VPN tunnel for inspection and policy enforcement. For example, you can allow all Salesforce traffic to go through the VPN tunnel using the ***Salesforce.com** destination domain. By including all Salesforce traffic in the VPN tunnel, you can provide secure access to the entire Salesforce domain and subdomains.

When you configure a split tunnel to exclude traffic based on the application process name or destination domain and port (optional), all traffic for that specific application or domain is sent directly to the physical adapter on the endpoint without inspection. For example, you can exclude all Skype traffic from the VPN tunnel using the **C:\Program Files (x86)\Skype\Phone\Skype** application process name.

Use the following steps to configure a split tunnel for public applications or video streams:

- Configure a split tunnel to include or exclude public applications based on the destination domain:

© 2020 Palo Alto Networks, Inc.

1. Configure a GlobalProtect gateway.

   Select **Network** > **GlobalProtect** > **Gateways** to modify an existing gateway or **Add** a new one.
2. Enable split tunneling.
   1. On the **Agent** > **Tunnel Settings** tab, enable **Tunnel Mode** to enable split tunneling.
   2. Configure the tunnel parameters for the GlobalProtect app.
3. Configure a split tunnel to include or exclude SaaS or public cloud applications based on the destination domain and port (optional).

   This feature supports both IPv4 and IPv6 traffic.

   1. On the **Agent** > **Client Settings** tab, select an existing client setting or **Add** a new one.
   2. Disable the **No direct access to local network** option (**Split Tunnel** > **Access Route**). If enabled, this setting disables split tunneling on Windows, Linux, and macOS networks.
   3. (Optional) **Add** the SaaS or public cloud applications that you want to route to GlobalProtect through the VPN connection using the destination domain and port (**Split Tunnel** > **Domain and Application** > **Include Domain**). You can add up to 200 entries to the list. For example, add **`*.office365.com`** to allow all Office 365 traffic to go through the VPN tunnel.
   4. (Optional) **Add** the SaaS or public cloud applications that you want to exclude from the VPN tunnel using the destination domain and port (**Split Tunnel** > **Domain and Application** > **Exclude Domain**). You can add up to 200 entries to the list. For example, add **`*.engadget.com`** to exclude all Engadget traffic from the VPN tunnel.
   5. Click **OK** to save your client settings.
4. Save the gateway configuration.
   1. Click **OK** to save the gateway configuration.
   2. **Commit** your changes.

- Configure a split tunnel to include or exclude public applications based on the application process name:
  1. Configure a GlobalProtect gateway.

     Select **Network** > **GlobalProtect** > **Gateways** to modify an existing gateway or add a new one.
  2. Enable split tunneling.
     1. On the **Agent** > **Tunnel Settings** tab, enable **Tunnel Mode** to enable split tunneling.
     2. Configure the tunnel parameters for the GlobalProtect app.
  3. Configure a split tunnel to include or exclude SaaS or public cloud applications based on the application process name.

     This feature supports both IPv4 and IPv6 traffic.

     1. On the **Agent** > **Client Settings** tab, select an existing client setting or **Add** a new one.
     2. Disable the **No direct access to local network** option (**Split Tunnel** > **Access Route**). This setting disables split tunneling on Windows, Linux, and macOS networks.
     3. (Optional) **Add** the SaaS or public cloud applications that you want to route to GlobalProtect through the VPN connection using the application process name (**Split Tunnel** > **Domain and Application** > **Include Client Application Process Name**. You can add up to 200 entries to the list. For example, add **`/Application/Safari.app/Contents/MacOS/Safari`** to allow all Safari-based traffic to go through the VPN tunnel on macOS endpoints.

© 2020 Palo Alto Networks, Inc.

4. (Optional) **Add** the SaaS or public cloud applications that you want to exclude from the VPN tunnel using the application process name (**Split Tunnel** > **Domain and Application** > **Exclude Client Application Process Name**). You can add up to 200 entries to the list. For example, add **`/Applications/Microsoft Lync.app/Contents/MacOS/MicrosoftLync`** to exclude all Microsoft Lync application traffic from the VPN tunnel.
5. Click **OK** to save your client settings.

4. Save the gateway configuration.
   1. Click **OK** to save the gateway configuration.
   2. **Commit** your changes.

- Configure a split tunnel to exclude video streaming traffic:
  1. Configure a GlobalProtect gateway.

     Select **Network** > **GlobalProtect** > **Gateways** to modify an existing gateway or add a new one.
  2. Enable split tunneling.
     1. On the **Agent** > **Tunnel Settings** tab, enable **Tunnel Mode** to enable split tunneling.
     2. Configure the tunnel parameters for the GlobalProtect app.
  3. Configure a split tunnel to exclude video streaming traffic from the VPN tunnel.

     All video traffic types are redirected for the following video streaming applications:

     - YouTube
     - Dailymotion
     - Netflix

     If you exclude any other video streaming applications from the VPN tunnel, only the following video traffic types are redirected for those applications:

     - MP4
     - WebM
     - MPEG

     The App-ID functionality on the firewall identifies the video stream before traffic can be split tunneled.

     If the physical adapter on a Windows or macOS endpoint supports only IPv4 addresses, the endpoint user cannot access the video streaming applications that you exclude from the VPN tunnel when you configure the GlobalProtect gateway to assign IPv6 addresses to the virtual network adapters on the endpoints that connect to the gateway. In this case, ensure that the IP pools used to assign IP addresses to the virtual network adapters on these endpoints do not include any IPv6 addresses (Network > GlobalProtect > Gateways > Agent > Client IP Pool or Client Settings > IP Pools).

     If you exclude video streaming traffic from the VPN tunnel (Network > GlobalProtect > Gateways > <gateway-config> > Agent > Video Traffic), do not include web browser applications, such as Firefox or Chrome, in the VPN tunnel (Network > GlobalProtect > Gateways > <gateway-config> > Agent > Client Settings > <client-setting> > Split Tunnel > Domain and Application). This ensures that there is no conflicting logic in the split tunnel configuration and that your users can stream videos from web browsers.

     To exclude Sling TV app traffic from the VPN tunnel, use application-based split tunneling (Network > GlobalProtect > Gateways > <gateway-config> > Agent > Client

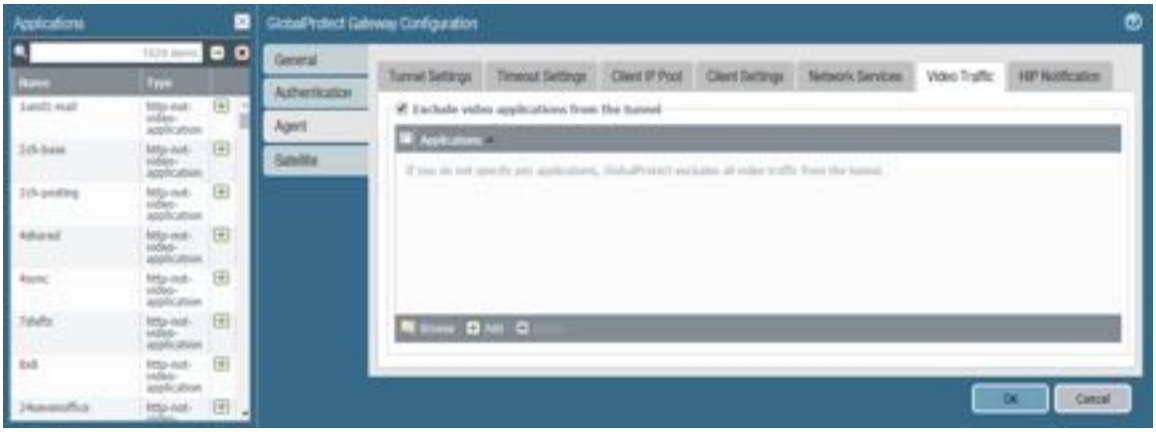
© 2020 Palo Alto Networks, Inc.

Settings > <client-setting-config> > Split Tunnel > Domain and Application > Exclude Client Application Process Name).

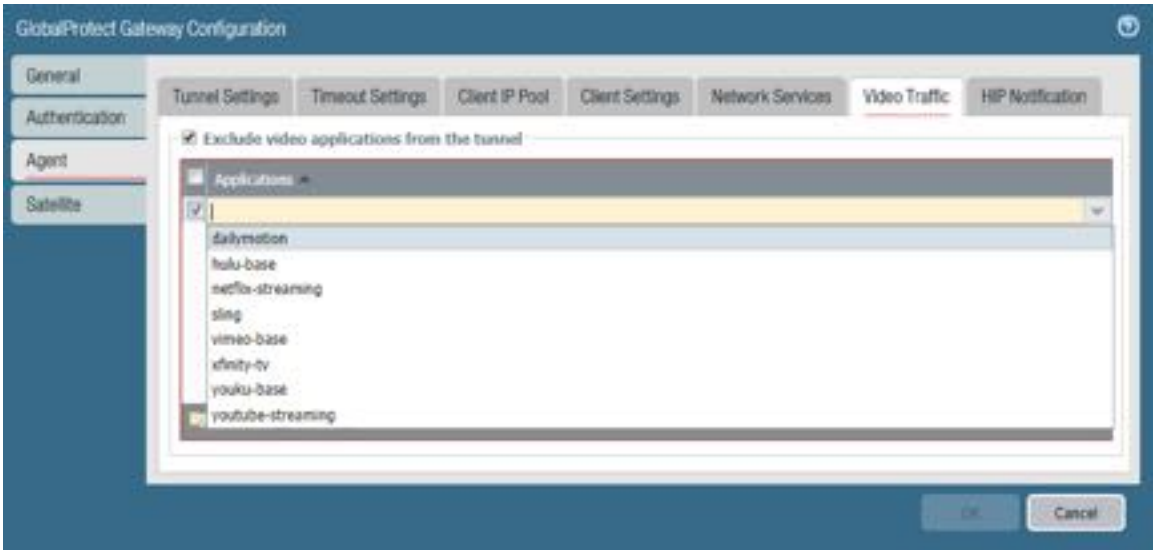
1. On the **Agent** > **Video Traffic** tab, enable the option to **Exclude video applications from the tunnel**.

   If you enable this option but do not select specific video streaming applications to exclude from the VPN tunnel, all video streaming traffic is excluded.

2. (Optional) **Browse** the **Applications** list to view all of the video streaming applications that you can exclude from the VPN tunnel. Click the add icon ( ) for the application(s) that you want to exclude. For example, click the add icon for **directv** to exclude DIRECTV video streaming traffic from the VPN tunnel.



3. (Optional) **Add** the video streaming applications that you want to exclude from the VPN tunnel using the **Applications** drop-down—a shortened version of the **Applications** list that contains some of the most popular video streaming applications. For example, select **youtube-streaming** from the **Applications** drop-down to exclude all YouTube-based video streaming traffic from the VPN tunnel.



4. Save the gateway configuration.
   1. Click **OK** to save the gateway configuration.
   2. **Commit** your changes.

© 2020 Palo Alto Networks, Inc.

Kerberos Authentication Support for macOS

Software Support: Starting with GlobalProtect™ App 4.1 and with PAN-OS® 8.0 and later releases

OS Support: macOS 10.10 and later releases

The GlobalProtect™ app for Mac endpoints now supports Kerberos V5 single sign-on (SSO) for GlobalProtect portal and gateway authentication. Kerberos SSO maintains a seamless logon experience by providing accurate User-ID™ information without user interaction. Networks that support Kerberos SSO require end users to log in only during initial network access. After the initial login, end users can access any Kerberos-enabled service in the network (such as webmail) without having to log in again until the SSO session expires (the SSO session duration is established by the Kerberos administrator). This authentication method helps identify users for user and HIP policy enforcement.

If you enable both Kerberos SSO and an external authentication service (such as RADIUS), GlobalProtect attempts SSO first. You can configure GlobalProtect to fall back to an external authentication service when SSO fails or you can configure GlobalProtect to use only Kerberos SSO for authentication.

In this implementation, the GlobalProtect portal and gateway act as Kerberos service principals and the GlobalProtect app acts as a user principal that authenticates end users with a Kerberos service ticket from the Key Distribution Center (KDC).

The following items must be in place for the GlobalProtect app for macOS endpoints to support Kerberos SSO:

- A Kerberos infrastructure, which includes a KDC with an authentication server (AS) and a ticket-granting service (TGS).

  GlobalProtect supports the following KDCs:

  - Microsoft Active Directory on Windows Server 2008 R2
  - Microsoft Active Directory on Windows Server 2012
  - MIT Kerberos V5

  The KDC must be reachable from the endpoint on which the Globalprotect app is running. In most instances, the KDC is reachable only from inside the enterprise network, which means the GlobalProtect app can use Kerberos authentication only when the endpoint is internal. However, if the KDC is reachable from outside the enterprise network (from the Internet), the GlobalProtect app can use Kerberos authentication when the endpoint is external.

  If the user certificate store contains at least one certificate that is issued by the same CA as the certificate used for pre-logon tunnel establishment, you can also use Kerberos authentication with pre-logon to enable the GlobalProtect app to use Kerberos authentication when the endpoint is external.

  When an end user attempts to access protected network resources using Kerberos authentication, the AS grants the user a Ticket to Get Tickets (TGT), which is a service request used to generate service tickets from the TGS. The service ticket is then used to authenticate the end user and establish a service session.
- A Kerberos service account for each GlobalProtect portal and gateway.

  Service accounts are required for creating Kerberos keytabs, which are files that contain the principal name and password of each GlobalProtect portal or gateway.

© 2020 Palo Alto Networks, Inc.

- PAN-OS 8.0 or a later release.
- macOS version 10.10 or a later release.

STEP 1 | Create a Kerberos keytab file.

1. Log in to the KDC using your Kerberos service account credentials.
2. Open a command prompt and then enter the following command:

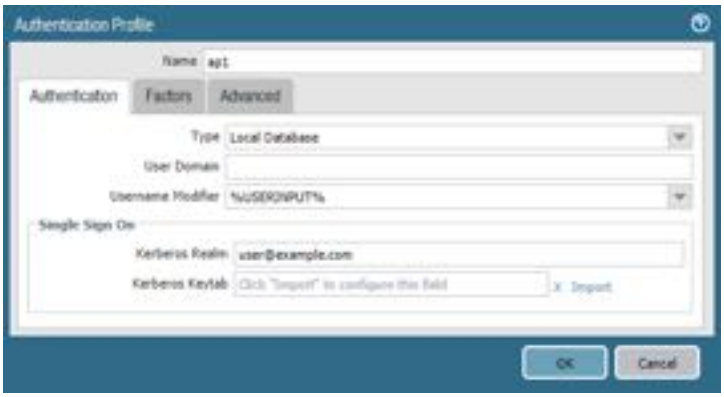
```
ktpass /princ <principal_name> /pass <password> /crypto <algorithm> /ptype
 KRB5_NT_PRINCIPAL /out <file_name>.keytab
```

The <principal_name> and <password> are the principal name and password of the GlobalProtect portal or gateway. The <algorithm> must match the algorithm in the service ticket issued by the TGS, which is determined by the Kerberos administrator. If the GlobalProtect portal or gateway is running in FIPS or CC mode, the algorithm must be aes128-cts-hmac-sha1-96 or aes256-cts-hmac-sha1-96. If the portal or gateway is not running in FIPS or CC mode, you can also use des3-cbc-sha1 or arcfour-hmac.

STEP 2 | Create a server profile for Kerberos authentication.

STEP 3 | Import the Kerberos keytab file to an authentication profile.

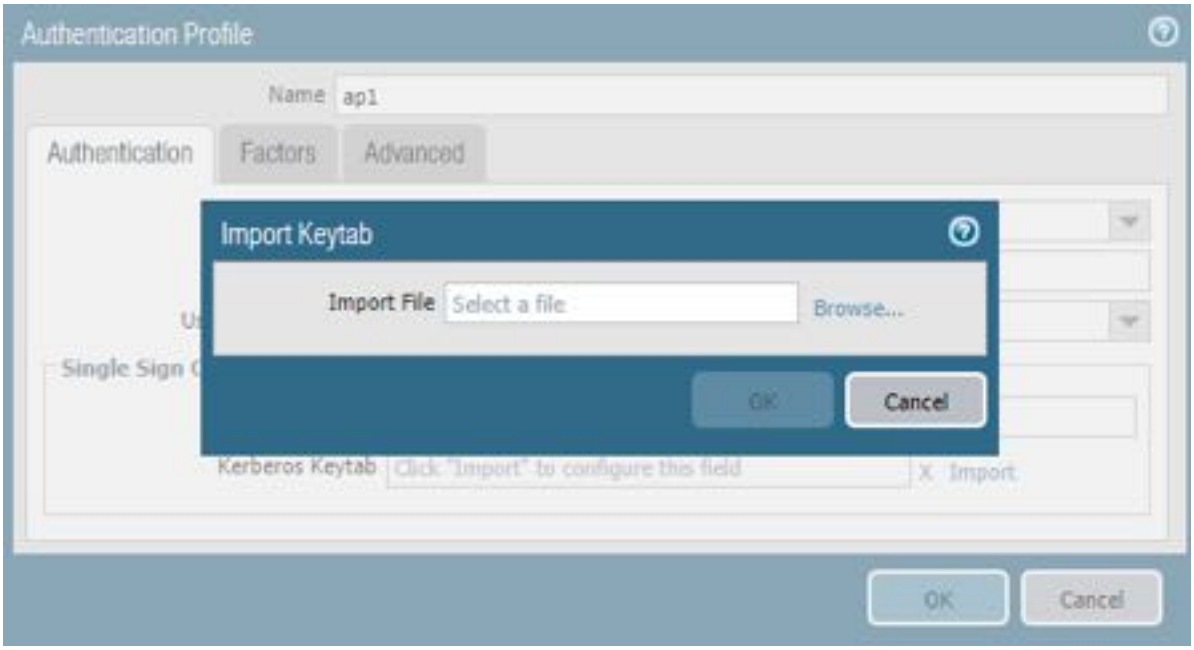
1. Select **Device** > **Authentication Profile**.
2. Select an existing authentication profile or **Add** a new one.



3. In the **Single Sign-On** area, enter the **Kerberos Realm** (up to 127 characters), which is the domain to which the end user belongs. For example, a user with the account name user@EXAMPLE.LOCAL belongs to the EXAMPLE.LOCAL realm.
4. **Import** a **Kerberos Keytab** file to the authentication profile.

   When the Import Keytab dialog opens, **Browse** to and select the keytab file, and then click **OK**.

© 2020 Palo Alto Networks, Inc.



5. Click **OK** to save your changes.

STEP 4 | Assign the authentication profile to an internal gateway. If the Kerberos authentication infrastructure is deployed in an external gateway, such as a DMZ, you can also assign the authentication profile to an external gateway.

1. Select **Network** > **GlobalProtect** > **Gateways** to modify an existing gateway or **Add** a new one.



2. Select an existing **SSL/TLS Service Profile** for securing the gateway, or **Add** a new service profile (**Network** > **GlobalProtect** > **Gateways** > *<gateway-config>* > **Authentication**).
3. **Add** a **Client Authentication** configuration (**Network** > **GlobalProtect** > **Gateways** > *<gateway-config>* > **Authentication**), and then configure the following settings:
   - **Name**—Name of the client authentication configuration.

© 2020 Palo Alto Networks, Inc.

- **OS**—Operating systems on which the gateway can be accessed.
- **Authentication Profile**—Authentication profile to which your Kerberos keytab file was imported.
- (Optional) **Username Label**—Custom username label for GlobalProtect gateway login.
- (Optional) **Password Label**—Custom password label for GlobalProtect gateway login.
- (Optional) **Authentication Message**—Message that is displayed when end users authenticate to the gateway.

4. Click **OK** to save your changes.

STEP 5 | Assign the authentication profile to the GlobalProtect portal.

1. Select **Network** > **GlobalProtect** > **Portals**.
2. Select an existing portal or **Add** a new one.



3. Select an existing **SSL/TLS Service Profile** for securing the portal, or **Add** a new service profile (**Network** > **GlobalProtect** > **Portals** > *<portal-config>* > **Authentication**).
4. **Add** a **Client Authentication** configuration (**Network** > **GlobalProtect** > **Portals** > *<portal-config>* > **Authentication**), and then configure the following settings:

   - **Name**—Name of the client authentication configuration.
   - **OS**—Operating systems on which the portal can be accessed.
   - **Authentication Profile**—Authentication profile to which your Kerberos keytab file is imported.
   - (Optional) **Username Label**—Custom username label for GlobalProtect portal login.
   - (Optional) **Password Label**—Custom password label for GlobalProtect portal login.
   - (Optional) **Authentication Message**—Message that is displayed when end users log in to the portal.

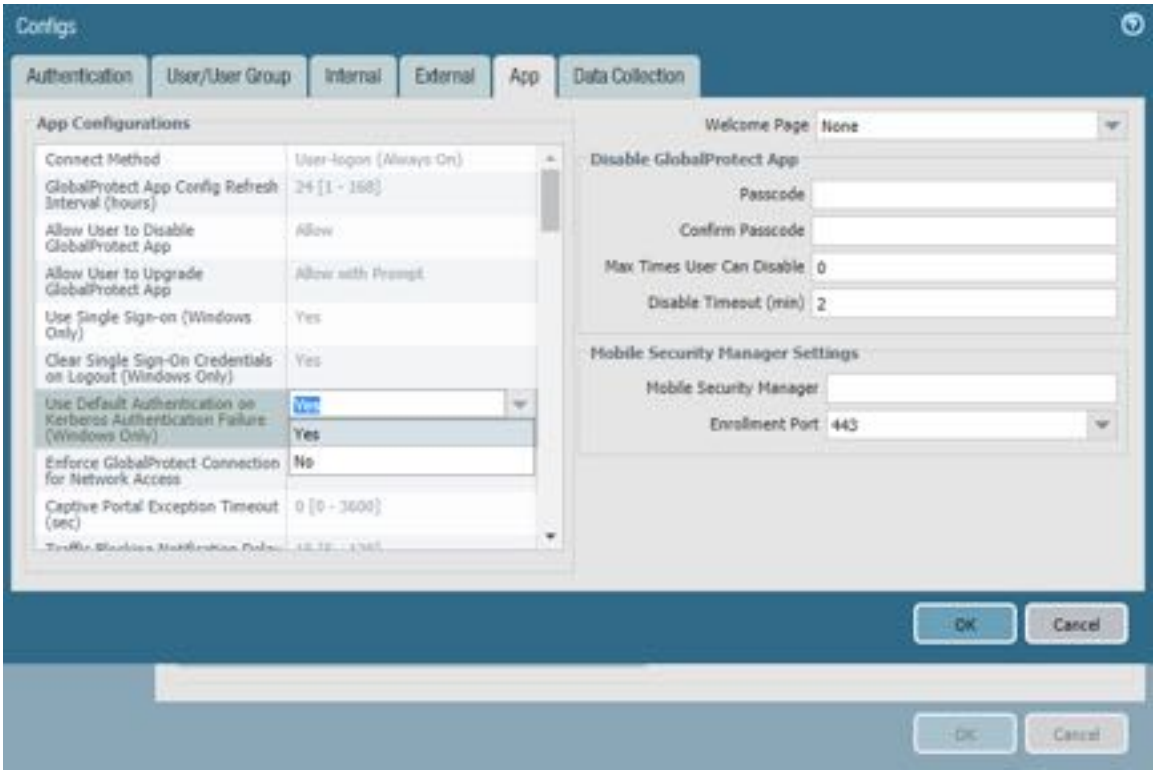
5. Click **OK** to save your changes.

STEP 6 | Configure the GlobalProtect app behavior for Kerberos authentication failure.

1. Select **Network** > **GlobalProtect** > **Portals**.
2. Select a portal configuration.
3. Select the agent configuration that you want to modify, or **Add** a new one (**Network** > **GlobalProtect** > **Portals** > *<portal-config>* > **Agent**).

© 2020 Palo Alto Networks, Inc.

4. In the **App Configurations** area (**Network** > **GlobalProtect** > **Portals** > *<portal-config>* > **Agent** > *<agent-config>* > **App**), select one of the following options for **Use Default Authentication on Kerberos Authentication Failure**:
   - **Yes**—Enables authentication to fall back so that when Kerberos authentication fails, GlobalProtect authenticates users through the default authentication method.
   - **No**—GlobalProtect can use only Kerberos to authenticate users.



5. Click **OK** to save your changes.
6. Click **OK** to complete your configuration.

STEP 7 | **Commit** the configuration.

© 2020 Palo Alto Networks, Inc.

SAML SSO for GlobalProtect on Chromebooks

Software Support: Starting with GlobalProtect™ App 4.1 and with PAN-OS® 8.0 and later releases

OS Support: Google Chrome OS 45 and later releases

The GlobalProtect app for Chromebooks (Chrome OS) now supports Security Assertion Markup Language (SAML) single sign-on (SSO). If you configure SAML as the authentication standard for Chromebooks, end users authenticate to GlobalProtect by leveraging the same login they use to access their Chromebook applications. This feature enables end users to connect to GlobalProtect automatically without having to re-enter their credentials on the GlobalProtect app.

In this implementation, Google acts as the SAML service provider while the GlobalProtect app authenticates users directly to their organization's SAML identity provider.



*GlobalProtect currently supports only the **Post** SAML HTTP binding method.*

To use SAML SSO for Chrome applications, end users must install the **SAML SSO for Chrome Apps** Google extension on their Chromebooks. This extension allows users to access multiple Chrome applications during a single login session by sending the user's SAML SSO cookies to all applications that are whitelisted by an administrator.

Use the following steps to configure SAML SSO for Chrome applications:

STEP 1 | Set up SAML authentication for GlobalProtect.

- Create a server profile with settings for access to the SAML authentication service.
- Create an authentication profile that refers to the SAML server profile.

STEP 2 | Configure a GlobalProtect gateway.

1. Specify a SAML authentication profile for Chrome gateway users.
   - On the **Authentication** tab of the GlobalProtect gateway configuration, select a SAML **Authentication Profile** or create a new SAML profile for the gateway. This profile is used to authenticate endpoints seeking access to the gateway.
2. (Optional) Select a **Certificate Profile** for client gateway authentication. The client certificate must be pre-deployed or deployed using the Simple Certificate Enrollment Protocol (SCEP).

STEP 3 | Define the GlobalProtect client authentication configurations on the GlobalProtect portal.

1. Specify a SAML authentication profile for the portal.
   - On the **Authentication** tab of the GlobalProtect portal configuration, select a SAML **Authentication Profile** or create a new SAML profile for the portal. This profile is used to authenticate endpoints seeking access to the portal.
2. (Optional) Select a **Certificate Profile** for client portal authentication. A valid client certificate must be pre-deployed on all Chromebooks if you configure the **Certificate Profile**.

STEP 4 | Install the **SAML SSO for Chrome Apps** extension from Google. This extension enables SAML SSO for Chrome applications.

Launch the Chrome **Web Store** and install the **SAML SSO for Chrome Apps** extension.

STEP 5 | Configure SAML SSO for Chrome apps.

© 2020 Palo Alto Networks, Inc.

- For GlobalProtect to support SAML SSO, you must add the GlobalProtect application ID (**nicidmbokaedpmoegdbcebhnchpegcdc**) to the whitelist in the **SAML SSO for Chrome Apps** extension configuration file.

© 2020 Palo Alto Networks, Inc.

OPSWAT SDK V4 Support

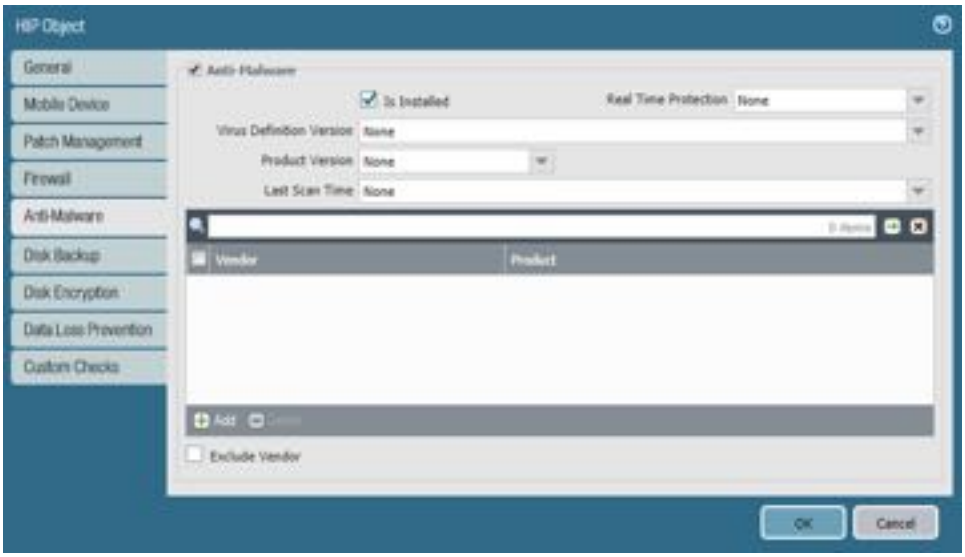
Software Support: Starting with GlobalProtect™ app 4.1 and with PAN-OS® 8.1 and later releases.

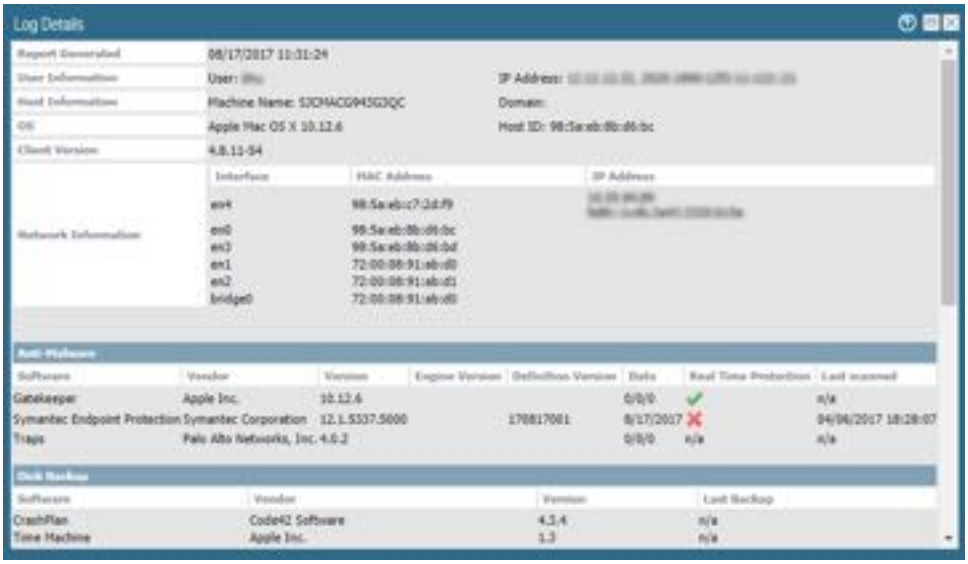
OS Support: Windows 7 and later releases and macOS 10.12 and later releases.

GlobalProtect is now integrated with OPSWAT SDK V4 to detect and assess the endpoint state and the third-party security applications running on the endpoint. OPSWAT is a security tool leveraged by the Host Information Profile (HIP) to collect information about the security status of the endpoints in the network, which is used for policy enforcement on the GlobalProtect gateway. This integration follows the end-of-life (EoL) announcement for OPSWAT SDK V3, which is the OPSWAT SDK version that GlobalProtect previously supported.

With this migration from OPSWAT SDK V3 to OPSWAT SDK V4, GlobalProtect introduces the following changes to the HIP Match log and HIP Object:

- The **Antivirus** and **Anti-Spyware** categories merged to form the new **Anti-Malware** category.
- Vendor and product names are based on OPSWAT SDK V4.





© 2020 Palo Alto Networks, Inc.

OPSWAT SDK is unable to detect the following Anti-Malware information for the Gatekeeper security feature on macOS endpoints:

- *Engine Version*
- *Definition Version*
- *Date*
- *Last Scanned*

© 2020 Palo Alto Networks, Inc.

GlobalProtect Tunnel Preservation on User Logout

Software Support: PAN-OS® 8.1 and later releases

OS Support: Windows 10



This feature requires Content Release version _____ and later.

You can now configure GlobalProtect to preserve the existing VPN tunnel when users log out of their endpoint. With this enhancement, you can specify a logoff timeout period to indicate the amount of time for which the GlobalProtect session remains active during user logout. For example, if certain updates (such as Group Policy Object updates) require remote endpoints to be connected to the network and users to log out then log back in to the endpoint, you can configure GlobalProtect to preserve the VPN tunnel for a specified period of time after user logout.

Consider the following GlobalProtect connection behaviors when you configure GlobalProtect to preserve the VPN tunnel:

- If the same user logs out and then logs back in to an endpoint within the specified timeout period in either Always On or On-Demand mode, GlobalProtect remains connected without requiring any user interaction (including portal and gateway authentication). If the user does not log back in within the specified timeout period, the tunnel disconnects and he or she must reestablish the GlobalProtect connection.
- If a user logs out of an endpoint and then a different user logs in to the same endpoint in either Always On or On-Demand mode, the existing tunnel is renamed for the new user only if the new user authenticates to GlobalProtect successfully within the specified timeout period. If the new user does not log in and authenticate successfully within the specified timeout period, the existing tunnel disconnects and a new GlobalProtect connection must be established. If the new user is in Always On mode, GlobalProtect attempts to establish a new connection automatically. If the new user is in On-Demand mode, he or she must establish a new GlobalProtect connection manually.

Use the following steps to configure GlobalProtect to preserve the VPN tunnel following user log out:

STEP 1 | Configure a GlobalProtect gateway.

STEP 2 | Enable tunneling and then configure the tunnel parameters.

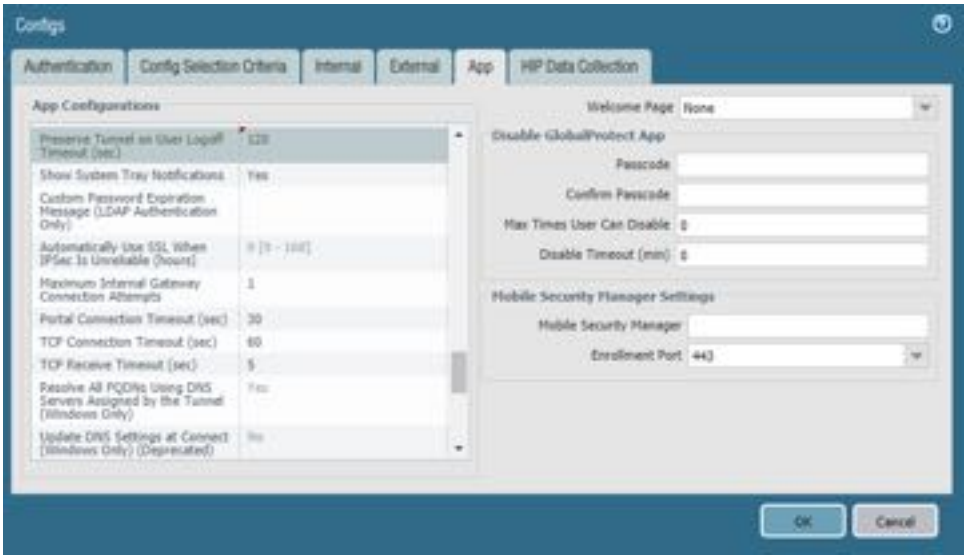
STEP 3 | Set up access to the GlobalProtect portal.

STEP 4 | Define the GlobalProtect agent configurations.

STEP 5 | Customize the GlobalProtect app.

- Specify a **Preserve Tunnel on User Logoff Timeout** value (range is 0 to 600 seconds; default is 0 seconds). This value indicates the amount of time during which GlobalProtect preserves the VPN tunnel after users log out of their endpoint. If you accept the default value of **0**, GlobalProtect does not preserve the tunnel following user logout.

© 2020 Palo Alto Networks, Inc.



STEP 6 | **Commit** your changes.

© 2020 Palo Alto Networks, Inc.

Automatic Launching of Web Browser in Captive Portal Environment

Software Support: PAN-OS® 8.1 and later releases

OS Support: Windows 7 and later releases and macOS 10.11 and later releases

This feature requires Content Release version _____ and later.

You can now configure GlobalProtect to launch your default web browser automatically upon captive portal detection so that users can log in to the captive portal seamlessly. With this enhancement, you can specify the URL of the website that you want to use for the initial connection attempt that initiates web traffic when the default web browser launches. The captive portal then intercepts this website connection attempt and redirects the default web browser to the captive portal login page. If a user logs in to the captive portal successfully and the internet becomes reachable within the configured **Captive Portal Exception Timeout period** (**Network** > **GlobalProtect** > **Portals** > *<portal-config>* > **Agent** > *<agent-config>* > **App**), GlobalProtect establishes a connection and permits traffic on the network.

Redirects to the captive portal login page are dependent on the captive portal solution.

Depending on the endpoint operating system (OS), some endpoints continue to retry the OS-based captive portal check during the configured Captive Portal Exception Timeout period. If both the GlobalProtect and OS-based captive portal checks are successful, the endpoint displays both the browser-based captive portal login page and OS-based captive portal assistant. Palo Alto Networks does not have control over the OS behavior.

Use the following steps to configure GlobalProtect to launch your default web browser automatically upon captive portal detection:

STEP 1 | Configure a GlobalProtect gateway.

STEP 2 | Enable tunneling and then configure the tunnel parameters.

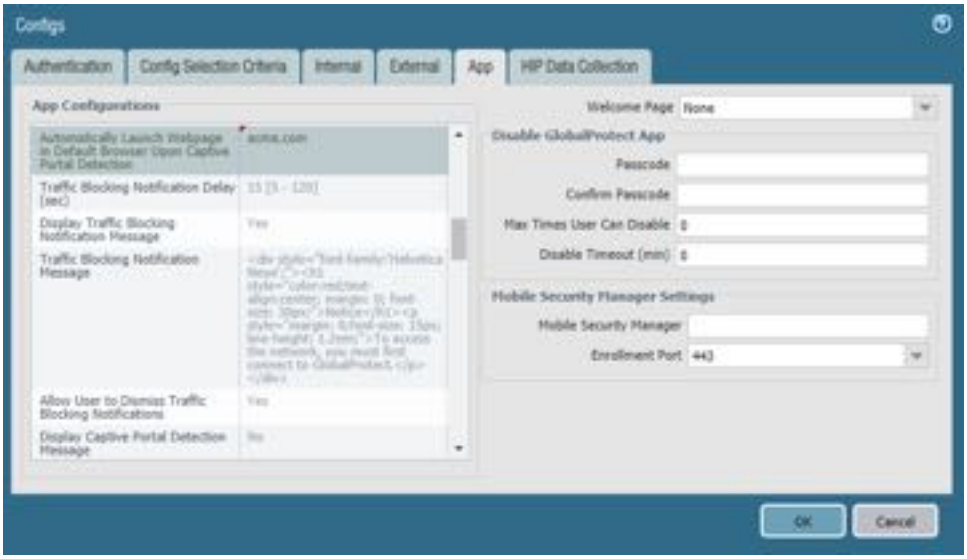
STEP 3 | Set up access to the GlobalProtect portal.

STEP 4 | Define the GlobalProtect agent configurations.

STEP 5 | Customize the GlobalProtect app.

- In the **Automatically Launch Webpage in Default Browser Upon Captive Portal Detection** field, enter the fully qualified domain name (FQDN) or IP address of the website that you want to use for the initial connection attempt that initiates web traffic when the default web browser launches (maximum length is 256 characters). If this field is empty (default), GlobalProtect does not launch the default web browser automatically upon captive portal detection.

© 2020 Palo Alto Networks, Inc.



STEP 6 | **Commit** your changes.

Management Features

- Rule Usage Tracking
- Configuration Table Export
- Reporting Engine Enhancements
- Enhanced Application Logging

© 2020 Palo Alto Networks, Inc.

Rule Usage Tracking

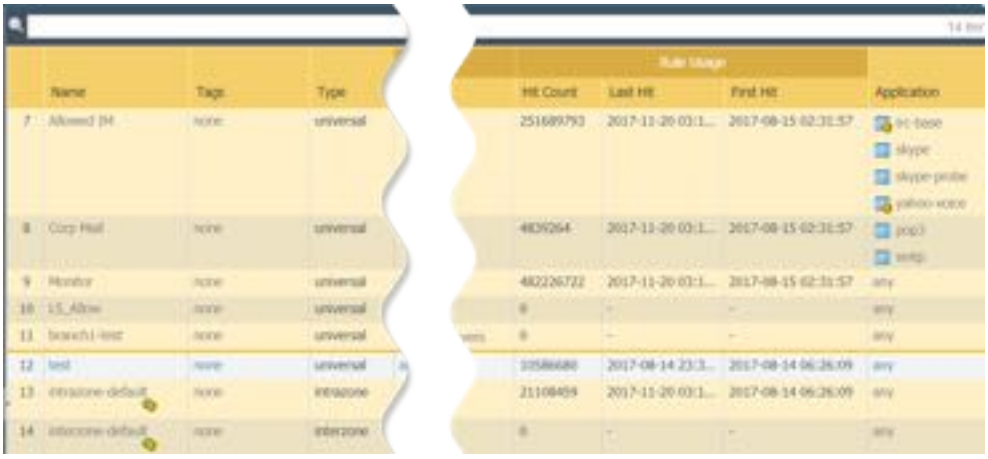
The Panorama and firewall web interfaces now display the hit count for traffic that matches a policy rule to help keep your firewall policies up to date as your environment and security needs change over time. To prevent attackers from exploiting over-provisioned access, such as when a server is decommissioned or when you no longer require temporary access to a service, the rule usage tracking feature helps you - identify and remove unused rules. Additionally, this feature provides the ability to validate rule additions and rule changes and to monitor the time frame when a rule was used. For example, when you migrate port-based rules to app-based rules, you create an app-based rule above the port-based rule and then you check for any traffic that matches the port-based rule. After migration, the hit-count data helps you determine whether it is safe to remove the port-based rule by confirming that traffic is matching the app-based rule instead of the port-based rule.

On the firewall, rule usage tracking allows you to view rule usage hit count and the last timestamp of the last hit. On Panorama, the rule usage tracking data allows you to view whether a policy rule pushed to firewalls in a specific device group has traffic matches. The rule usage tracking data gives you the information you need to determine whether a rule is effective for access enforcement. For more information, see Monitor Policy Rule Usage.

STEP 1 | Launch the firewall or Panorama web interface.

- **On a firewall**

1. Launch the web interface and select **Policies**.
2. View the rule usage statistics for each policy rule. The following information is displayed:
   - **Hit Count**—The number of times traffic matched the criteria you defined in the policy rule. Persists through reboot, dataplane restarts, and upgrades unless you manually reset or rename the rule.
   - **Last Hit**—The most recent timestamp for when traffic matched the rule.
   - **First Hit**—The first instance when traffic was matched to this rule.



- **On Panorama**

1. Launch the web interface and select **Policies**.
2. Determine whether the rule is being used (Rule Usage column). The policy rule usage status is one of the following:

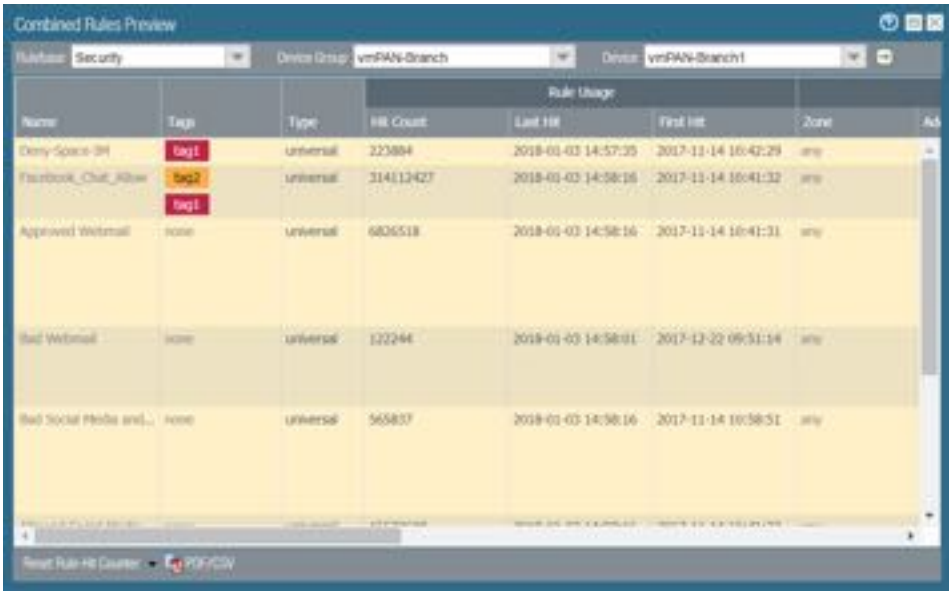
   The Rule Usage column displays rule usage for each appliance in the device group. The rule usage information displayed persists through reboot, dataplane restarts, and upgrades.

© 2020 Palo Alto Networks, Inc.

- **Used**—When all appliances in the device group—to which you pushed the policy rule—have traffic matches for the policy rule.
- **Partially Used**—When some of the appliances in the device group —to which you pushed the policy rule— have matches for the policy rule.
- **Unused**—When no appliances in the device group—to which you pushed the policy rule—have traffic matches for the policy rule.



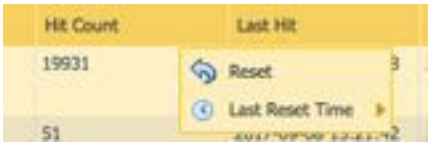
3. **Preview Rules** to view and select a specific firewall managed by Panorama to view the firewall-specific policy rule usage data. If needed, you can reset the firewall hit-count data for individual rules. Panorama retrieves rule usage information from managed firewalls every five minutes.



STEP 2 | Reset the rule usage tracking count data.

You can reset the rule hit count data to validate an existing rule or to gauge rule usage within a specified period of time. Policy rule hit-count data is not stored on the firewall or Panorama so after you clear the hit count using the reset option, that data is no longer available.

1. Identify any rules you need to reset and navigate to the Hit Count column.
2. Select **Reset** from the drop-down. If you previously reset a rule policy hit count, you can also view the **Last Reset Time** from the drop-down.

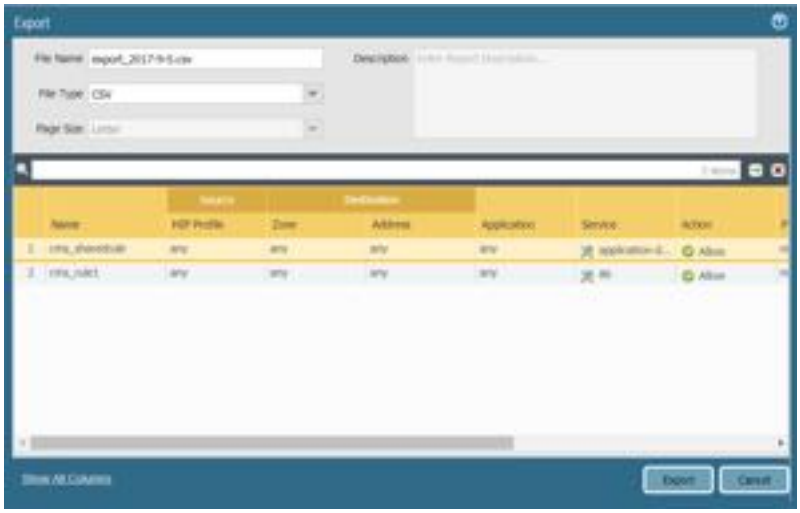


© 2020 Palo Alto Networks, Inc.

Configuration Table Export

You can export policy rules, objects, and IPS signatures from Panorama™ and firewalls to demonstrate regulatory compliance to external auditors, to conduct periodic reviews of the firewall configuration, and to generate reports on firewall policies. This prevents you from having to give auditors direct access to Panorama or the firewall, taking screen shots, or using the XML API to generate configuration reports. From the web interface, you can export the configuration table data for policies, objects, network, firewall, and Panorama configurations, as well as the Signature exceptions in the Antivirus, Anti-Spyware, and Vulnerability Protection Security profiles, in either a PDF or CSV file.

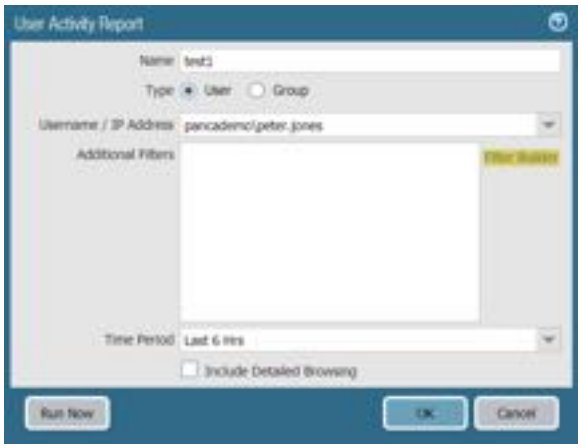
Configuration table export works like a print function—you cannot import generated files back in to Panorama or the firewall. The data that you view on the web interface is exported to **PDF/CSV** and you can use these export formats to apply filters that match your report criteria and search within PDF reports to quickly find specific data. Additionally, when you export the configuration table data, a system log is generated to record the event. For more information on how to export your configurations as a PDF or CSV, see Export Configuration Table Data.



© 2020 Palo Alto Networks, Inc.

Reporting Engine Enhancements

The reporting engine has been enhanced to provide better context on network events and user activity. These enhancements give you greater visibility and control of the log data and reports you generate so that you can create intelligent policies.

Reporting Engine Enhancements	Description
Accurate log directionality	Firewall writes logs from the perspective of who initiated the network interaction or event. For example, if you are downloading a file from a website, the threat log describes the source (SRC) of the network interaction as the individual requesting the download and describes the destination (DST) as the website from which the source is downloading the file. The client is the entity that initiates the request while the server is the entity that receives the request.
Overlay of commits and content updates in the ACC	User events in the **ACC**, such as commits and content updates, are represented as a dotted line on all time-trended line graphs. This helps you to correlate commits and upgrades with suspicious network events.
Filters for User Activity Reports	Use the **Filter Builder** to build custom filters (**Monitor** > **PDF Reports** > **User Activity Report**).  

© 2020 Palo Alto Networks, Inc.

Reporting Engine Enhancements	Description
On Demand Report Scheduling	Select **Run Now** for a scheduled report or **Pick up Later** to return at a later time to retrieve the report instead of waiting for a an on-demand report to generate.  You can also run a **Background Report** to export the generated report in your desired format (**Monitor** > **Manage Custom Report**). You can save only one report for pick up at a time. 
Custom reporting enhancements	Custom reporting is now more user friendly: • Completions are added for user, user-group, and any other items that have possible completions. • Column stickiness is added between databases to keep as many columns as possible when switching between databases.

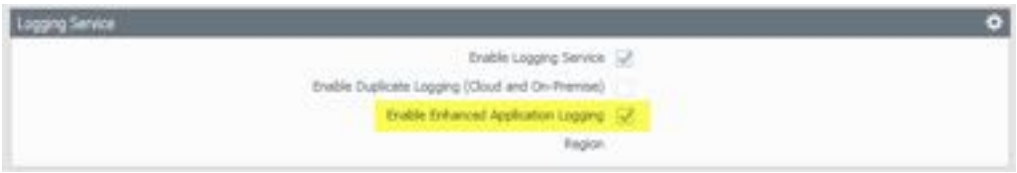
© 2020 Palo Alto Networks, Inc.

Enhanced Application Logging

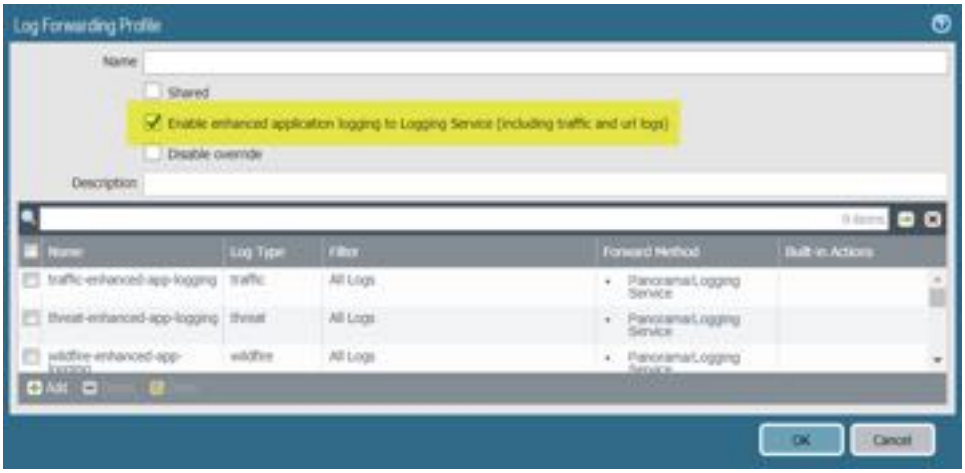
Enhanced application logging allows the firewall to collect data specifically intended to increase visibility into network activity for the Palo Alto Networks apps and services. For example, with this extended network visibility, Magnifier—the Palo Alto Networks behavior analytics service—can better understand your normal network behavior, in order to recognize unusual behavior that might indicate attacker reconnaissance or lateral movement.

The data this feature collects is designed strictly for Palo Alto Networks apps and services to consume and process (you cannot view enhanced application logs). Only turn on this feature in consultation with your Systems Engineer (SE) and after also enabling the Palo Alto Networks logging service.

To turn on **Enable Enhanced Application Logging**, select **Device** > **Setup** > **Management** > **Logging Service** on the firewall web interface:



Then, update the log forwarding profiles that are attached to your security policy rules to **Enable enhanced application logging to Logging Service** (**Objects** > **Log Forwarding**).



Notice that when you enable enhanced application logging in a Log Forwarding profile, match lists that specify the log types required for enhanced application logging are automatically added to the profile. If the Log Forwarding profile that you've updated is not yet attached to a security policy rule, be sure to add it to security policy rules to trigger log generation and forwarding for the traffic matched to those rules (select **Policies** > **Security** > **Actions** > **Log Forwarding** and select the Log Forwarding profile enabled with enhanced application logging).

Networking Features

- Tunnel Content Inspection Logging
- Dynamic IP Address Support for Destination NAT
- FQDN Support for IKE Gateway Peer IP Address
- Configuration Capacity Improvements
- Refresh of Default Trusted CAs
- ARP Cache Timeout

© 2020 Palo Alto Networks, Inc.

Tunnel Content Inspection Logging

To easily view tunnel inspection logs for cleartext tunnels and to separate tunnel inspection logs from traffic logs, you can now configure tunnel content inspection logging and log forwarding. You can configure tunnel content inspection to log tunnel sessions at the start of a session, the end of a session, or both. These log settings override the Security policy rule settings that would otherwise control the tunnel inspection logs. You can also override log forwarding settings in Security policy rules that control traffic logs by configuring tunnel inspection log settings to store tunnel logs separately from traffic logs. The tunnel inspection logs store the outer tunnel (GRE, non-encrypted IPSec, or GTP-U) sessions and the traffic logs store the inner traffic flows. This allows you to easily report on tunnel activity (as opposed to inner content activity) using the ACC and reporting features.

You can create a Log Forwarding profile to specify where to send tunnel inspection logs. A Log Forwarding profile for tunnel inspection is separate from a Log Forwarding profile specified in a Security policy rule, which applies to traffic logs.

When you view a detailed tunnel inspection log, the log now includes the name of the Tunnel Inspection policy rule that applied to the session captured in the log, which makes it easier to track information about non-encrypted tunnel traffic.

STEP 1 | Specify logging of sessions that match a tunnel inspection policy rule and configure log forwarding.

1. Select **Policies** > **Tunnel Inspection** and select a Tunnel Inspection policy rule.
2. Select **Inspection** > **Monitor Options**.
3. Select **Override Security Rule Log Setting**.
4. Select **Log at Session Start** and **Log at Session End**.
5. Select a **Log Forwarding** profile to determine where the firewall forwards tunnel logs for sessions that match the Tunnel Inspection policy rule.

STEP 2 | View tunnel inspection logs.

1. Select **Monitor** > **Logs** > **Tunnel Inspection**.
2. Click Detailed Log View ( ) to see details about a tunnel inspection log.

© 2020 Palo Alto Networks, Inc.

Dynamic IP Address Support for Destination NAT

Destination NAT is enhanced so that you can translate the original destination address to a destination host or server that has a dynamic IP address that is associated with an FQDN and can be resolved by DNS. It is especially helpful to Configure Destination NAT Using Dynamic IP Addresses in cloud deployments, which typically use dynamic IP addressing across multiple servers. Each time the host or server in the cloud receives a new (dynamic) IP address, you don't have to manually update the NAT policy rule by continuously querying the DNS server, nor do you need to use a separate external component to update the DNS server with the latest FQDN-to-IP address mapping. Dynamic IP (with session distribution) supports IPv4 addresses only. The dynamic IP translation type for destination NAT is in addition to the static, one-to-one translation that continues to be supported in this and earlier releases.

If an FQDN in the translated destination NAT address resolves to more than one IP address, the firewall automatically distributes translated sessions among those addresses (based on a round-robin algorithm) to provide improved session distribution.

Using the **Dynamic IP (with session distribution)** destination address type also allows you to translate multiple original destination IP addresses to multiple translated destination IP addresses. A many-to-many translation means, for example, that three original destination IP addresses and four translated destination IP addresses can result in 12 possible destination NAT translations using a single NAT rule.

*You can configure the frequency at which the firewall refreshes an FQDN (*Use Case 1: Firewall Requires DNS ResolutionforManagement Purposes*).*

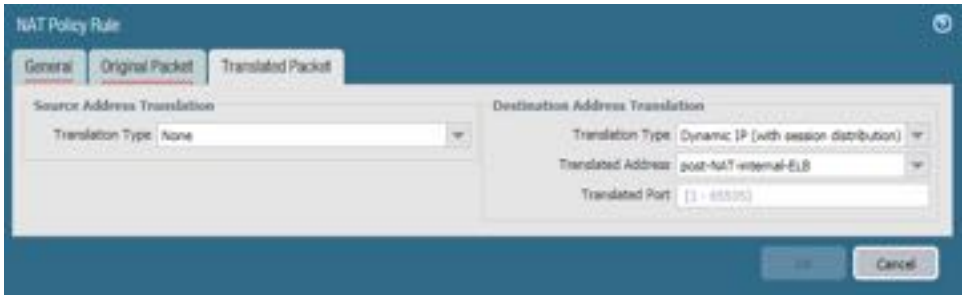
STEP 1 | Create an address object using the FQDN of the ELB or server to which you want to translate the address.



STEP 2 | Create the destination NAT policy.

1. Specify the original packet to use the publicly routed IP address of the service hosted behind the firewall.
2. Configure the **Translation Type** for the translated packet as **Dynamic IP (with session distribution)**.
3. Enter the FQDN address object (that you created) as the **Translated Address**.

© 2020 Palo Alto Networks, Inc.



4. Click **OK**.

STEP 3 | **Commit** your changes.

© 2020 Palo Alto Networks, Inc.

FQDN Support for IKE Gateway Peer IP Address

When you configure an IPSec tunnel with an IKE gateway peer, you can now configure that address as an FQDN or an address object that uses an FQDN. Using an FQDN for the peer address saves you from repeatedly reconfiguring peer addresses in several scenarios. An FQDN prevents IKE exchange problems that arise when many branch offices use a DHCP-assigned address on their external interface and that dynamic address changes. Similarly, FQDNs are a benefit in cloud environments where AWS and Azure use dynamic addresses as IKE termination points.

Another use case is when you have several satellite offices with multiple hub locations and VPN connectivity between firewalls at the satellites and hub gateway. You can configure each satellite office to use an FQDN for the IKE peer address object, so that if one hub goes down, the DNS server for that FQDN then resolves the FQDN to the IP address for the second hub. You don't have to reconfigure the IKE peer to use the IP address of the second hub.

STEP 1 | Set up an IKE gateway.

Perform the first two steps to Set Up an IKE Gateway (define the gateway and establish the local endpoint of the tunnel).

STEP 2 | Specify the IKE peer IP address for the peer at the far end of the tunnel (gateway) as an FQDN. You can enter the FQDN string directly or use the FQDN in an address object.

STEP 3 | Continue to Set Up an IKE Gateway, resuming with the step where you specify how the peer is authenticated.

© 2020 Palo Alto Networks, Inc.

Configuration Capacity Improvements

Various firewall models support larger configuration capacities in PAN-OS® 8.1 than in earlier PAN-OS releases. These increases ease your migration to Palo Alto Networks® firewalls and your deployment growth. The increased capacities are:

Capacity Description	All Firewall Models
Characters per Rule Name	63

PA-220 firewalls support more Security policy rules:

Capacity Description	PA-220
Security Policy	500

PA-3000 Series firewalls support the following capacities:

Capacity Description	PA-3020	PA-3050	PA-3060
Address Groups	750	1,500	1,500
Service Groups	375	375	375
Service Entries per Service Group	1,000	1,000	1,000
Zones	40	40	60

PA-5220 firewalls support the following increased capacities:

Capacity Description	PA-5220
DHCP Relay Agents	2,048 (On the PA-5220 firewall, you can have a maximum of 500 DHCP servers; the maximum number of DHCP relay agents is 2,048 minus the number of configured DHCP servers.)
FQDN Address Objects	6,144
Zones	2,500

PA-5250, PA-5260, and PA-7000 Series firewalls support the following increased capacities:

Capacity Description	PA-5250 and PA-5260	PA-7000 Series
Address Objects	160,000	160,000

© 2020 Palo Alto Networks, Inc.

Capacity Description	PA-5250 and PA-5260	PA-7000 Series
Address Groups	80,000	80,000
Service Objects	12,000	12,000
Service Groups	6,000	6,000
Service Entries per Service Group	2,500	2,500
FQDN Address Objects	6,144	6,144
Zones	17,000	4,000
Security Policy	65,000	65,000
Tunnel Inspection Policy	8,500	8,500
DHCP Relay Agents	4,096*	4,096*
	* On these models, you can have a maximum of 500 DHCP servers; the maximum number of DHCP relay agents is 4,096 minus the number of configured DHCP servers.	

PA-820 and PA-850 firewalls support the following increased capacities:

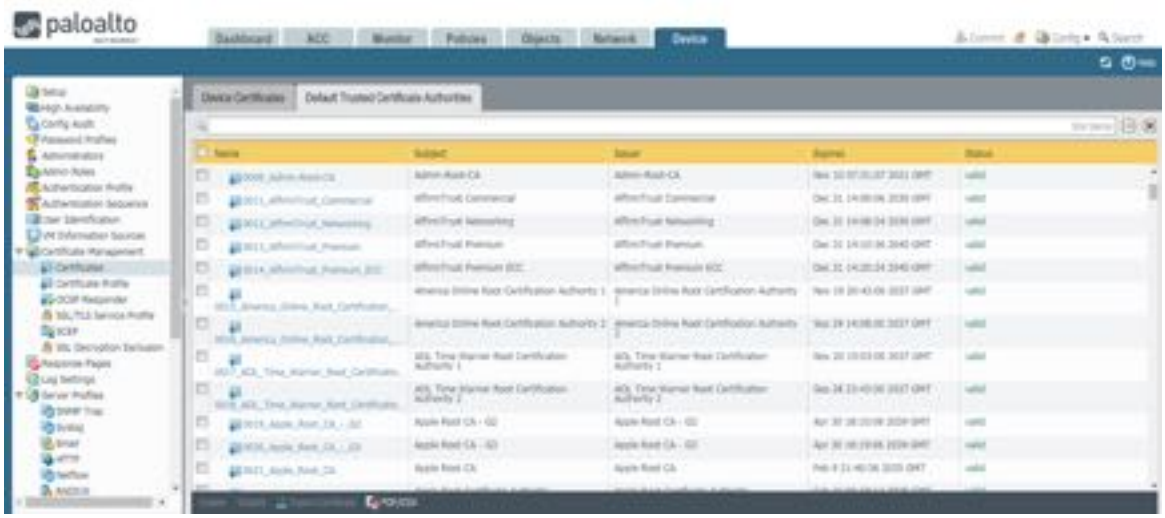
Capacity Description	PA-820	PA-850
Maximum SSL Inbound Certificates	75	100
SSL Certificate Cache Entries	1,000	2,000

© 2020 Palo Alto Networks, Inc.

Refresh of Default Trusted CAs

The certificate authorities (CAs) that the firewall trusts by default are updated in PAN-OS 8.1; new trusted root CAs are added and expired CAs are removed. The pre-installed list of CAs includes the most common and trusted certificate providers responsible for issuing the certificates the firewall requires to secure connections to the internet. Because the firewall trusts these CAs by default, the only additional CAs you might want to add are any trusted enterprise CAs that your organization requires.

To view and manage the list of CAs that the firewall trusts by default, select **Device** > **Certificate Management** > **Certificates** > **Default Trusted Certificate Authorities**. For details on the different keys and certificates that the firewall uses, and how to obtain and manage them, seeCertificate Management.



© 2020 Palo Alto Networks, Inc.

ARP Cache Timeout

You can change the length of time that the firewall keeps IP address-to-hardware address mappings in its ARP cache to suit your environment.

STEP 1 | Access the CLI and specify how many seconds the firewall keeps ARP entries in its cache. Use the operational CLI command **`set system setting arp-cache-timeout`** <***value***>; range is 60 to 65,535, default is 1,800.

STEP 2 | View the ARP cache timeout setting with the operational CLI command **`show system setting arp-cache-timeout`**.

User-ID Features

> Support for Multiple Username Formats

© 2020 Palo Alto Networks, Inc.

Support for Multiple Username Formats

The firewall can now identify a user even if different User-ID sources send usernames in different formats. For example, a single user may have multiple usernames that are represented in different formats (**`jane.doe@domain.com`**, **`DOMAIN\jdoe`**, **`jdoe`**)

The usernames are matched based on the user attributes that the firewall reads from the LDAP-compliant directory. You can specify which user attributes to collect from the directory using the Group Mapping profile.

Because the firewall now supports multiple user attributes, you should specify an attribute as the *Primary Username* for users. The primary username represents the user in the logs, reports, and in the policy configuration.

If your User-ID sources send usernames without an associated domain and your usernames are unique, you can also configure the firewall to not consider the domain when matching users. If you enable this option and the firewall finds more than one matching username, an error displays to indicate the username is not unique.

STEP 1 | Select **Device** > **Server Profiles** > **LDAP** and **Add** an LDAP server profile.

STEP 2 | Select **Device** > **User Identification** > **Group Mapping Settings** and **Add** Group Mapping using the LDAP server profile you added in the previous step.

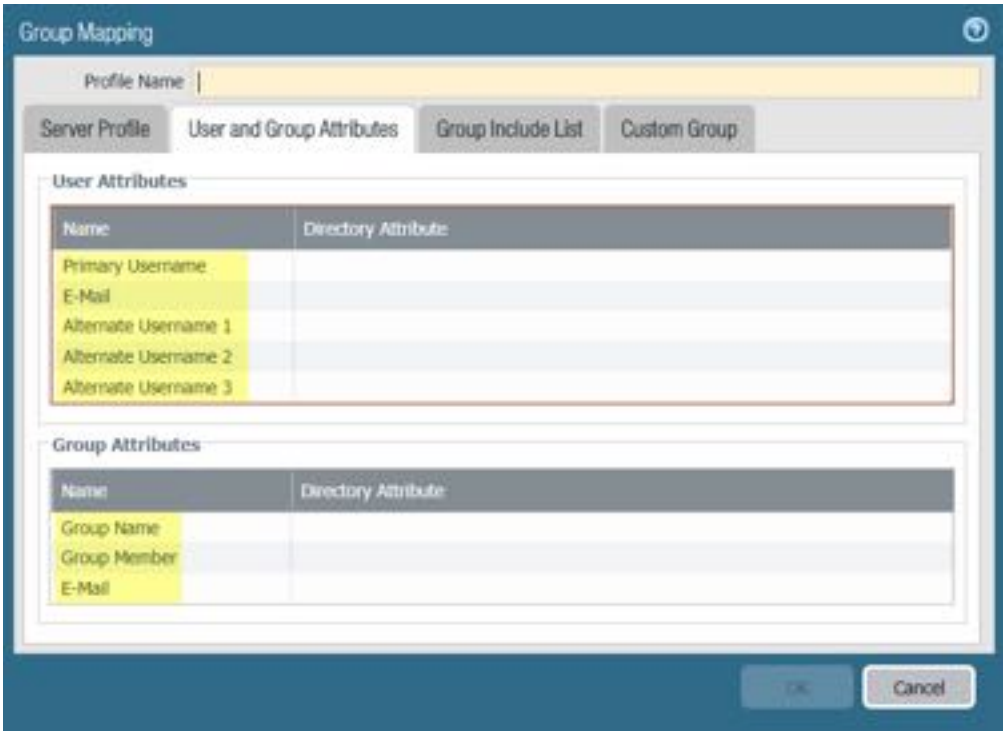
STEP 3 | Specify the **Primary Username** that will identify users in reports and logs and optionally specify the **Directory Attribute** for users or groups, then **Commit** your changes.

When you select the Server Profile **Type**, the firewall auto-populates the values for the user and group attributes. Based on the user information that your User-ID sources send, you may need to configure the correct attributes. For more information, refer to Map Users to Groups.

- **For users:**
  1. Select **Device** > **User Identification** > **Group Mapping Settings** > **Add** > **User and Group Attributes** and specify a **Primary Username** (for example, **`userPrincipalName`** or **`sAMAccountName`**).

     If the Primary Username is in User Principal Name (UPN) format, it will not be normalized in the `domain\username` format as in previous versions. For example, if the Primary Username is received in the UPN format, it will be displayed as `username@domain`, not `domain\username`.
  2. (Optional) Specify additional alternate **User Attributes** to identify users, such as an **E-Mail** or up to three **Alternate Usernames**.
- **For groups:**
  1. Select **Device** > **User Identification** > **Group Mapping** > **Add** > **User and Group Attributes**.
  2. Specify **Group Attributes** such as the **Group Name**, **Group Member**, or **E-Mail**.

© 2020 Palo Alto Networks, Inc.



STEP 4 | Use the groups and usernames that the group mapping profile collects to Enable User- and Group-Based Policy.

(Optional) If your User-ID sources only send the username and the username is unique across the organization, select **Device** > **User Identification** > **Palo Alto Networks User-ID Agent Setup** > **Cache** and **Edit** to **Allow matching usernames without domain** to allow the firewall to check if unique usernames collected from the LDAP server during group mapping match the users associated with a policy and avoid overwriting the domain in your source profile. This option is disabled by default.

Before enabling this option, configure group mapping for the LDAP group containing the User-ID source (such as GlobalProtect or Captive Portal) that collects the mappings. After you commit the changes, the User-ID source populates the usernames without domains. Only usernames collected during group mapping can be matched without a domain.
If your User-ID sources send user information in multiple formats and you enable this option, verify that the attributes collected by the firewall have a unique prefix. To ensure users are identified correctly if you enable this option, all attributes for group mapping should be unique. If the username is not unique, the firewall logs an error in the Debug logs.



© 2020 Palo Alto Networks, Inc.

STEP 5 | Map users based on information from User-ID sources by configuring User-ID to gather IP-user mappings from sources using the PAN-OS integrated User-ID agent or the Windows User-ID Agent.

STEP 6 | Verify the user mapping is successful:

- To verify the Group Mapping configuration, select the **Group Include List** to confirm the firewall has fetched all of the groups.
- To verify all the user attributes have been captured correctly, use the **`show user user-attributes user all`** command.
- Verify the usernames are displayed correctly in the **Source User** column of the Monitor tab.
- Select **Monitor > Logs > User-ID** and check the **User Provided by Source** column to verify the users are mapped to the correct username.

© 2020 Palo Alto Networks, Inc.

WildFire Features

- WildFire Appliance-to-Appliance Encryption

For more information about the features introduced in the WildFire cloud, see the WildFire What's New Guide.

© 2020 Palo Alto Networks, Inc.

WildFire Appliance-to-Appliance Encryption

You can now encrypt WildFire® communications between appliances deployed in a cluster. Prior to 8.1 and by default, WildFire appliances send data using cleartext when communicating with management appliances as well as WildFire cluster peers. You can use either predefined or custom certificates to authenticate connections between WildFire appliance peers using the IKE/IPsec protocol. The predefined certificates meet current FIPS/CC/UACPL-approved certification and compliance requirements. If you want to use custom certificates instead, you must select a FIPS/CC/UACPL-compliant certificate or you will not be able to import the certificate.

You can configure WildFire appliance-to-appliance encryption locally using the WildFire CLI or centrally through Panorama. Keep in mind, all WildFire appliances within a given cluster must run a version of PAN-OS that supports encrypted communications.



If the WildFire appliances in your cluster uses FIPS/CC mode, encryption is automatically enabled using predefined certificates.

Before configuring WildFire appliance-to-appliance encryption, be sure to review your existing WildFire secure communications configuration. If you previously configured the WildFire appliance and the firewall for secure communications using a custom certificate, you can use that custom certificate and the requisite DNS name for configuring secure communications between WildFire appliances.



It is imperative that you use the correct, matching DNS name in the register firewall to: field in Panorama. Failure to do so will prevent appliance-to-appliance encryption from working as intended.

The following tables describe the high-level tasks involved in configuring WildFire appliance-to-appliance encryption. For detailed instructions on these tasks, refer to the WildFire Administration Guide for the full installation procedure.

Table 2: WildFire Appliance-to-Appliance Encryption Installation Tasks Using Panorama

Configuration Using Custom Certificates through Panorama	Configuration Using Predefined Certificates through Panorama
1. In the **Panorama > Managed WildFire Clusters > WildFire Cluster** page: • Configure the DNS name used for authentication found in the custom certificate. • Enable customize secure server communication and configure the SSL/TLS Service Profile and certificate profile to define the custom certificate for encrypted communication between WildFire peers. • Import or generate a custom certificate. If you are generating a custom certificate, be sure to use the same DNS stated in the certificate. 2. In the **Device > Setup > Management > Secure Communication Settings** page,	1. In the **Panorama > Managed WildFire Clusters > WildFire Cluster** page: • Enable **Secure Cluster Communication**. • Enable **HA Traffic Encryption**.

© 2020 Palo Alto Networks, Inc.

Configuration Using Custom Certificates through Panorama	Configuration Using Predefined Certificates through Panorama
configure the firewall secure communication settings to use the custom certificate. This enables the WildFire cluster to communicate with the firewall using encryption. 3. In the **Panorama > Managed WildFire Clusters > WildFire Cluster** page: • Make sure you configure the appliance to use **Custom Certificates Only**, so that it does not use the predefined certificate. • Enable**Secure Cluster Communication**. • Enable **HA Traffic Encryption**.	

Table 3: WildFire Appliance-to-Appliance Encryption Installation Tasks Using the CLI

Configuration Using Custom Certificates through the WildFire CLI	Configuration Using Predefined Certificates through the CLI
To configure the WildFire appliance for encrypted communications, you must enable and configure the following on the active-controller in 2-node clusters. If your cluster has 3 or more nodes, you must also duplicate the configuration on the server nodes.	
1. In the **Panorama > Managed WildFire Clusters > WildFire Cluster** page: • Configure the DNS name used for authentication found in the custom certificate. • Enable customize secure server communication and configure the SSL/TLS Service Profile and certificate profile to define the custom certificate for encrypted communication between WildFire peers. • Import or generate a custom certificate. If you are generating a custom certificate, be sure to use the same DNS stated in the certificate. 2. Configure the firewall secure communication settings to use the custom certificate. This enables the WildFire cluster to communicate with the firewall using encryption. 3. From the WildFire appliance cluster active-controller CLI: • Make sure you configure the appliance to use custom certificates only, so that it does not use the predefined certificate. • Enable secure cluster communication.	1. From the WildFire appliance cluster active-controller CLI: • Enable secure cluster communication. • Enable HA traffic encryption.

© 2020 Palo Alto Networks, Inc.

Configuration Using Custom Certificates through the WildFire CLI	Configuration Using Predefined Certificates through the CLI
• Enable HA traffic encryption.	

© 2020 Palo Alto Networks, Inc.

Feature	PA-5260 	PA-5250 	PA-5220 
Performance			
**1. Firewall throughput measured with App-ID and User-ID features enabled utilizing 64KB HTTP transactions. 2. Threat prevention throughput measured with App-ID, User-ID, IPS, antivirus and anti-spyware features enabled utilizing 64KB HTTP transactions. 3. New sessions per second measured with 4KB HTTP transactions. Additionally, for VM models please refer to hypervisor, cloud specific data sheet for associated performance.*			
App-ID firewall throughput	68 Gbps	39 Gbps	18 Gbps
Threat prevention throughput	30 Gbps	20 Gbps	9 Gbps
IPSec VPN throughput	24 Gbps	16 Gbps	8 Gbps
Connections per second	462,000	348,000	171,000
Sessions			
Max sessions (IPv4 or IPv6)	32,000,000	8,000,000	4,000,000
Policies			
Security rules	65,000	65,000	20,000
Security rule schedules	256	256	256
NAT rules	16,000	8,000	6,000
Decryption rules	5,000	5,000	2,000
App override rules	4,000	4,000	2,000
Tunnel content inspection rules	8,500	8,500	2,000
Policy based forwarding rules	2,000	2,000	2,000
Captive portal rules	8,000	8,000	8,000
DoS protection rules	2,000	2,000	1,000
Security Zones			
Max security zones	17,000	17,000	2,500
Objects (addresses and services)			
Address objects	160,000	160,000	40,000
Address groups	80,000	80,000	4,000
Members per address group	2,500	2,500	2,500

Service objects	12,000	12,000	2,000
Service groups	6,000	6,000	250
Members per service group	2,500	2,500	500
FQDN address objects	6,144	6,144	6,144
Max IP addresses registered per system **Applies to IP addresses registered to dynamic address groups*	100,000	100,000	100,000
Tags per IP address	32	32	32
Security Profiles			
Security profiles	750	750	750
App-ID			
Custom App-ID signatures	6,000	6,000	6,000
Shared custom App-IDs	512	512	512
Custom App-IDs (virtual system specific)	6,416	6,416	6,416
User-ID			
User-IP mappings (management plane)	512,000	512,000	512,000
User-IP mappings (data plane)	512,000	512,000	512,000
Active and unique groups used in policy	10,000	10,000	10,000
Number of agents	100	100	100
Monitored servers per agent	100	100	100
Maximum terminal services agents	1,000	1,000	1,000
SSL Decryption			
Max SSL inbound certificates	2,000	1,200	600
SSL certificate cache (forward proxy)	32,000	24,000	16,000
Max concurrent decryption sessions	3,200,000	800,000	400,000
SSL Port Mirror	Yes	Yes	Yes
SSL Decryption Broker	Yes	Yes	Yes
HSM Supported	Yes	Yes	Yes
URL Filtering			
Total entries for allow list, block list and custom categories	100,000	100,000	100,000
Max custom categories	2,849	2,849	2,849
Max custom categories (virtual system specific)	500	500	500
Dataplane cache size for URL filtering	100,000	100,000	100,000
Management plane dynamic cache size	1,000,000	1,000,000	1,000,000

Interfaces

Mgmt - out-of-band	10/100/1000, RJ45 console	10/100/1000, RJ45 console	10/100/1000, RJ45 console
Mgmt - 10/100/1000 high availability	NA	NA	NA
Mgmt - 40Gbps high availability	1	1	1
Mgmt - 10Gbps high availability	NA	NA	NA
Traffic - 10/100/1000	NA	NA	NA
Traffic - 100/1000/10000	4	4	4
Traffic - 1Gbps SFP	0/16	0/16	0/16
Traffic - 10Gbps SFP+	0/16	0/16	0/16
Traffic - 10Gbps XFP	NA	NA	NA
Traffic - 40Gbps QSFP	4x40/100	4x40/100	4X40
802.1q tags per device	4,094	4,094	4,094
802.1q tags per physical interface	4,094	4,094	4,094
Max interfaces (logical and physical)	5,120	4,096	4,096
Maximum aggregate interfaces	8	8	8

Virtual Routers

Virtual routers	225	125	20

Virtual Wires

Virtual wires	2,048	2,048	2,048

Virtual Systems

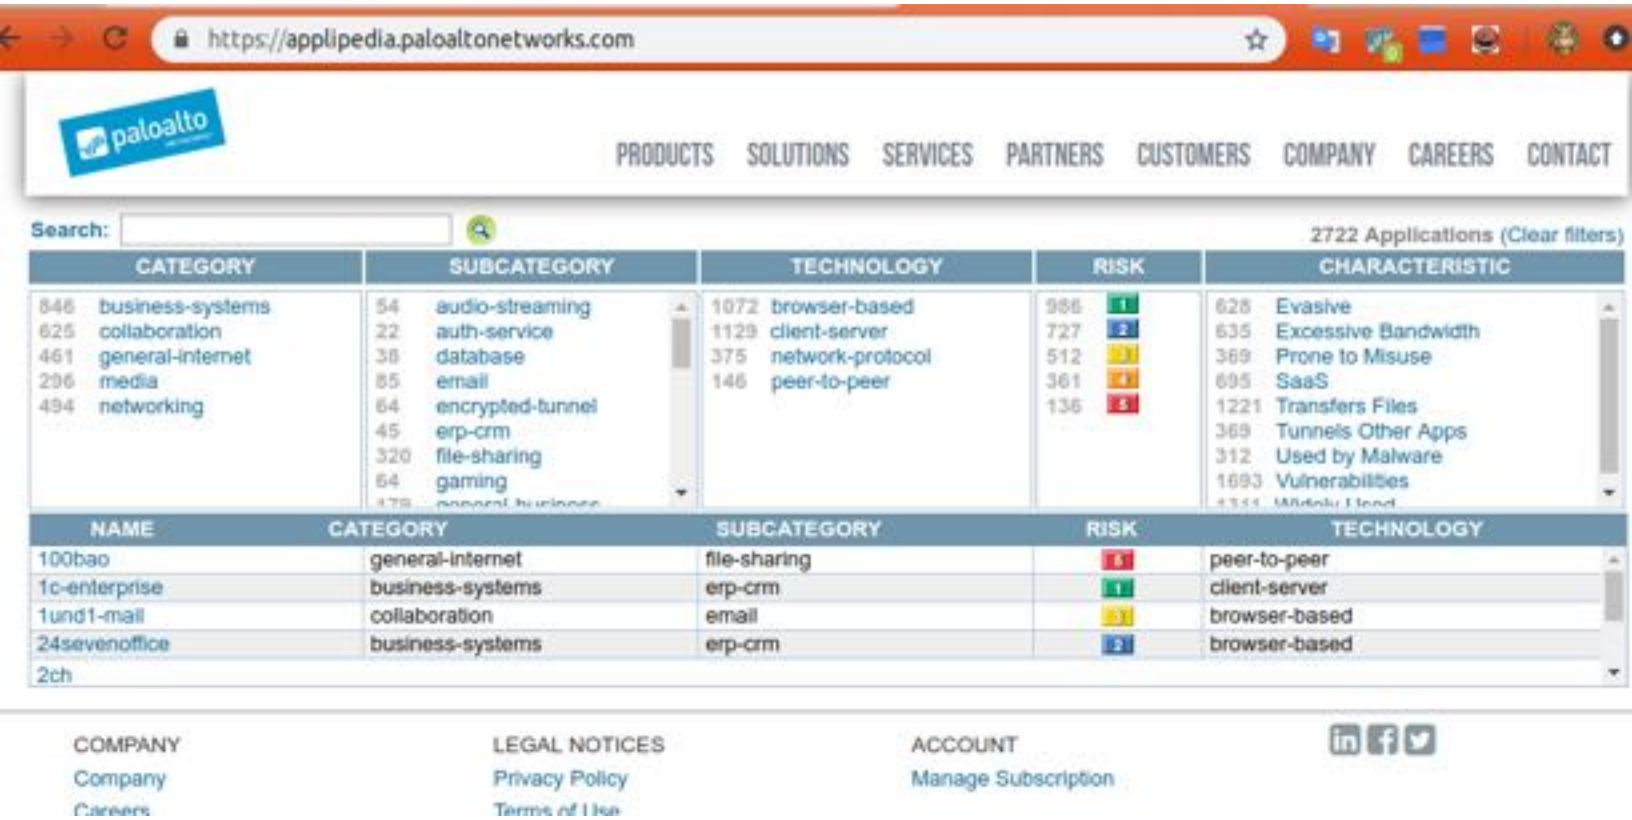
Base virtual systems	25	25	10
Max virtual systems **Additional licenses are required for virtual system capacities above the base virtual systems capacity*	225	125	20

Routing

IPv4 forwarding table size **Entries shared across virtual routers*	100,000	100,000	100,000
IPv6 forwarding table size **Entries shared across virtual routers*	100,000	100,000	100,000
System total forwarding table size	200,000	200,000	200,000
Max route maps per virtual router	50	50	50
Max routing peers (protocol dependent)	1,000	1,000	1,000

Static entries - DNS proxy	1,024	1,024	1,024
Bidirectional Forwarding Detection (BFD) Sessions	1,024	1,024	1,024
L2 Forwarding			
ARP table size per device	128,000	128,000	128,000
IPv6 neighbor table size	128,000	128,000	128,000
MAC table size per device	128,000	128,000	128,000
Max ARP entries per broadcast domain	128,000	128,000	128,000
Max MAC entries per broadcast domain	128,000	128,000	128,000
NAT			
Total NAT rule capacity	16,000	8,000	6,000
Max NAT rules (static) **Configuring static NAT rules to full capacity requires that no other NAT rule types are used.*	16,000	8,000	6,000
Max NAT rules (DIP) **Configuring DIP NAT rules to full capacity requires that no other NAT rule types are used.*	16,000	8,000	4,000
Max NAT rules (DIPP)	8,000	6,000	4,000
Max translated IPs (DIP)	320,000	160,000	64,000
Max translated IPs (DIPP) **DIPP translated IP capacity is proportional to the DIPP pool oversubscription value. The capacity shown here is based on an oversubscription value of 1x.*	8,000	6,000	4,000
Default DIPP pool oversubscription **Source IP and source port reuse across concurrent sessions*	8	8	8
Address Assignment			
DHCP servers	500	500	500
DHCP relays **Maximum capacity represents total DHCP servers and DHCP relays compbined*	4,096	4,096	2,048
Max number of assigned addresses	64,000	64,000	64,000
High Availability			
Devices supported	2	2	2
Max virtual addresses	4,096	4,096	4,096
QoS			
Number of QoS policies	4,000	4,000	4,000
Physical interfaces supporting QoS	12	12	12

Clear text nodes per physical interface	63	63	63
DSCP marking by policy	Yes	Yes	Yes
Subinterfaces supported	2,048	2,048	2,048
IPSec VPN			
Site to site	15,000	12,000	10,000
Max IKE Peers	5,000	4,000	3,000
GlobalProtect Client VPN			
Max tunnels (SSL, IPSec, and IKE with XAUTH)	60,000	30,000	15,000
GlobalProtect Clientless VPN			
Max SSL tunnels	10,000	5,000	2,500
Multicast			
Replication (egress interfaces)	2,000	2,000	1,000
Routes	4,000	4,000	4,000
Product Notes			
End-of-sale	NA	NA	NA


https://applipedia.paloaltonetworks.com
paloalto
PRODUCTS SOLUTIONS SERVICES PARTNERS CUSTOMERS COMPANY CAREERS CONTACT
Search:
2722 Applications (Clear filters)
CATEGORY
846 business-systems
625 collaboration
461 general-internet
296 media
494 networking
SUBCATEGORY
54 audio-streaming
22 auth-service
38 database
85 email
64 encrypted-tunnel
45 erp-crm
320 file-sharing
64 gaming
TECHNOLOGY
1072 browser-based
1129 client-server
375 network-protocol
146 peer-to-peer
RISK
986 1
727 2
512 3
361 4
136 5
CHARACTERISTIC
628 Evasive
635 Excessive Bandwidth
369 Prone to Misuse
695 SaaS
1221 Transfers Files
369 Tunnels Other Apps
312 Used by Malware
1693 Vulnerabilities
NAME
CATEGORY
SUBCATEGORY
RISK
TECHNOLOGY
100bao
general-internet
file-sharing
5
peer-to-peer
1c-enterprise
business-systems
erp-crm
1
client-server
1und1-mail
collaboration
email
3
browser-based
24sevenoffice
business-systems
erp-crm
3
browser-based
2ch
COMPANY
Company
LEGAL NOTICES
Privacy Policy
ACCOUNT
Manage Subscription

Platform Specifications and Features Summary

Performance and Capacities[1]	PA-7080 System[3]	PA-7050 System	PA-5060	PA-5050	PA-5020
Firewall throughput (App-ID enabled)	200 Gbps	120 Gbps	20 Gbps	10 Gbps	5 Gbps
Threat prevention throughput	100 Gbps	60 Gbps	10 Gbps	5 Gbps	2 Gbps
IPSec VPN throughput	80 Gbps	48 Gbps	4 Gbps	4 Gbps	2 Gbps
New sessions per second	1,200,000	720,000	120,000	120,000	120,000
Max sessions	40,000,000	24,000,000	4,000,000	2,000,000	1,000,000
Virtual systems (base/max[2])	25/225	25/225	25/225	25/125	10/20
Hardware Specifications	**PA-7080 System[3]**	**PA-7050 System**	**PA-5060**	**PA-5050**	**PA-5020**
Interfaces supported NPC option 1[3]	Up to (20) QSFP+, (120) SFP+	Up to (12) QSFP+, (72) SFP+	(12) 10/100/1000, (8) SFP, (4) 10 SFP+		(12) 10/100/1000, (8) SFP
Interfaces supported NPC option 2[3]	Up to (120) 10/100/1000, (80) SFP, (40) SFP+	Up to (72) 10/100/1000, (48) SFP, (24) SFP+			
Management I/O	(2) 10/100/1000, (2) QSFP+ high availability, (1) 10/100/1000 out-of-band management, (1) RJ45 console		(2) 10/100/1000 high availability, (1) 10/100/1000 out-of-band management, (1) RJ45 console		
Rack mountable?	19U, 19" standard rack	9U, 19" standard rack or 14U, 19" standard rack with optional Airduct kit	2U, 19" standard rack		
Power supply	4x2500W AC (2400W / 2700) expandable to 8	4x2500W AC (2400W / 2700W)	Redundant 450W AC or DC		
Redundant power supply?	Yes		Yes		
Disk drives	2TB RAID1		120GB or 240GB SSD, RAID Optional		
Hot swap fans	Yes		Yes		

Performance and Capacities[1]	PA-3060	PA-3050	PA-3020	PA-500	PA-200
Firewall throughput (App-ID enabled)	4 Gbps	4 Gbps	2 Gbps	250 Mbps	100 Mbps
Threat prevention throughput	2 Gbps	2 Gbps	1 Gbps	100 Mbps	50 Mbps
IPSec VPN throughput	500 Mbps	500 Mbps	500 Mbps	50 Mbps	50 Mbps
New sessions per second	50,000	50,000	50,000	7,500	1,000
Max sessions	500,000	500,000	250,000	64,000	64,000
Virtual systems (base/max[2])	1/6	1/6	1/6	N/A	N/A
Hardware Specifications	**PA-3060**	**PA-3050**	**PA-3020**	**PA-500**	**PA-200**
Interfaces supported[3]	(8) 10/100/1000, (8) SFP, (2) 10 SFP+	(12) 10/100/1000, (8) SFP		(8) 10/100/1000	(4) 10/100/1000
Management I/O	(1) 10/100/1000 out-of-band management,(2) 10/100/1000 high availability, (1) RJ-45 console			(1) 10/100/1000 out-of-band management, (1) RJ-45 console port	
Rack mountable?	1.5U, 19" standard rack	1U, 19" standard rack		1U, 19" standard rack	1.75" H x 7"D x 9.25"
Power supply	Redundant 400W AC	250W AC		180W	40W
Redundant power supply?	Yes	No		No	No
Disk drives	120GB SSD			160GB	16GB SSD
Hot swap fans	No			No	No

Performance and Capacities[4]	VM-1000-HV	VM-300	VM-200	VM-100
Firewall throughput (App-ID enabled)	1 Gbps	1 Gbps		
Threat prevention throughput	600 Mbps	600 Mbps		
IPSec VPN throughput	250 Mbps	250 Mbps		
New sessions per second	8,000	8,000		
Max sessions	250,000	250,000	100,000	50,000
Virtualization Specifications	**VM-1000-HV**	**VM-300**	**VM-200**	**VM-100**
HyperVisor	• VMware ESXi 5.5/6.0, NSX Manager 6.0/6.1/6.2 (Required for NSX integrated solution) • VMware ESXi 5.1/5.5/6.0 (Stand alone) • KVM on CentOS/RHEL 6.5 and Ubuntu 12.04 LTS • Citrix Xen Server on SDX 10.1 • Amazon AWS	• VMware ESXi 5.1/5.5/6.0 • KVM on CentOS/RHEL 6.5 and Ubuntu 12.04 LTS • Citrix Xen Server on SDX 10.1 • Amazon AWS		
Network drivers	• VMware ESXi: VMXNet 3 • Citrix NetScaler SDX: Igbvf version 2.0.4, Ixgbevf version 2.7.12 • KVM: virtIO, e1000, SR-IOV and PCI passthrough for Intel 82576 based 1G NIC, Intel 82599 based 10G NIC, Broadcom 57112 and 578xx based 10G NIC • Amazon Web Services: proprietary			
Dedicated CPU cores	2, 4 or 8			
Dedicated Memory (Minimum)	4GB			
Dedicated Disk drive capacity (Min/Max)	40GB/2TB			

(1) Performance and capacities are measured under ideal testing conditions with App-ID enabled and PAN-OS 7.0. (2) Adding virtual systems to the base quantity requires a separately purchased license. (3) QSFP, SFP, SFP+ and XFP transceivers are sold separately. PAN-OS 7.0 is required on the PA-7080 and when NPC option 1 is used in the PA-7050. (4) VM-Series performance and capacities are measured under ideal testing conditions with App-ID enabled and PAN-OS 7.0 using 4 CPU cores.

Platform Specifications and Features Summary

Key Features	Supported Across All Platforms
Firewall	
Applications identified (App-IDs)	2,020+
Customizable App-IDs	✓
Policy-based control by user, group or IP address, application, application category, subcategory, technology, risk factor or characteristic	✓
Scheduled policies	✓
Per policy SSL decryption & inspection (inbound and outbound)/Per policy SSH control (inbound and outbound)	✓
IPSec VPN (site-to-site) and SSL VPN (GlobalProtect)	✓
User control: Active Directory, LDAP, eDirectory, Exchange, Citrix and Microsoft Terminal Services, XML API (User-ID)	✓
Threat Prevention (Subscription Required)	
Stream-based protection against viruses (including those embedded in HTML, Javascript, PDF and compressed), spyware, worms	✓
Drive-by download protection	✓
Behavioral botnet detection	✓
Modern Malware Protection (WildFire Subscription Required)	
Dynamic analysis and detection of unknown malware within PE, PDF, Java, Android APK, Adobe Flash, all Microsoft Office file types and webpages.	✓
Automated signature generation and delivery for discovered malware	✓
Inline control of malware infection and command/control traffic	✓
URL Filtering (Subscription Required)	
On-box customizable URL filtering database	✓
Customizable categories, allow, block lists and block pages	✓
Safe search (Google, Bing, Yahoo)	✓
File and Data Filtering	
Bidirectional control over the unauthorized transfer of more than 60 file types	✓
Bidirectional control over the transfer of Social Security Numbers, Credit Card Numbers, custom data patterns	✓
QoS	
Policy-based traffic shaping (priority, guaranteed, maximum) per application, per user, per tunnel and more	✓
Policy-based diffserv marking	✓
VPN Connectivity	
GlobalProtect (IPSec/SSL VPN)	✓
IPSec VPN (Site-to-Site)	✓
Mobile Security (GlobalProtect Mobile Security Manager, Gateway & App subscription or licenses required)	
Monitor and report device state for policy enforcement, such as jailbroken devices	✓
Manage device settings, configure key device functions, detect Android malware	✓
Deliver mobile threat prevention and policy enforcement based on apps, users, content, device and device state	
Networking	
Dynamic routing (BGP, OSPF, RIP)	✓
Tap mode, virtual wire, layer 2, and layer 3	✓
802.1Q VLAN tagging (layer 2, layer 3)	✓
Network address translation (NAT)	✓
DHCP server/ DHCP relay (up to 3 servers)	✓
IPv6 L2, L3, tap, virtual wire (transparent mode), App-ID, User-ID, Content-ID, WildFire and SSL decryption	✓
Multicast	✓
High Availability	
Active/Active, Active/Passive	✓ **(PA-200 supports HA Lite only)**
Configuration and session synchronization	✓
Interface and IP tracking	✓
Link and path failure monitoring	✓
Management and Visibility Tools	
Integrated web interface, CLI and centralized management (Panorama)	✓
Graphical summary of applications, URL categories, threats and data (ACC)	✓
View, filter, export traffic, threat, URL, and data filtering logs	✓
Syslog and SNMPv2/v3	✓
XML-based REST API	✓
Automation features that facilitate dynamic policy updates (Dynamic Object Groups VM-monitoring)	✓
Fully customizable reporting	✓
Shared policies and objects (Panorama)	✓
Role-based administration (Individual Device or Panorama)	✓

APP-ID

A foundation for visibility and control in the Palo Alto Networks Security Platform

App-ID uses multiple identification techniques to determine the exact identity of applications traversing your network – irrespective of port, protocol, evasive tactic, or encryption. Identifying the application is the very first task performed by App-ID, providing you with the knowledge and flexibility needed to safely enable applications and secure your organization.

App-ID™ is a patented traffic classification technology that identifies applications traversing the network, irrespective of port, protocol, evasive tactic or encryption (TLS/SSL or SSH).

- Facilitates a more complete understanding of the business value and associated risk of the applications traversing the network.
- Enables creation and enforcement of safe application enablement policies.
- Brings application visibility and control back to the firewall, where it belongs.

As the foundational element of our enterprise security platform, App-ID provides visibility and control over applications – even those that try to evade detection by masquerading as legitimate traffic, hopping ports or sneaking through the firewall using encryption (TLS/SSL or SSH).

In the past, unapproved or non-work-related applications on your network left you with two choices—either block everything in the interest of data security, or enable everything in the interest of business. These choices left little room for compromise.

App-ID enables you to see the applications on your network and learn how they work, their behavioral characteristics, and their relative risk. When used in conjunction with User-ID™, you can see exactly who is using the application based on their identity, not just an IP address. Armed with this information, your security team can use positive security model rules to allow the applications that enable the business, controlling them as needed to improve your security posture.

Firewall Traffic Classification: Applications, not Ports

Stateful inspection, the basis for most of today's firewalls, was created at a time when applications could be controlled using ports and source/destination IPs. The strict adherence to port-based classification and control is foundational and cannot be turned off. Even when augmented by "after the fact" classifiers, applications cannot be effectively controlled.

Palo Alto Networks® recognized that applications had evolved to where they can easily slip through the firewall and chose to develop App-ID, an innovative firewall traffic classification technique that does not rely on any one single element like port or protocol to identify applications. Instead, App-ID uses multiple mechanisms to determine what the application is. The application's identity then becomes the basis for firewall policy. App-ID has been created to be highly extensible and, as applications continue to evolve, application detection mechanisms can be added to App-ID or updated

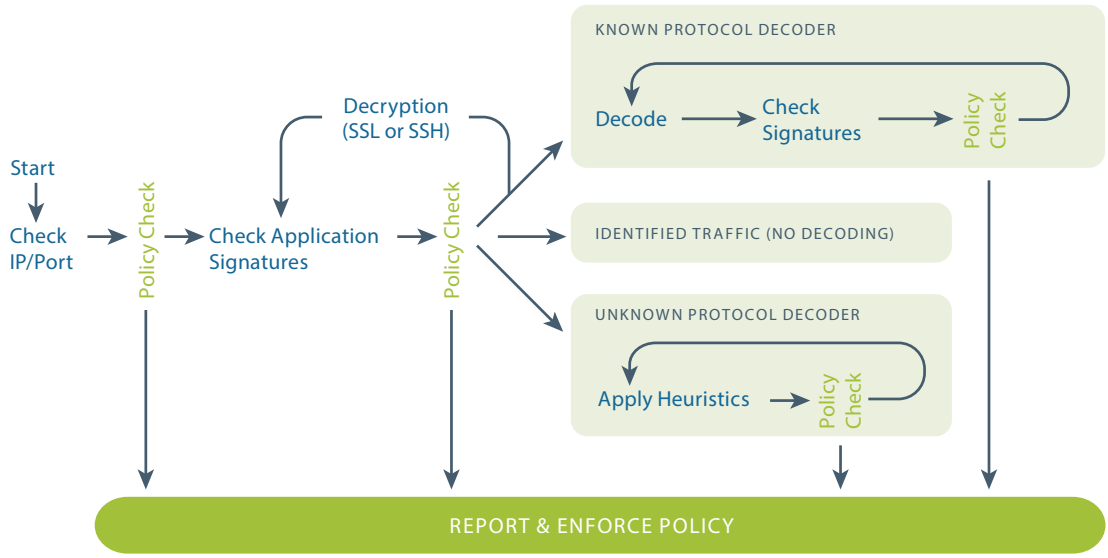


Figure 1: How App-ID classifies traffic.

as a means of keeping pace with the ever-changing application landscape.

App-ID Traffic Classification Technology

Using as many as four different techniques, App-ID determines what the application is as soon as the traffic hits the firewall appliance, irrespective of port, protocol, encryption (TLS/SSL or SSH) or other evasive tactic. The number and order of identification mechanisms used to identify the application will vary depending on the application. The general flow for App-ID is as follows:

- **Application Signatures:** Signatures are used first to look for unique application properties, and related transaction characteristics, to correctly identify the application regardless of the protocol and port being used. The signature also determines if the application is being used on its default port or a non-standard port (for example, RDP across port 80 instead of port 3389, its standard port). If the identified application is allowed by security policy, further analysis of the traffic is done to identify more granular applications as well as scan for threats.
- **TLS/SSL and SSH Decryption:** If App-ID determines that TLS/SSL encryption is in use and a decryption policy is in place, the traffic is decrypted and then passed to other identification mechanisms as needed. If no policy is in place, then decryption is not employed. A similar approach is used with SSH to determine if port forwarding is in use as a means to tunnel traffic over SSH. Such tunneled traffic is identified as *ssh-tunnel* and can be controlled via security policy.
- **Application and Protocol Decoding:** Decoders for known protocols are used to apply additional context-based signatures to detect other applications that may be tunneling inside of the protocol (e.g., Yahoo! Messenger used across HTTP). Decoders validate that the traffic conforms to the protocol specification and they provide support for NAT traversal and opening dynamic pinholes for applications such as VoIP or FTP. Decoders for popular applications are used to identify the individual functions within the application as well (e.g., webex-file-sharing). In addition to identifying applications, decoders identify files and other content that should be scanned for threats or sensitive data.
- **Heuristics:** In certain cases, evasive applications still cannot be detected even through advanced signature and protocol analysis. In those situations, it is necessary to apply additional heuristic, or behavioral analysis to identify certain applications, such as peer-to-peer file sharing or VoIP applications that use proprietary encryption. Heuristic analysis is used as needed, with the other App-ID techniques discussed here, to provide visibility into applications that might otherwise elude positive identification. The actual heuristics used are specific to an application and include checks based on such things as the packet length, session rate, and packet source.

With App-ID as the foundational element of our security platform, your security team can regain visibility into, and control over, the applications traversing your network.

App-ID: Dealing with Custom or Unknown Applications

New applications are added to the App-ID database weekly, yet nearly every network will still have cases where unknown application traffic is detected. There are typically three scenarios where unknown traffic will be detected: a commercially available application unknown to App-ID, an internal custom application, or a threat.

- **Unknown Commercial Applications:** Using visibility tools, you can quickly determine if the traffic is a commercial off-the-shelf (COTS) application. If it is a COTS application, you can capture and submit traffic packets to Palo Alto Networks for App-ID development. The new App-ID is developed, tested, and added to the database for all users in the form of a weekly update.
- **Internal or Custom Applications:** If the application is internal, or custom, you

can create a custom App-ID using a set of available protocol and application decoders. Once the custom App-ID is developed, your internal application is classified and inspected in the same manner as applications with standard App-IDs. Custom App-IDs are managed in a separate database on the device, ensuring they are not impacted by the weekly (commercial) App-ID updates.

- **Threats:** Once the commercial and internal applications have been addressed, the third possible source of unknown traffic is threats. Here too, you can quickly determine the risk levels using the behavioral botnet report or other forensics tools to isolate the characteristics and apply appropriate policy control.

  Even after attempts to identify, some traffic in the system may remain unknown. Because our firewall supports a positive enforcement model, the remaining unknown traffic can be blocked (by default) or allowed but tightly controlled by policy if desired. Alternative offerings (e.g., Intrusion Prevention Systems) are based on negative control and will allow unknown traffic to pass through without providing any semblance of visibility or control.

How App-ID Works: Identifying WebEx

When a user initiates a WebEx session, the initial connection is an encrypted communication. With App-ID, the device sees the traffic and the signatures determine that it is using TLS/SSL. The decryption engine and protocol decoders are then initiated to decrypt the TLS/SSL and detect that it is HTTP traffic. Once the decoder has the HTTP stream, App-ID can apply contextual signatures and detect that the application in use is WebEx. At this point the session traffic becomes known as WebEx traffic by the firewall. Visibility (e.g., ACC in the user interface) and control of the WebEx traffic via security policy are enabled.

If the end user were to initiate the WebEx Desktop Sharing feature, this "mode-shift" from conferencing to remote access would be detected by App-ID. Again, visibility to this specific application function would be provided – and policy control over WebEx Desktop Sharing would be possible (distinct from general WebEx use).

Application Identity: The Heart of Policy Control

Identifying the application is the first step in learning more about the traffic traversing your network. Learning what the application does, the ports it uses, its underlying technology, and its behavioral characteristics is the next step toward making a more informed decision about how to treat the application. Once a complete picture of usage is gained, you can apply policies with a range of responses. Examples include:

- Allow or deny
- Allow but scan for exploits, viruses and other threats
- Allow based on schedule, users or groups
- Control file or sensitive data transfer
- Decrypt and inspect
- Apply traffic shaping through QoS
- Apply policy-based forwarding
- Allow a subset of application functions
- Any combination of the above

With App-ID as the foundational element of our firewalls, you can restore visibility and control over the applications and traffic traversing your network.



Figure 2: Application Function Control – maximize productivity by safely enabling the application itself (Microsoft SharePoint) or individual functions.

Application Function-Level Controls

To many customers, safe application enablement means striking an appropriate security policy balance by enabling some application functions while blocking others. Examples include:

- Allowing Microsoft® SharePoint® Documents, but blocking the use of SharePoint Administration.
- Block Facebook®-mail, -chat, -posting and -apps, but allow Facebook itself, effectively only allowing users to browse Facebook.
- Enable the use of MSN®, but disable the use of MSN-file transfer – or only allow certain file types to be transferred.

Using an application hierarchy that includes the base application and supporting functions, App-ID makes it easy for you to choose which applications to allow overall, while blocking or controlling functions within the application. Figure 2 shows SharePoint as the base application, and the individual functions within.

Controlling Multiple Applications: Dynamic Filters and Groups

There are cases where you may want to control larger groups of applications in bulk, as opposed to controlling them individually. The two mechanisms that address this policy requirement are dynamic filters and application groups.

Dynamic Filters: A dynamic filter is a set of applications that is created based on any combination of the filter criteria: category, subcategory, behavioral characteristic, underlying technology or risk factor. Security policies (e.g. deny, allow, scan) can be applied to dynamic filters. The security policy is then enforced for application traffic that matches the filter criteria.

As new App-IDs are introduced and delivered to the firewall via weekly updates, dynamic filters are automatically updated for those applications that meet the filter criteria. This helps minimize administrative effort associated with security policy management.

Figure 3 below shows a snapshot view of Palo Alto Networks online application database. Here you can browse the current database of App-IDs, including an interactive view of applications based on the same criteria that can be used in dynamic filters.

- Application Groups: An application group is defined as a static list of applications. An example would be a group of remote management applications such as RDP, Telnet, and SSH. In a typical organizational scenario, each of these applications is used by support and IT personnel, yet employees who fall outside of these groups are also known to use them as a means to access their home networks. An application group can be created, with an associated security policy that allows use only by support and IT personnel (supported by User-ID). As new employees join the organization, they need only be added to the appropriate directory group. No updates are needed to the security policy itself.



Figure 3: Browse up-to-date application research and analysis at the Palo Alto Networks Application Research Center, https://applipedia.paloaltonetworks.com

Expanding the List of Applications

The list of App-IDs is expanded weekly with 3 to 5 new applications typically added based on input from customers, partners, and market trends. When you find unidentified applications on your network, you can capture the traffic and then submit the information for App-ID development. Once a new App-ID is developed and tested, it is added to the list as part of the weekly content updates.

Summary

App-ID is a powerful and differentiated core capability of the Palo Alto Networks Next-Generation Firewall, enabling advanced visibility and granular control of traffic in your network. With this visibility and control, you can evaluate what is taking place in your environment, and then define policies that: ensure appropriate use, reduce attack surface, and stop threats. Ultimately, App-ID is a key foundational element in enabling superior risk management for your organization.

4401 Great America Parkway
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com

© 2015 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.

pan-tb-app-id-092115



Advanced Threat Prevention Deployment

Tech Note

PAN-OS 7.0

Version 3.0 - Revision A ©2015, Palo Alto Networks, Inc. www.paloaltonetworks.com

Contents

Summary....4
The attack Kill-Chain....5
Zero Trust Network Security....5
Palo Alto Networks Firewall Deployment....6
Palo Alto Networks Advanced Threat Prevention....7
Application Identification (App-ID)....8
- Kill-Chain....9
- Recommendations....9
- Use Case....10
- Observations....11

User Identification (User-ID)....14
- Recommendations....14
- Use Case....14
- Observations....15

URL Filtering....17
- Kill-Chain....17
- Recommendations....18
- Use Case....18
- Observations....20

Vulnerability Protection....22
- Kill-Chain....23
- Recommendations....23
- Use Case....24
- Observations....24

Anti-Spyware....26
- Kill-Chain....27
- Recommendations....28
- Use Case....28
- Observations....28

Antivirus....29
- Kill-Chain....29
- Recommendations....30
- Use Case....30

©2015, Palo Alto Networks, Inc.

Observations....30
File Blocking....32
Kill-Chain....32
Recommendations....32
Use Case....32
Observations....33
WildFire Advanced Malware Prevention....35
Kill-Chain....35
Recommendations....35
Use Case....36
Observations....38
DoS Protection....41
Kill-Chain....41
Recommendations....42
Use Case....43
Observations....44
Zone Protection....46
Kill-Chain....46
Recommendations....46
Use Case....47
Observations....47
Automated Correlation Engine....49
Recommendations....49
Observations....49
Behavioral Botnet Report....50
Recommendations....50
Observations....51
Appendix A: Performance....52
Appendix B: Default Action for Threat Prevention Signatures....53
Appendix C: Guidelines for the Vulnerability Protection Analysis Phase....54
Appendix D: Slow HTTP Test Output....56
Appendix E: Evasions....61
Revision History....62

©2015, Palo Alto Networks, Inc.

Summary

Over the past years cybersecurity threats have evolved from single-vector attacks to the current incarnations of advanced persistent threat actors and sophisticated multi-stage malware. In order to provide efficient protection, today's security solutions need to understand and address all stages of the malware kill-chain.

Palo Alto Networks' next-generation Firewall offers Application Identification, URL Filtering, Vulnerability Protection, Antivirus, Anti-Spyware, WildFire and DoS Protection technologies capable of detecting and preventing a vast range of threats and malicious communications over the network. Additionally each of these technologies can assist in breaking the malware kill-chain through automated coordination and interaction between the Threat Intelligence Cloud and all layers of defense.

This document provides a general overview and use case information regarding the deployment of security policies and profiles on the Palo Alto Networks next-generation firewall in order to maximize the benefit of an integrated approach to threat prevention.

©2015, Palo Alto Networks, Inc.

The attack Kill-Chain

The term “Kill-Chain” was originally coined by Lockheed Martin and describes the different stages an attacker needs to successfully progress in order to achieve its desired objective. The mitigation of just one of these stages disrupts the chain and the adversary. It is paramount to understand that mitigation requires proactive preventative measures to be in place. An approach that relies mainly on detection and reactive measures will never result in timely disruption of the kill-chain. Palo Alto Networks firewalls offer several key technologies that can assist in disrupting the attack kill-chain. Moreover, with the automation of threat intelligence collection that is provided by WildFire and the Threat Intelligence Cloud, each firewall has the ability to provide automated protection against known and unknown threats through timely updates of its threat signatures.



For a more detailed overview of the kill-chain principle, the following document is available:
http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf

Zero Trust Network Security

The concept of Zero Trust Network Security is gaining adoption as the preferred foundation for network security architectures throughout organizations worldwide. Palo Alto Networks’ next-generation firewall is an obvious choice when it comes to selecting the Zero Trust segmentation platform that will provide the majority of the security functionality needed to deliver on the Zero-Trust operational objectives.

For further details on the zero-trust model, the following document is available:
http://csrc.nist.gov/cyberframework/rfi_comments/040813_forrester_research.pdf

The following page discusses a zero-trust approach to network segmentation:
https://www.paloaltonetworks.com/solutions/initiative/network-segmentation.html

©2015, Palo Alto Networks, Inc.

Palo Alto Networks Firewall Deployment

Palo Alto Networks next-generation firewalls can be deployed in many different ways. The following deployment modes are available:

- Monitor or visibility mode (tap or span)
- Inline transparent mode (virtual wire)
- L2 Bridging
- L3 Routing

Deployment modes are configured on an interface level and can be combined on the same firewall appliance. The threat prevention capabilities discussed in this document can be enabled for each of these deployment modes.

For detailed information on the different deployment modes and an overview of available scenarios, please refer to the following document:

https://live.paloaltonetworks.com/docs/DOC-2561

Palo Alto Networks Advanced Threat Prevention

Palo Alto Networks next-generation firewalls offer a wide range of threat prevention functionalities that interact with our Threat Intelligence Cloud through a feedback loop mechanism to offer automated protection against today's advanced threats. By integrating all these functionalities in a single solution, Palo Alto Networks firewalls offer a unified approach in both security policy design as well as reporting.

Each of these functionalities can protect against specific types of threats and can be configured through firewall rules, security profiles and zone settings. When combined, these functionalities provide an excellent mechanism to disrupt an adversary's kill-chain and at the same time improve visibility and control over the network.

This document will cover each of these functionalities and provide guidance on how to create a next-generation firewall policy by means of a use case. The use case will be built around an example company and focus on reducing the attack surface for all users and systems present on the network.

	Delivery	Exploitation	Installation	Command and Control	Actions on Objectives	
App-ID	Block High-Risk Applications Decrypt SSL			Block C2 on non-standard ports	Prevent data exfiltration and lateral movement	**Coordinated intelligence to detect and block active attacks based on signatures, source and behaviors.**
URL Filtering	Block known malicious sites and links			Block Malware and fast-flux domains		
Vulnerability Protection		Block exploits			Prevent lateral movement	
Anti-Spyware				Block spyware and C2 traffic		
Antivirus			Block malicious files		Prevent lateral movement	
File Blocking			Block unwanted file types and prevent drive-by downloads		Prevent data exfiltration and lateral movement	
WildFire Threat Intelligence	Detect new malicious sites		Detect unknown malware	Detect new C2 traffic		
DoS/Zone Protection		Prevent L3/L4 evasions			Prevent DoS attacks and lateral movement	

Application Identification (App-ID)

In today's networking environments, Application Identification is a critical component towards building an effective security policy. Over the last ten years, use of both the Internet and internal networking has increased dramatically resulting in an abundance of applications now accessing network services as part of day-to-day business operations.

The following trends are observed upon closer inspection of network application behavior:

- Many Internet bound networking applications are designed to access the Internet over port 80 using the HTTP protocol or they use port 80 as a fallback port in case their regular port is blocked.
- SSL encryption is used more and more in order to securely tunnel applications through the firewall.
- Internal applications often use multiple and/or dynamic ports to facilitate communication between end points.
- Using traditional firewalls, it becomes difficult for security administrators to maintain visibility and distinguish between applications based on port or protocol alone.

Applications by themselves can be used as a launch platform for attacks and carry threats inside a company's network. Application identification and control helps in reducing the attack surface for your organization, which is defined as the sum of all possible exploitable targets. All systems, services, applications and users on your network are a potential target for cybercriminals. By creating security policies based on true application identification rather than port or protocol alone, it is possible to reduce the risk your systems, services, applications and users are exposed to. This is done by only allowing those applications that are required for day-to-day business use and by consequently blocking all others.

Here are some examples of common applications as seen in today's company networks:

http-proxy	ms-ds-smb	linkedin-base	snmpv1	snmp-trap	ping	active-directory
ssl	facebook-base	http-video	outlook-web	google-video-base	telnet	dailymotion
web-browsing	citrix	adobe-update	unknown-tcp	radius	livelink	genesys
dns	mssql-db	ms-sms	youtube-base	ms-update	yammer	sharepoint-documents
twitter-base	rtmpt	symantec-av-update	snmp-base	rss	photobucket	webdav
ldap	facebook-social-plugin	google-maps	limelight	google-translate-base	ms-netlogon	yahoo-mail
netbios-ns	snmpv2	msrpc	google-safebrowsing	google-translate-manual	google-picasa	tumblr-base
ftp	kerberos	ntp	hotmail	netlog	ike	ssh
flash	netbios-dg	ms-exchange	blackberry	flickr	gmail-base	xing
google-analytics	sharepoint-base	ipsec-esp	unknown-udp	netbios-ss	soap	dostupest

Some applications will also contain a number of application functions. For example, the application Facebook has application functions such as facebook-base, facebook-chat, facebook-mail, facebook-posting, etc. In order to enforce proper security policies, it may be required to only allow certain applications functions instead of the complete application.

©2015, Palo Alto Networks, Inc.

A full overview of each supported application is available via the Applications view under the Objects tab or via the Applipedia web page at: http://apps.paloaltonetworks.com/applipedia/

More information on App-ID can be found here:
https://www.paloaltonetworks.com/products/technologies/app-id.html

Additional technical documentation on App-ID is available here:
https://www.paloaltonetworks.com/documentation/70/pan-os/pan-os/app-id.html

Unknown applications

One specific use case of application identification is the ability of the App-ID engine to identify and control unknown applications. Unknown applications are applications that are not part of the current App-ID application database and are identified as unknown-udp, unknown-tcp or unknown-p2p.

With regards to App-ID as a threat prevention mechanism, the ability to identify and control unknown applications becomes important as a generic approach to block malware communication. Malware will often use a proprietary protocol to communicate with outside command-and-control servers and in many cases this traffic is identified as unknown.

Kill-Chain

	Delivery	Exploitation	Installation	Command and Control	Actions on Objectives
App-ID	Block High-Risk Applications Decrypt SSL			Block C2 on non-standard ports	Prevent data exfiltration and lateral movement

App-ID can play an important role during several stages of the Kill-Chain. It can disrupt the Delivery stage by blocking user access to applications that are commonly used for targeted delivery such as social media sites, personal webmail and file-sharing applications. During the application identification process, it can also provide SSL decryption to avoid inspection bypass for encrypted content and applications.

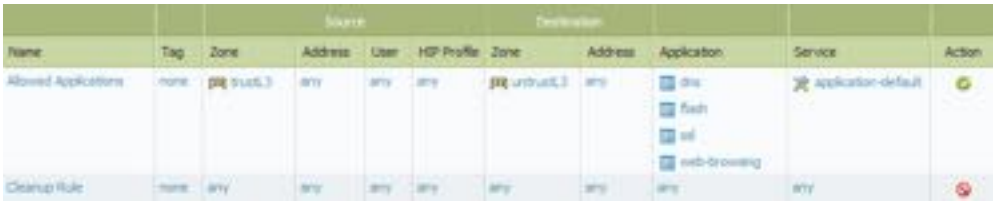
One particular feature of App-ID is its ability to classify unknown traffic as unknown-tcp, unknown-udp or unknown-p2p. This can be leveraged to block connections where the underlying application cannot be identified, as is often the case with custom outbound C2 channels. As such, App-ID can assist in disrupting the kill-chain at the Command-and-Control stage.

Finally, App-ID can reenforce a zero-trust architecture by limiting application usage based on user groups. This can aid in limiting lateral movement capabilities and data exfiltration through specific applications.

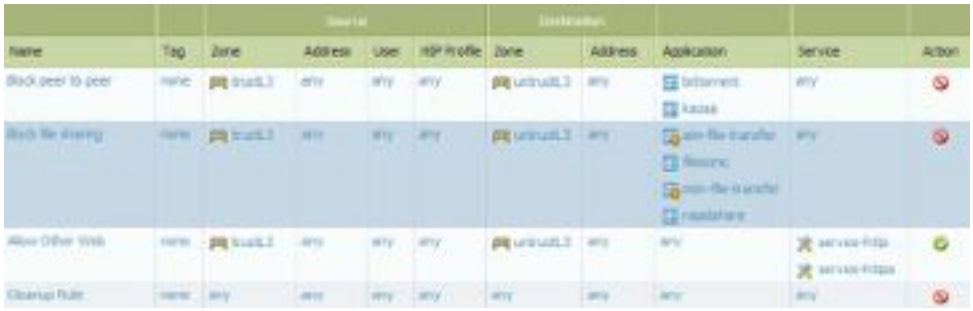
Recommendations

The key principle to apply when defining an application based firewall policy is to build a policy that uses a positive enforcement approach. Positive enforcement means 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. A negative enforcement approach requires you to keep track of any new applications and constantly adapt your policy to block them. This would be a non-stop task and would still leave a high risk gap. A positive enforcement approach only requires you to define an allowed list of applications and have the firewall block everything else with a cleanup rule at the end of the rule base.

Here's an example of the recommended approach showing a positive enforcement firewall policy.



And here's an example of a negative enforcement firewall policy.



As you can see, it will be much easier to maintain a positive enforcement firewall policy over time as you will only need to add those applications required for day-to-day business operations.

When switching from a port-based to an application-based firewall security policy, the most important task is to determine what applications should be allowed for day-to-day business operations. If no well-defined list of allowed applications exists, it is advisable to first configure the Palo Alto Networks firewall as a traditional port-based firewall or to deploy it in monitoring mode in order to obtain a list of all applications running across your network. This list should then be examined in order to maintain the required applications and have any non-desired applications removed from the final policy.

L7 Evasions

App-ID can play an important role in securing your network from L7 evasions. For more information on how to correctly configure a firewall to protect against evasion techniques, please refer to Appendix E: Evasions.

Unknown applications

Given that the App-ID application database now covers most Internet applications, any observation of unknown traffic should be investigated. In order to apply better control over unknown applications and potentially related malware, it is advised to create custom applications or configure application override policies where needed to identify and allow known-good applications that are initially detected as unknown. Combining this approach with a positive enforcement firewall policy will give you strict control over unknown applications, both known-good and bad ones.

More information on handling unknown applications can be found here:
https://live.paloaltonetworks.com/docs/DOC-2007

User-based access control

Besides explicitly defining what applications are allowed, access to applications should also be restricted on a user group level as opposed to using IP addresses in the source fields. The following topic will discuss User-ID in detail.

Use Case

The use case section for each topic will create an example policy that describes how to enable and deploy each feature in an actual security policy.

This first part of the use case will focus on controlling application access for external hosts to public services such as web or mail servers for our example company. In our use case, access policies for two services will be created: web and e-mail.



DMZ Inbound Security Policy

To allow outbound access from the DMZ for update purposes, an additional rule will be created.



DMZ Outbound Security Policy

> ***Note:*** *the above rule is very generic and just serves for clarification purposes. In a real environment, the allowed destinations should be limited to those required for update purposes where possible.*

By defining access by application rather than port or protocol, any traffic to the public services that does not match the application will be denied. This approach reduces the attack surface for these services by disallowing any applications that may have been inadvertently been left enabled on the server, but are not required for proper operation.

Observations

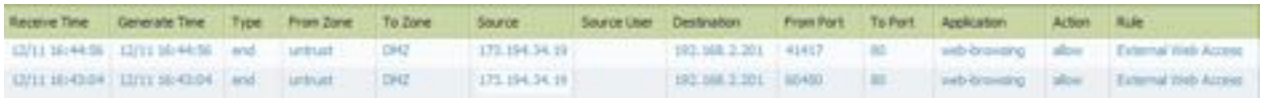
Logs

An overview of all application activity can be consulted via the App Scope Network Monitor, under the Monitor tab. Drill-down analysis can be performed via the Application Command Center (ACC).



Network Monitor

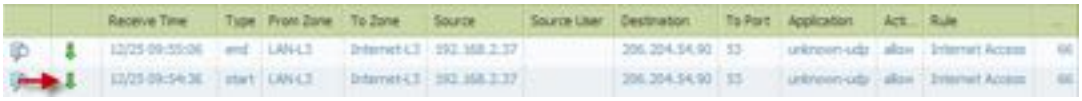
Traffic that matches a rule will generate a log entry in the firewall traffic log if logging is enabled for that rule. A rule can have logging enabled on session end (default), but also on session start.



Traffic Log

Unknown applications

Any new instances of traffic identified as unknown-udp, unknown-tcp or unknown-p2p will trigger a packet capture that is included in the corresponding traffic log. These packet captures may provide an initial indication about the type of traffic.



Traffic Log

Application dependency

Certain applications will depend on other applications for proper operation. In order to allow applications, it is required that the parent applications are allowed as well. The most common example of such application dependency is the reliance of web-based applications on applications ssl and web-browsing. Within one particular application it is also possible that applications functions depend on a parent application. For example, the application function facebook-posting depends on the parent application facebook-base.

The application web-browsing is a very generic application that encompasses any http-based traffic that is not recognized as a specific application. If the Palo Alto Networks firewall is positioned in between host systems and a web-proxy, any proxied http traffic is classified as application http-proxy instead of web-browsing.

As of 5.0, application dependency has been removed for applications that can be identified within a pre-determined point in the session, and use any of the following protocols: HTTP, SSL, MSRPC, RPC, t.120, RTSP, RTMP, and NETBIOS-SS. Custom applications based on HTTP, SSL, MS-RPC, or RTSP can also be allowed in security policy without explicitly allowing a parent application. For example, if you want to allow Java software updates, which use HTTP, you no longer have to allow web-browsing.

Should you want to allow only specific applications that still have an explicit dependency on web-browsing and at the same time block all or selected generic web-browsing, you can apply a URL Filtering profile for this purpose. For more information on URL filtering, see the dedicated topic on this.

User Identification (User-ID)

User Identification allows you to create policies and perform reporting based on users and groups rather than individual IP addresses. The Palo Alto Networks next-generation firewall supports the following enterprise directory services:

- Active Directory
- LDAP
- eDirectory

In addition, there is a User-ID agent for Citrix and Microsoft Terminal Services environments and it is possible to feed 3rd party authentication information to a syslog listener on the firewall or User-ID agent. This allows integration with other network components that have already authenticated a user, e.g. wireless access solutions.

When using Active Directory, User-ID can be performed transparent to the end-user by deploying one or more User-ID agents that can monitor user logon events on domain controllers and Exchange servers and perform IP-address-to-user name mapping. This information is then passed on to the firewall(s) that will perform user-to-group mapping and match user IDs and/or user groups to the right policy rules.

More information on User-ID can be found here:
https://www.paloaltonetworks.com/products/technologies/user-id.html

Additional technical documentation on User-ID is available here:
https://www.paloaltonetworks.com/documentation/70/pan-os/pan-os/user-id.html
https://live.paloaltonetworks.com/docs/DOC-6591

Recommendations

When creating policies based on users and groups, it is advised that you make separate rules per user or user group rather than combine users or groups into one rule, even if they have the same access privileges. This approach will make overall access management easier for each group. It will also provide more rule-based filtering possibilities when building custom reports.

In case the User-ID agent is used to provide IP-address-to-user name mapping, it is advised that you deploy multiple agent instances for redundancy purposes. In addition, you can define a Captive Portal rule for any users on the network that are not authenticated transparently, e.g. because their host machine is not part of the AD domain.

Use Case

In this particular part of the security policy design, the following list of applications required per user group has been defined:

	Sales	Marketing	IT	Human Resources	Production	Management
web-browsing	•	•	•	•		•
linkedin	•	•		•		•
facebook-base		•		•		
facebook-posting		•				
twitter		•				
google-analytics		•				

salesforce	•					•
webex	•					•
google-maps	•					
adobe-update			•			
paloalto-update			•			
ms-update			•			
symantec-av-update			•			

Translating this allow-list into a firewall security policy results in the following rule base:



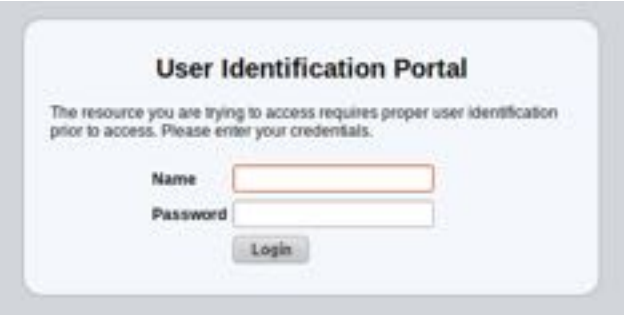
Internal Users Security Policy

The policy as shown above will reduce the attack surface for these user groups and at the same time improve workforce performance by blocking access to non-work related applications.

Observations

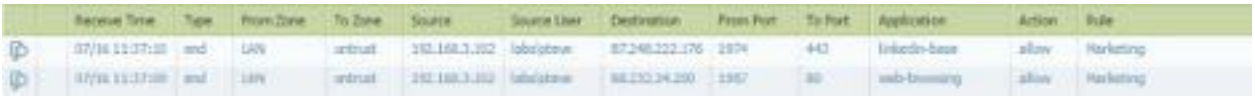
User identification

Most User-ID mechanisms can identify a user in a transparent way meaning the user will not have to explicitly re-authenticate against the firewall. In the case where none of these mechanisms apply, a Captive Portal page can be presented to the user.



Logs

Log entries can be consulted from the Monitor tab, under Logs -> Traffic. As you can see from the screenshot below, the user name will be present in the log entry.



Traffic Log

User Notification

When access to a web-based application is blocked, the user will receive a notification in the browser window. The message and page layout can be customized.



URL Filtering

The URL Filtering feature can block access to specific URLs, web sites and web site categories, or generate an alert when URLs are accessed. You can also define a “block list” of URLs that are always blocked (or generate alerts) and an “allow list” of URLs that are always allowed.

Once firewall access rules have been enabled, a URL Filtering profile can be applied to those rules that allow web access for internal hosts and users in order to further reduce the overall attack surface of your company. In its most basic form this can be done by blocking access to those sites classified as being undeniably malicious. Other possibilities include blocking access to those web categories or sites that pose an increased risk because they serve a large audience and as such are a favorite source platform for malware propagation by cybercriminals. Examples include file-sharing sites, user forums or social media sites.

URL Filtering can be enabled either by placing categories directly in the security rule or by defining and enabling a URL Filtering profile per rule.

There is one default profile available for URL Filtering, which blocks access to the following categories:

- Abused-drugs
- Adult
- Gambling
- Hacking
- Malware
- Phishing
- Questionable
- Weapons

Custom profiles can be created to filter categories following the company’s security or acceptable-use policies.

More information on URL filtering can be found here:
https://www.paloaltonetworks.com/products/features/url-filtering.html

Additional technical documentation on URL Filtering is available here:
https://www.paloaltonetworks.com/documentation/70/pan-os/pan-os/url-filtering.html

Kill-Chain

	Delivery	Exploitation	Installation	Command and Control	Actions on Objectives
URL Filtering	Block known malicious sites and links			Block Malware and fast-flux domains	

URL Filtering can be used to reduce the attack surface for employees as well as disrupt the malware kill-chain at the Delivery stage by blocking access to malicious websites and links pointing to malware deliverables. Additionally, URL Filtering can prevent access to websites, pages and domains that serve as a C2 servers.

In the case of PAN-DB, the malware category is continuously updated with malicious URLs obtained during WildFire Analysis of unknown malware. This provides near real-time protection for URL Filtering subscribers against fast-flux C2 servers.

Recommendations

Depending on your companies risk profile and security policy, you could consider either the Block or Continue action, besides Allow or Alert. The advantage of the Continue action is that it still allows access after an employee clicks a Continue button, while at the same time this action will block any unwanted or obfuscated visits to particular categories or unknown URLs. This offers great protection against drive-by exploits where a link is visited without the employee being aware of it.

The most obvious categories to block are *Malware* and *Phishing*, although other categories can be considered as well if they are not required for your business' daily operations. Additionally, the *Unknown* category can also be given a different action than allow or alert to reduce potential risk from unknown URLs or web-sites.

Container Pages

The default setting for a new URL Filtering profile is to only log container pages. This will create a log entry only for those URIs where the requested page or file matches certain mime-types. The default set includes the following mime-types:

- application/pdf
- application/soap+xml
- application/xhtml+xml
- text/html
- text/plain
- text/xml

When container page logging is enabled, there may not always be a correlated URL log entry for threats detected by Antivirus or Vulnerability Protection. If you want to maintain full logging for any URL that was requested, it is advised to disable container page logging. Alternatively, you can add additional container page mime-types under Device -> Setup -> Content-ID -> Content-ID features -> Container Pages.

Use Case

In our use case, we will define a URL filtering policy based on the following requirements:

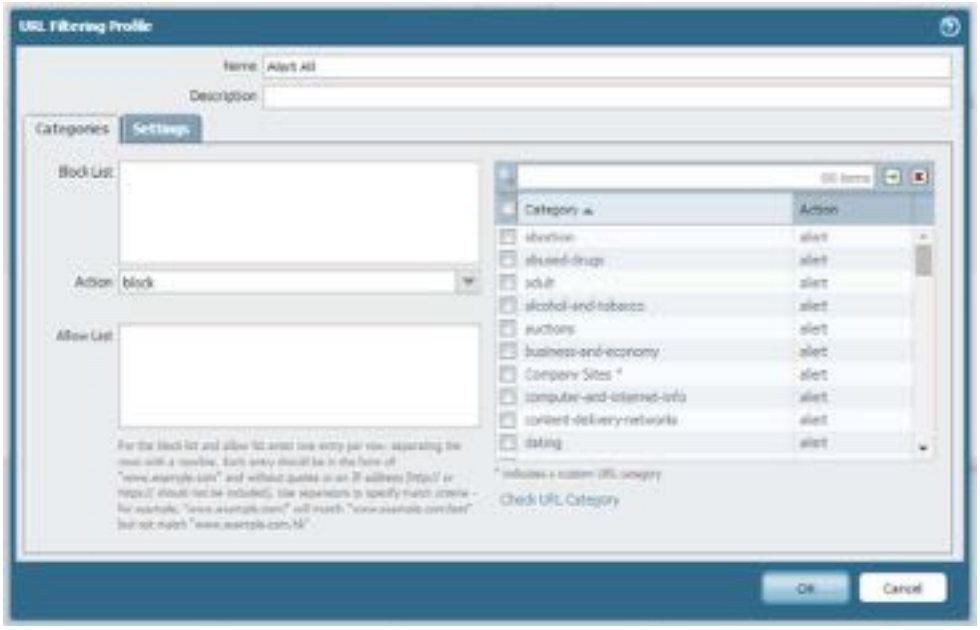
1. Access to the following external sites should be blocked for everyone using the company's Internet link. This includes employees and guests.
    - Hacking
    - Malware
    - Peer-to-peer
    - Phishing
    - Proxy-avoidance-and-anonymizers
    - Questionable
2. Access to any other external site should be monitored and alerted on for known users, but not blocked.
3. Access to any other external site should be blocked for systems where the user is not identified.
4. Access to the company web site should be allowed for everyone, both internal and external. No URL filtering should be performed on this type of traffic.

The first requirement can be accomplished by creating a firewall block rule for these categories. By placing this rule before any other rules that allow outbound access you can assure that no person or system behind the firewall can access any of these categories.



Restricted URL Categories Security Policy

The second requirement can be accomplished by creating a URL Filtering profile that alerts on all categories and applying this profile to the access rule that is already in place for each user group. One additional firewall rule will be created to match any remaining authenticated users that do not match the existing rules.

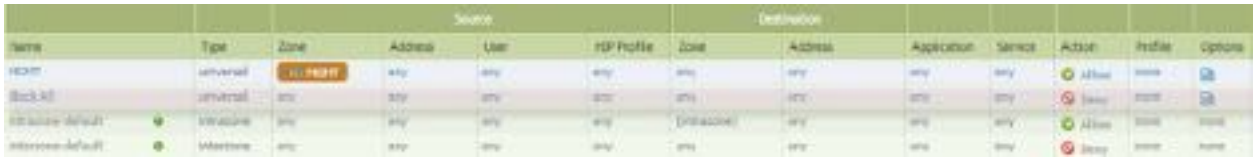


***Note**: As long as the user group access rules are positioned below the generic 'Restricted Sites' block rule there is no need to enable blocking for these categories in the Alert All URL Filtering profile. If you want to have additional safety built-in or if you do not wish to use a general rule for these 'Restricted Sites' as described under the first requirement, then you can enable blocking for those categories in the URL profile as well.*



Security Policy with URL Filtering profile enabled

The third requirement is accomplished by adding a general block rule that blocks any traffic not explicitly allowed. This rule is then placed at the bottom of the rule base.



Security Policy with Block All rule configured

> ***Note:*** *A "Block All" rule overrides the default intrazone and interzone rules and can potentially block intrazone connections that are otherwise allowed. Make sure you add the necessary rules above the Block All rule to explicitly allow this traffic should this be the case. As of PAN-OS 6.1 the default intrazone and interzone rules are editable and can be configured to also log traffic. As such, they may offer a good alternative to a more generic Block All rule without changing the intrazone default behavior.*

The fourth requirement is accomplished by the 'External Web Access' rule that was placed at the top of the rule base under the Application Identification section of this use case.

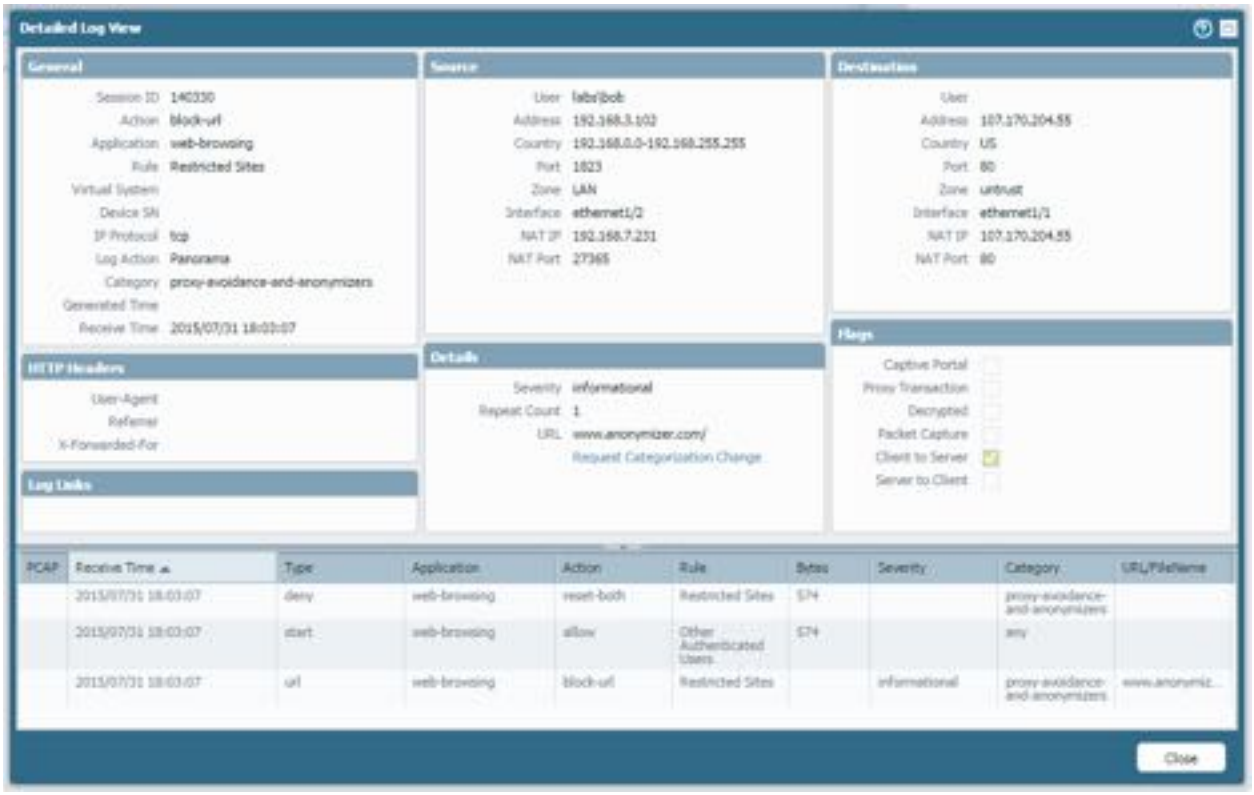
Observations

Logs

URL Filtering log entries can be consulted via the URL Filtering log view, under the Monitor tab. The detailed view for a particular URL will also reference related logs associated with that URL log entry.



URL Filtering Log



URL Filtering Log – Detailed Log View

User Notification

When access to a web site is blocked, the user will receive a notification in the browser window. The message and page layout can be customized.



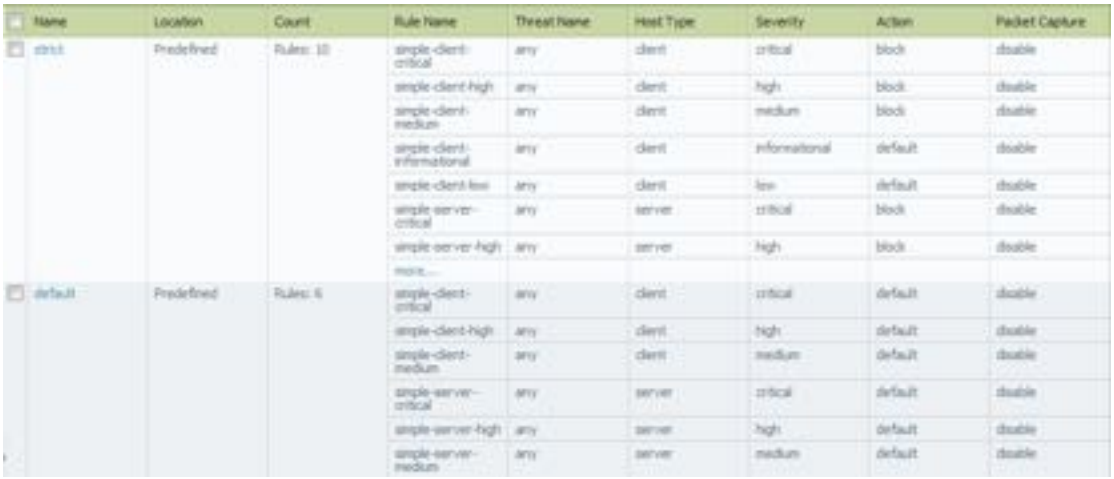
Vulnerability Protection

The Vulnerability Protection feature detects and prevents network-borne attacks against vulnerabilities on client and server systems. Vulnerabilities can be system and service specific or generic and are not bound to a specific port, but to a protocol or application.

Predefined Vulnerability Protection Profiles

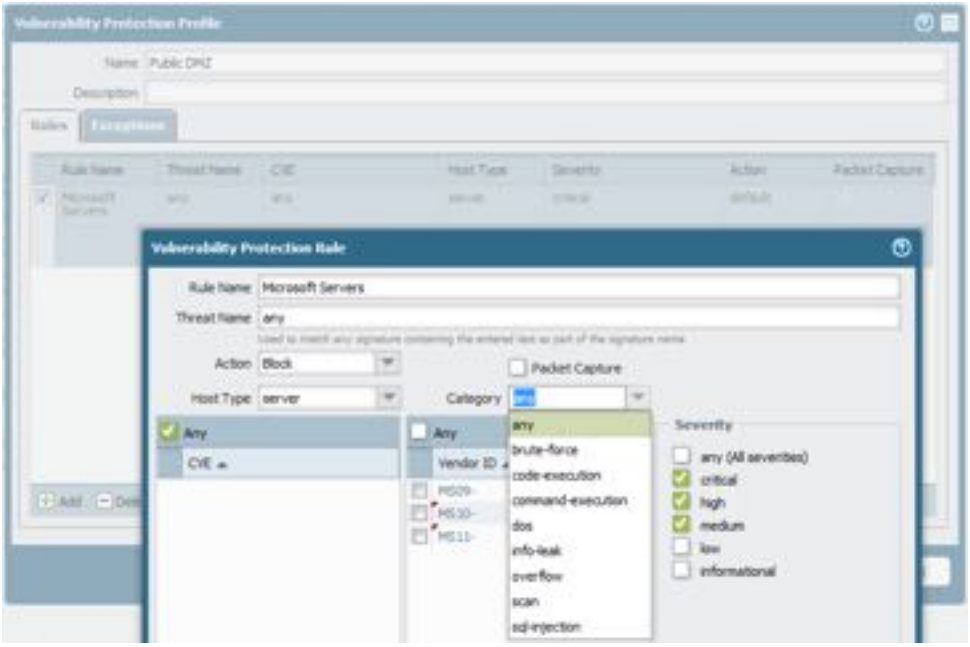
Currently there are two predefined profiles available for the Vulnerability Protection feature:

- The 'default' profile applies the default action to all client and server critical, high, and medium severity vulnerabilities. It does not detect low and informational vulnerability protection events.
- The 'strict' profile applies the block response to all client and server critical, high and medium severity spyware events and uses the default action for low and informational vulnerability protection events.



Custom Vulnerability Protection Profiles

In addition to the predefined Vulnerability Protection profiles, you can create custom profiles tailored to the environment you want to protect. A custom profile can contain one or more rules and exceptions that define which vulnerability protection signatures to include.



Custom profiles also allow you to enable packet captures of matching traffic. This can be used for evidence gathering or troubleshooting purposes.

More information on Vulnerability Protection can be found here:
https://www.paloaltonetworks.com/products/features/ips.html

Additional technical documentation on Threat Prevention is available here:
https://www.paloaltonetworks.com/documentation/70/pan-os/pan-os/threat-prevention.html

Kill-Chain

	Delivery	Exploitation	Installation	Command and Control	Actions on Objectives
Vulnerability Protection		Block exploits			Prevent lateral movement

Vulnerability Protection can assist in breaking the attack kill-chain at the Exploitation stage by preventing client and server vulnerabilities from being exploited. It can also provide the same protection during lateral movement attempts where vulnerabilities on internal hosts can provide a means to compromise additional systems.

Recommendations

When deploying vulnerability protection, special care should be taken to avoid a negative impact on the protected traffic. While vulnerability protection signatures are developed with great care and are submitted to extensive regression tests, some of the signatures are generic in nature and can trigger on traffic coming from misconfigured services or faulty applications. This is especially true for applications that have been custom-built such as in-house developed web applications. Because of this, it is generally not a good idea to simply turn on blocking for large groups of signatures without prior examination of those signatures and the potential impact they may have on the network.

If time and circumstances permit, it is advised to include an analysis phase in the vulnerability protection deployment timeline. In particular for environments where service availability is critical, such a phase will be required to assure proper functionality of the infrastructure once vulnerability protection is made fully operational.

©2015, Palo Alto Networks, Inc.

In general, it is advised to start with a profile that uses the default action for each signature. Alternatively, it is possible to deploy a custom vulnerability protection profile in alert-only or monitoring mode first, to obtain a clear picture on how blocking-mode may affect the infrastructure. Such a profile will have the action set to 'alert' for each signature. For a more detailed description of the steps involved to deploy an alert-only profile during an analysis phase, see Appendix C: Guidelines for the Vulnerability Protection Analysis Phase.

Use Case

In our use case, we will enable Vulnerability Protection for all traffic generated by internal users as well as outbound traffic for the DMZ. A clone of the default Vulnerability Protection profile will be applied to each rule.



Internal Users Security Policy with Vulnerability Protection profile enabled



DMZ Security Policy with Vulnerability Protection profile enabled

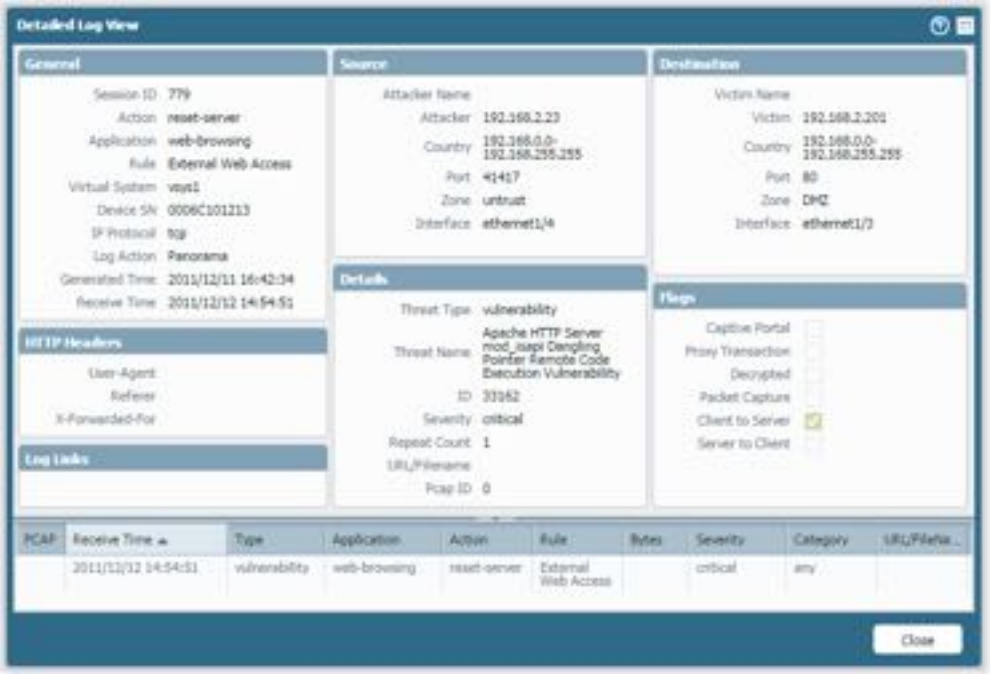
Observations

Logs

Details for Vulnerability Protection alerts can be consulted via the Threat log view, under the Monitor tab. The detailed view for a particular threat will also reference related logs associated with the threat log entry.

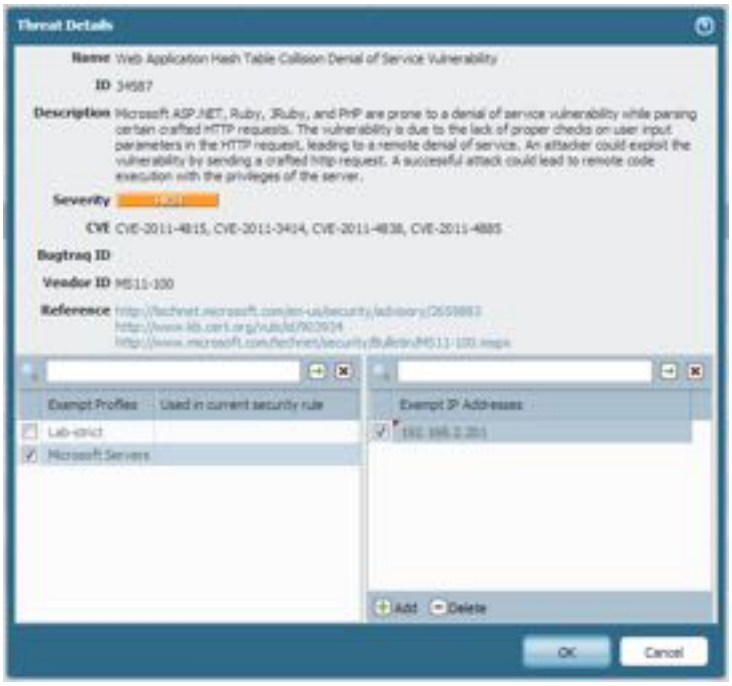


Threat Log



Threat Log – Detailed Log View

Starting with firmware 5.0, it is possible to create IP address-based exceptions for Vulnerability Protection events. This can be accomplished by clicking on the threat name in the Threat log view and adding the exempt IP addresses for selected exempt profiles in the pop up window.



This will create an exception for this threat/IP address combination in the selected Vulnerability Protection profiles.

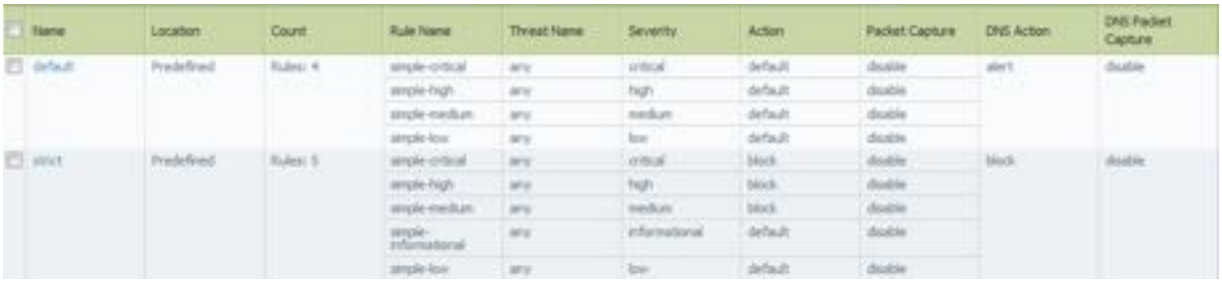
Anti-Spyware

The Anti-Spyware feature detects and prevents spyware and malware network communications. Unlike the Antivirus feature, the Anti-Spyware feature is not limited to specific protocols and can detect any type of phone-home communication.

Predefined Anti-Spyware Profiles

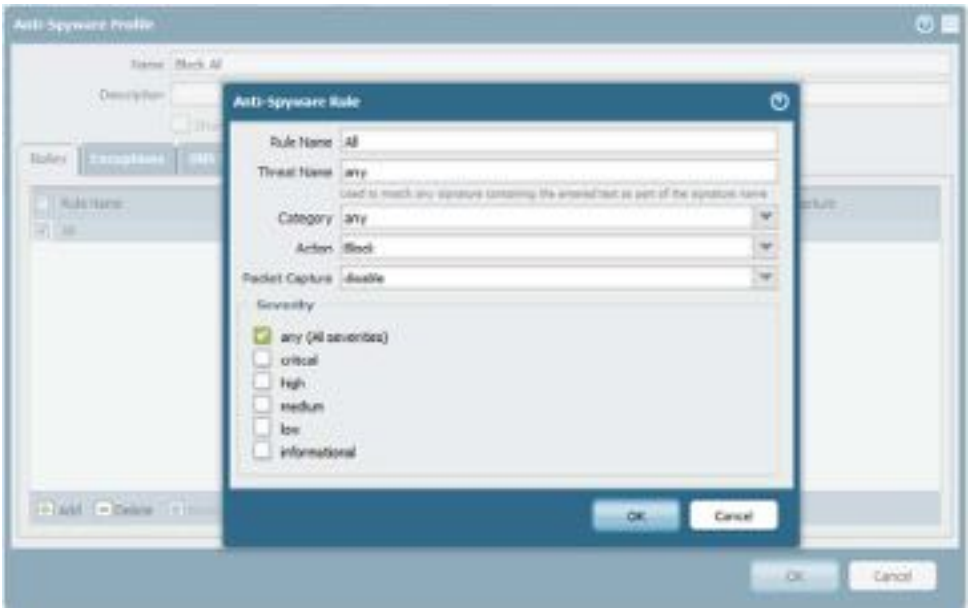
There are two predefined profiles available for the Anti-Spyware feature:

- The 'default' profile applies the default action to all critical, high, medium and low severity spyware events. It does not detect informational severity spyware events.
  As of 5.0, the 'default' profile applies the alert action to DNS queries that trigger DNS signatures for malware domains.
- The 'strict' profile applies the block response to critical, high and medium severity spyware events and uses the default action for low and informational severity spyware events.
  As of 5.0, the 'strict' profile applies the block action to DNS queries that trigger DNS signatures for malware domains.



Custom Anti-Spyware Profiles

In addition to the predefined Anti-Spyware profiles, you can create custom profiles tailored to the environment you want to protect. A custom profile can contain one or more rules and exceptions that define which Anti-Spyware signatures to include.



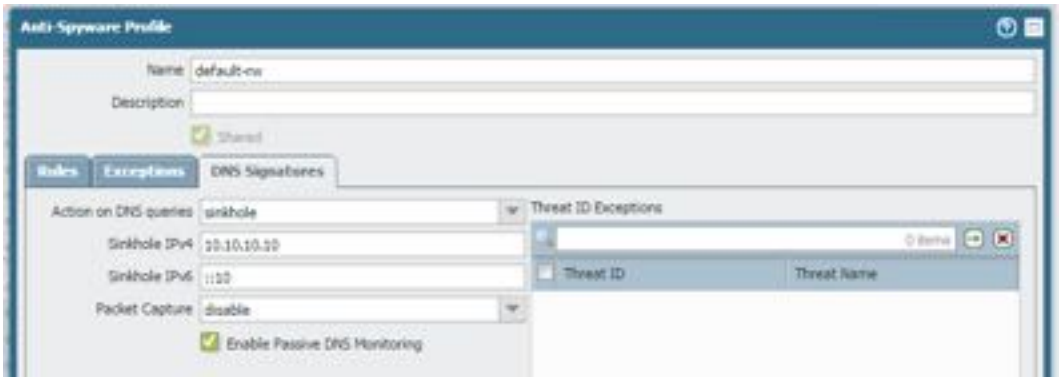
Custom profiles also allow you to enable packet captures of matching traffic. This can be used for evidence gathering or troubleshooting purposes.

DNS Sinkholing

The Anti-Spyware profile offers the ability to apply a sinkhole action to DNS queries made to malicious domains. The DNS sinkhole action provides administrators with a method of identifying infected hosts on the network using DNS traffic, even when the firewall is north of a local DNS server (i.e. the firewall cannot see the originator of the DNS query). When a threat prevention license is installed and an anti-spyware profile is enabled in a security profile, the DNS-based signatures will trigger on DNS queries directed at malware domains. 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 queries directed at malicious domains, so that clients attempting to connect to malicious domains (for command-and-control, for example) instead attempt connections to an IP address specified by the administrator. Infected hosts can then be easily identified in the traffic logs because any host that attempts to connect to the sinkhole IP are most likely infected with malware.

Passive DNS

The Anti-Spyware profile also offers the ability to enable 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.



More information on Anti-Spyware can be found here:
https://www.paloaltonetworks.com/products/technologies/content-id.html

Additional technical documentation on Anti-Virus is available here:
https://www.paloaltonetworks.com/documentation/70/pan-os/pan-os/threat-prevention.html

Kill-Chain

	Delivery	Exploitation	Installation	Command and Control	Actions on Objectives
Anti-Spyware				Block spyware and C2 traffic	

The Anti-Spyware feature can be used to disrupt the attack kill-chain at the Command-and-Control stage by preventing compromised hosts to establish a C2 channel to control systems outside the company network.

Recommendations

The predefined profiles will cover the majority of use cases. Because predefined profiles are read-only, it is advised to work with a custom clone of the predefined profiles, so exceptions can be defined when needed. In environments where blocking traffic is not allowed, a custom profile can also be used to only alert on spyware events.

Use Case

In our use case, we will enable Anti-Spyware for all traffic generated by internal users. A clone of the strict Anti-Spyware profile will be applied to each rule.



Internal Users Security Policy with Anti-Spyware profile enabled

In addition, it is advisable to also turn on Anti-spyware for any security rules that allow outbound traffic from the DMZ.

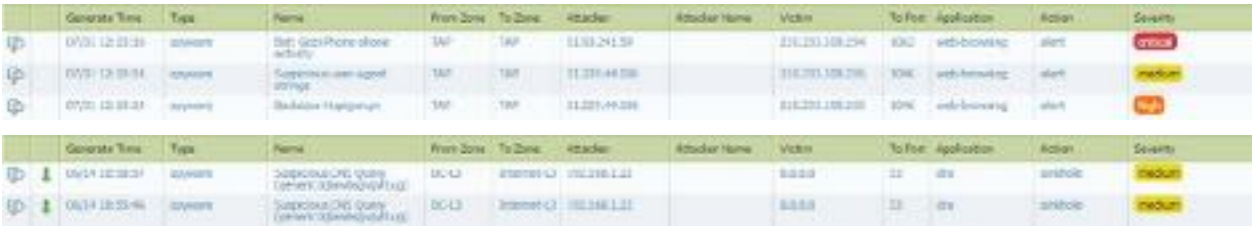


DMZ Security Policy with Anti-Spyware profile enabled

Observations

<u>Logs</u>

Details for Anti-Spyware alerts can be consulted via the Threat log view, under the Monitor tab. The detailed view for a particular threat will also reference related logs associated with the threat log entry.



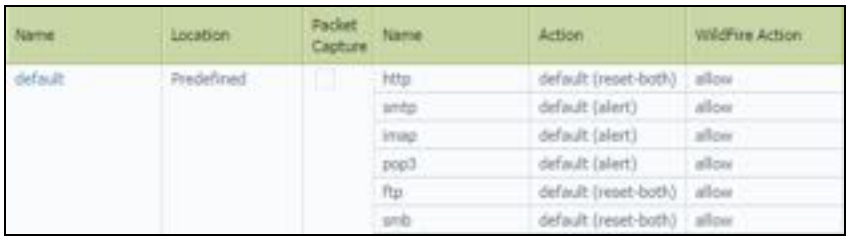
Antivirus

The Antivirus feature detects and prevents viruses and malware from being transferred over the following protocols:

- HTTP
- FTP
- SMTP
- IMAP
- POP3
- SMB

Files transferred by any application that uses any of the above protocols can be inspected by the Antivirus feature. Inspection is done through stream-based analysis, which means files are not cached or stored in their entirety on the firewall, but analyzed in real-time as they pass through the firewall.

There is one predefined profile - named 'default' - available for the Antivirus profile. This profile has the default action configured for each protocol. The default action differs for each protocol and follows the most up-to-date recommendation from Palo Alto Networks. The default action for each protocol is listed below.



Name	Location	Packet Capture	Name	Action	WildFire Action
default	Predefined		http	default (reset-both)	allow
			smtp	default (alert)	allow
			imap	default (alert)	allow
			pop3	default (alert)	allow
			ftp	default (reset-both)	allow
			smb	default (reset-both)	allow

> ***Note**: The reason why SMTP, POP3 and IMAP have the default action set to ALERT is because in most cases there is already a dedicated Antivirus gateway solution in place for these protocols. Specifically for POP3 and IMAP, it is not possible to clean files or properly terminate an infected file-transfer in-stream without affecting the entire session. This is due to shortcomings in these protocols to deal with this kind of situation.*

Custom profiles can be created and allow you to customize the action for each protocol.

More information on Antivirus can be found here:
https://www.paloaltonetworks.com/products/features/antivirus.html

Additional technical documentation on Anti-Virus is available here:
https://www.paloaltonetworks.com/documentation/70/pan-os/pan-os/threat-prevention.html

Kill-Chain

	Delivery	Exploitation	Installation	Command and Control	Actions on Objectives
Antivirus			Block malicious files		Prevent lateral movement

The Antivirus feature can be used to disrupt the attack kill-chain at the Installation stage by blocking malicious files and exploit tools as they traverse the network. It can also provide the same protection during lateral movement attempts where exploit kits or other tools are copied over to other compromised hosts.

Recommendations

The Antivirus feature can be used to disrupt the attack kill-chain at the Installation stage as well as prevent lateral movement by blocking malicious files and exploit tools as they traverse the network.

The default Antivirus profile settings can be used in most situations where dedicated SMTP, POP3 and/or IMAP-scanning solutions are also present. Because predefined profiles are read-only, it is advised to work with a custom clone of the predefined profiles, so exceptions can be defined when needed.

If no dedicated Antivirus gateway solution is present for SMTP, it is possible to define a custom Antivirus profile and apply the reset-both action to infected attachments. In such case, a 541 response will be sent back to the sending SMTP server to prevent it from resending the blocked message.

Typically, Antivirus signatures have an extremely low false positive rate. Should a false positive occur, it is possible to exclude specific Threat IDs from detection by defining Virus Exceptions in the Antivirus profile. For the same purpose, certain applications can also be excluded from being inspected. In such cases, it is best practice to create a specific profile and apply this profile to only the rules that are affected.

Use Case

In our use case, we will enable Antivirus for all traffic generated by internal users. A clone of the default Antivirus profile will be applied to each rule.



Internal Users Security Policy with Antivirus profile enabled

Additionally we will also enable Antivirus for outbound update traffic from the DMZ.



DMZ Security Policy with Antivirus profile enabled

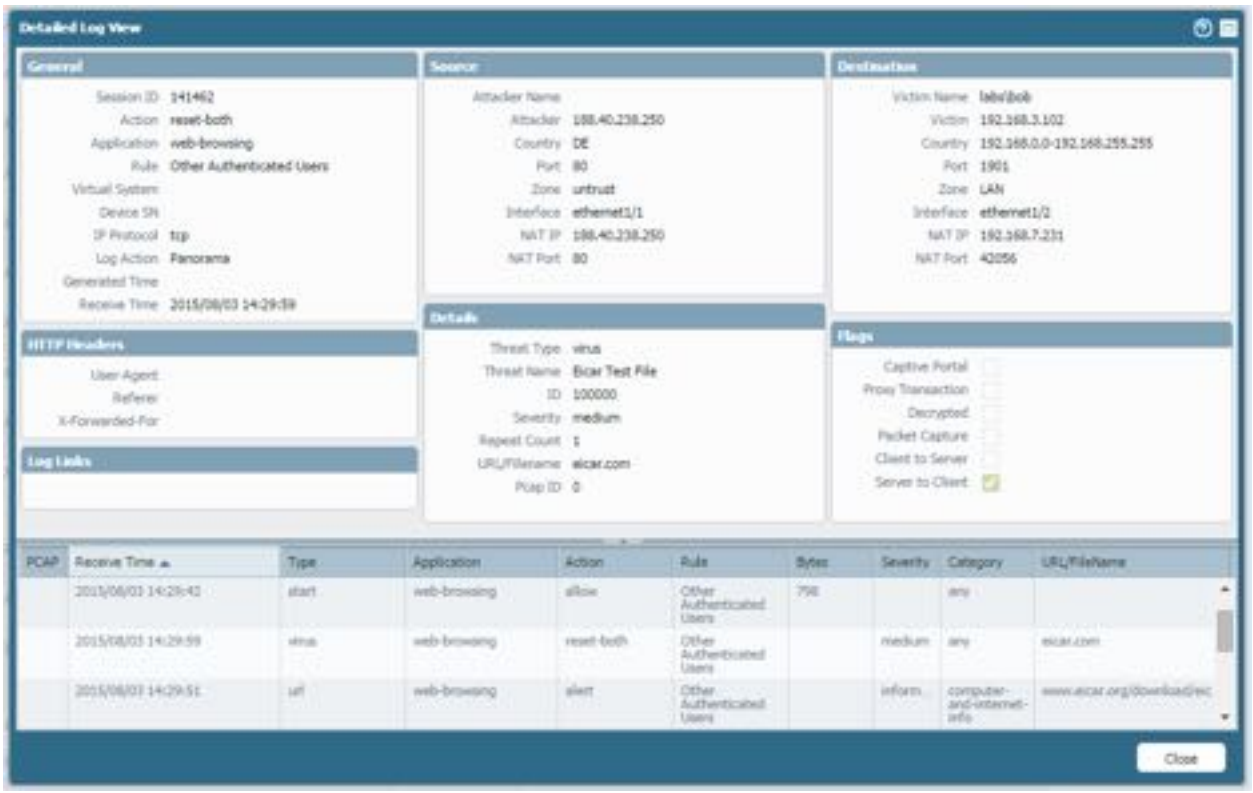
Observations

<u>Logs</u>

Details for each Antivirus alert can be consulted via the Threat log view, under the Monitor tab. The detailed view for a particular threat will also reference related logs associated with the threat log entry.



Threat Log



Threat Log – Detailed Log View

User Notification

Users who are attempting to download an infected file will be presented with a 'Virus Download Blocked' message in the browser. The message and page layout can be customized.



File Blocking

By using File Blocking profiles, it is possible to detect and prevent downloads or uploads of specific file types. As with all other security profiles, file blocking profiles can be enabled on a per rule basis and as such granular control can be applied to file transfers for specific network segments, users and user groups.

Kill-Chain

	Delivery	Exploitation	Installation	Command and Control	Actions on Objectives
File Blocking			Block unwanted file types and prevent drive-by downloads		Prevent data exfiltration and lateral movement

File Blocking profiles can be used to disrupt the attack kill-chain at the Installation stage by preventing downloads or uploads of specific file types such as executables. This can be particularly useful in blocking drive-by download attempts after exploitation where a user is not aware a file is being transferred in the background. File Blocking profiles can also help in preventing data leakage by preventing file uploads to public hosted services other than company approved services.

Recommendations

File blocking can be particularly useful in preventing users from downloading and installing unverified software on company assets and can also prevent drive-by downloads. No predefined file blocking profiles exist, but they are easy to create. If blocking files is too strict for your environment, you should consider using the Continue action instead. This will require a user to explicitly confirm a download and provides a simple but effective protection mechanism against drive-by downloads.

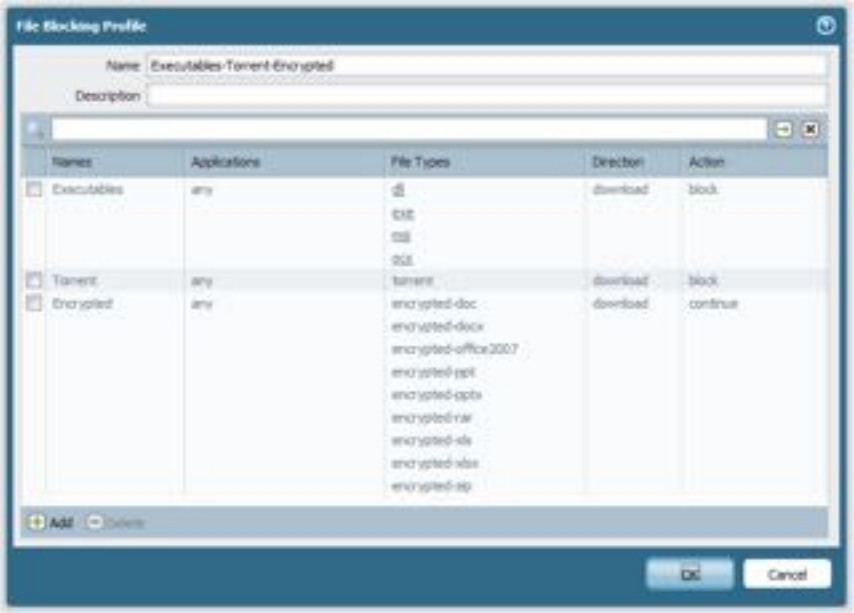
Use Case

In our use case, we will implement the following requirements with regards to file blocking for all user groups except the IT user group.

- Block executables
- Block torrent files
- Warn the user when downloading encrypted files

The file blocking profile looks as follows:

©2015, Palo Alto Networks, Inc.



This profile is then enabled for each user access rule except the IT user group.



Internal Users Security Policy with File Blocking profile enabled

Observations

Logs

File transfers in sessions matching a rule with a file blocking profile enabled will generate a log entry in the Data Filtering log view, under the Monitor tab.



Data Filtering Log

User notification

When a file download is blocked, the user will see a block or continue notification in the browser.


File Download Blocked
Access to the file you were trying to download has been blocked in accordance with company policy. Please contact your system administrator if you believe this is in error.
File name: tiny.exe


File Download Blocked
Access to the file you were trying to download has been blocked in accordance with company policy. Please contact your system administrator if you believe this is in error.
File name: gallery.zip
Please click Continue to download/upload the file.

WildFire Advanced Malware Prevention

Modern malware is at the heart of many of today's most sophisticated network attacks and is increasingly customized to avoid traditional security solutions. Palo Alto Networks has developed an integrated approach that addresses the full malware life cycle, which includes preventing infections, identifying zero-day malware (that is, malware that has not previously been identified by other antivirus vendors) or targeted malware (malware targeting a specific industry or corporation), as well as pinpointing and disrupting active infections.

The Palo Alto Networks WildFire engine exposes zero-day and targeted malware through direct observation in a virtual environment within the WildFire system. The WildFire feature also makes extensive use of Palo Alto Networks App-ID technology by identifying file transfers within all applications, not just email attachments or browser-based file downloads.

The key benefits of the Palo Alto Networks WildFire feature is that it can discover zero-day malware and can quickly generate signatures to protect against future infections of all of the malware it discovers. The firewall can provide instant alerts whenever malware is detected on your network by sending email alerts, syslog alerts, or SNMP traps. This allows you to quickly identify the user who downloaded the malware and eradicate it before it causes extensive damage or propagates to other users. In addition, every signature generated by WildFire is automatically propagated to all Palo Alto Networks firewalls protected with a Threat Prevention, URL Filtering and/or WildFire subscription, which provides automatic protection from malware even if it was not found in your network. Palo Alto Networks is currently discovering and generating signatures for thousands of zero-day malware every week and this number continues to grow.

WildFire is available as both a cloud-based as well as an on-premise solution: the WF-500.

More information on WildFire can be found here:
https://www.paloaltonetworks.com/products/technologies/wildfire.html

Additional technical documentation on WildFire is available here:
https://www.paloaltonetworks.com/documentation/70/wildfire/wf_admin.html

Kill-Chain

	Delivery	Exploitation	Installation	Command and Control	Actions on Objectives
WildFire Threat Intelligence	Detect new malicious sites		Detect unknown malware	Detect new C2 traffic	

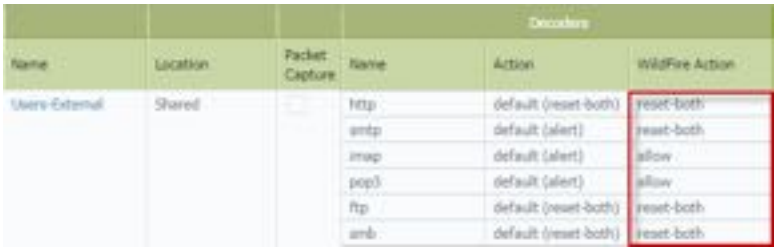
WildFire Analysis results are used to generate signatures for several protection mechanisms. Currently, WildFire can generate file and dns signatures as well as add malicious URLs to the PAN-DB malware category. As such it provides a very useful capability to automate the kill-chain disruption approach for unknown malware.

Recommendations

Due to the proliferation of unknown malware it is highly recommended to enable WildFire Analysis for files coming from locations that are not trusted. WildFire Analysis profiles can be configured to forward files to either the public-cloud WildFire or a private-cloud WF-500. WildFire is capable of analyzing files of type apk, ms-office, jar, flash, email-link and pdf in addition to executable file types (PE). To forward these file types to the WildFire analysis engine, a WildFire Analysis profile needs to be created.

> ***Note**: For pre-7.0 deployments, a file blocking profile with an action of forward or continue-and-forward needs to be created to forward files to WildFire.*

In addition to forwarding files for analysis, malicious file transfers can also be immediately blocked if a WildFire file signature is already available. To achieve this, configure the Antivirus profile with a reset-both action for the WildFire Action for applicable protocols.

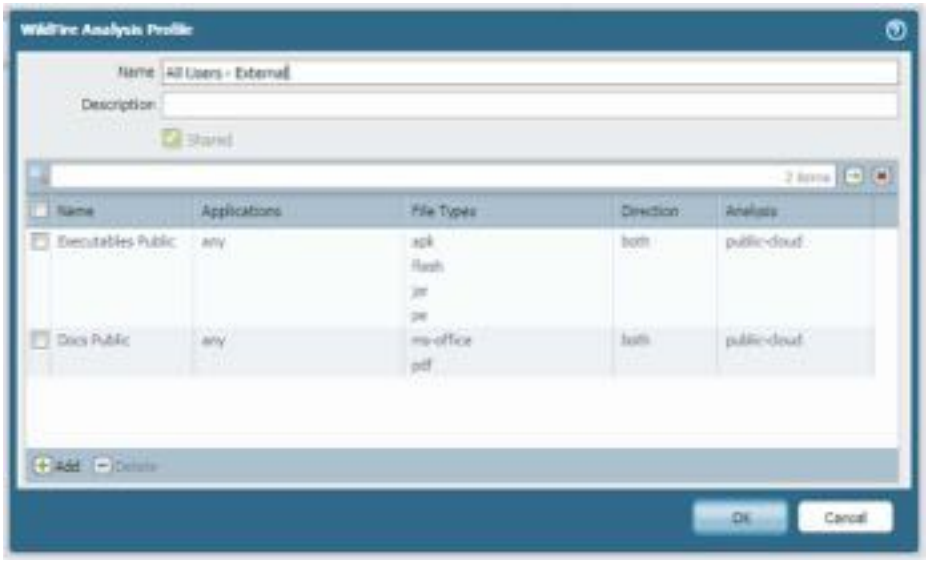


Antivirus profile WildFire Action

Use Case

In our use case, we will enable forwarding of files to the WildFire public cloud for all communication between users and the untrust zone. In addition, all communication between users and the datacenter will be subject to forwarding to the WildFire private cloud (WF-500). Finally, all inbound smtp traffic to the mail-relay in the DMZ will also be subject to forwarding.

Files originating from or destined to a public source are generally allowed to be forwarded to the WildFire public cloud. So the WildFire Analysis profile that will be used to forward files between users and the untrust zone looks as follows:



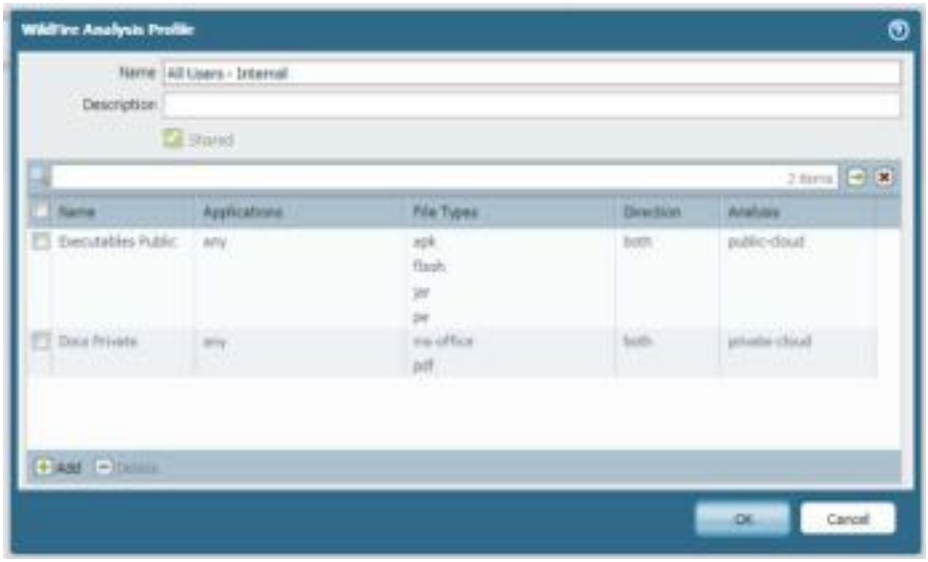
WildFire Analysis profile for user communication to untrust

This profile is enabled for all rules that allow users access to the untrust zone.



Security Policy for user communication to untrust with WildFire Analysis profile enabled

For files that contain sensitive information or are considered classified, there is the option to have these analyzed using the WildFire private cloud (WF-500). In our example use case, it is decided that executable files can be forwarded to the public cloud even if they originate from an internal source, but document files should be analyzed using the private cloud option. The WildFire Analysis profile looks like this:



WildFire Analysis profile for user communication to datacenter



Security Policy for user communication to the datacenter with WildFire Analysis profile enabled

Enabling WildFire Analysis for all files received via email requires a WildFire Analysis profile to be enabled on the rules that allow inbound smtp communication. Additionally, this profile can also be enabled for outbound communication. The WildFire Analysis profile for our use case looks as follows:



WildFire Analysis profile for communication between DMZ and untrust zones

This profile is then enabled on the security rules that allow smtp communication between the DMZ and untrust zones.

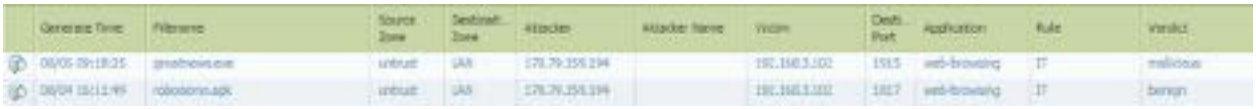


Security Policy for communication between DMZ and untrust zones with WildFire Analysis profile enabled

Observations

Logs

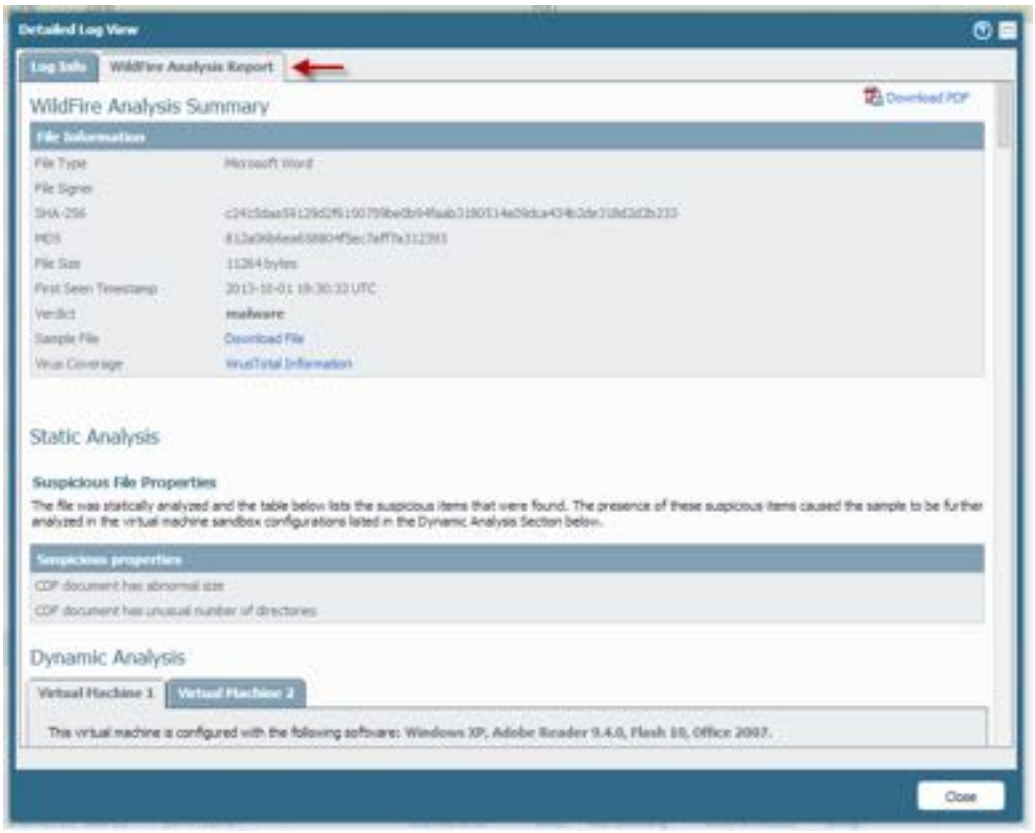
Once a file has been analyzed or verified against the WildFire DB, a log entry will be created in the WildFire Submissions log view, under the Monitor tab.



WildFire Submissions Log

Detailed reporting of the WildFire Analysis results are available through the Detailed Log View.

©2015, Palo Alto Networks, Inc.



The WildFire Analysis report can be downloaded as a pdf from this window and it is also possible to download the original sample file and pcap file of network traffic from the malware sample as seen by WildFire. Furthermore, the report includes dynamic analysis results from all virtual environments against which the sample was executed. Finally, the new report also allows you to report an incorrect verdict for each analysis directly from the report.

Signature Generation

If a file is found to be malicious during WildFire Analysis, protection signatures will be generated and made available as part of the Dynamic Updates mechanism. WildFire file signatures updates are generated every 15 minutes and Anti-Spyware DNS signatures are generated daily. Furthermore, malicious URLs are added to the malware URL Filtering category in PAN-DB within 15 minutes of analysis.

If a file is malicious and a WildFire file signature is available and installed, the file can be blocked by the Antivirus functionality and the firewall will generate a wildfire-virus log entry in the Threat Log.



User notification

When a file download is performed and the file is identified as malicious because it matches a WildFire signature, the file will be blocked according to the settings in the Anti- Virus profile and the user will receive a notification in his browser.


Virus Download Blocked
Download of the virus has been blocked in accordance with company policy. Please contact your system administrator if you believe this is in error.
File name: privacyn.exe

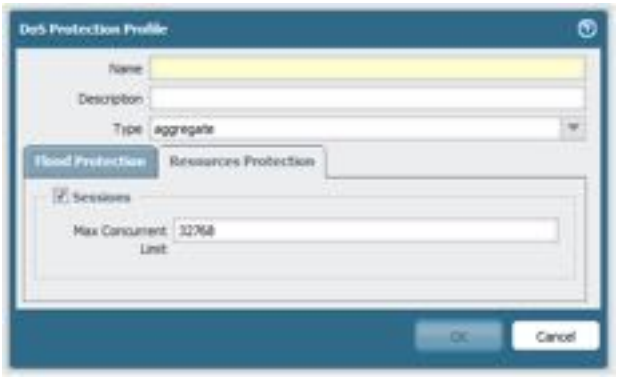
DoS Protection

The DoS Protection feature detects and prevents network-borne Denial-of-Service attacks. Two DoS protection mechanisms are available:

- **Flood Protection–**Detect and prevent flooding attacks where the network is flooded with packets which typically result in too many half-open sessions and/or services being unable to respond to each request. In this type of attack, the source address is often spoofed. Flood Protection can start blocking packets on an aggregate or classified basis as soon as a configurable threshold has been exceeded.
- **Resources Protection–**Detect and prevent session exhaustion attacks. This type of attack is typically performed using a large amount of source hosts (bots) to create as many fully established sessions as possible. It is more difficult to detect as the sessions may be used to send valid-looking requests to the target host. Resources Protection can limit the amount of available sessions on an aggregate or classified basis.

Both mechanisms can be used together in the same DoS Protection profile.





Kill-Chain

	Delivery	Exploitation	Installation	Command and Control	Actions on Objectives
DoS/Zone Protection		Prevent L3/L4 evasions			Prevent DoS attacks and lateral movement

While DoS Protection in itself does not provide any mechanism to disrupt the kill-chain, it can assist in reducing the impact of any cybercriminal action or objective that is aimed at impacting internal or public company services by performing a DoS or resource starvation attack against.

©2015, Palo Alto Networks, Inc.

Recommendations

Defining the correct DoS Protection profile(s) largely depends on what services require protection and what audience will be served. Because each environment is different, a certain amount of up-front analysis will need to be done. Some configuration examples are provided further down this guide.

Regardless of the environment, it is important to pay special attention to the following factors.

- Default threshold values
- Maximal Rate
- Aggregate vs Classified profiles
- SYN-cookie vs Random Early Drop

Default threshold values

The default DoS Protection threshold values do not represent best practice values. Threshold values should be configured based on actual session data for the environment (i.e. zones, hosts) where DoS Protection profiles will be applied. A proper method to obtain this baseline data is to configure Netflow reporting on the firewall to a Netflow analyzer.

> ***Note**: As of 4.0, SYN-cookie is active by default if selected as a protection mechanism with the default Activate Rate setting (Activate Rate = 0).*



SYN-cookie & Maximal Rate default values

Maximal Rate

When the Maximal Rate threshold is exceeded, any further packets that match the DoS Protection rule and classification criteria (in case of a classified profile) will be dropped for the block duration specified. The default value for SYN-cookie is 1.000.000 which will prevent it from being triggered in almost any environment. You should adjust this value to match your environment in the following scenarios:

- To protect against TCP SYN floods where Random Early Drop is used.
- To protect against UDP, ICMP or other IP floods.

- When SYN-cookie is used in a source-ip classification profile and there is never a need for a client to send more SYN pps than the maximal rate value. The SYN-cookie mechanism already provides protection by itself, so the maximal rate value will not provide additional protection to any service except for the firewall's SYN-cookie mechanism itself.

Aggregate vs Classified profiles

When configuring DoS thresholds for flood and resources protection it is important to understand the difference between aggregate and classified profile types.

- **Aggregate–**Applies 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.
- **Classified–**Applies the DoS thresholds configured in the profile to all packets matching the rule and classification criteria (source IP, destination IP or source-and-destination IP). For example, a classified rule with a SYN flood threshold of 10000 packets per second (pps) and source-ip as the classification criterion counts all packets per source IP address that hit that particular DoS rule.

From these definitions, you can already guess that an aggregate profile will be the better choice when protecting against attacks from the Internet. A classified profile is less appropriate here because there may be many clients residing behind a NAT device.

For internal clients not behind a NAT device, a source-ip classification profile may be a good fit from a resource limitation point of view. For example, in an educational environment where internal clients are allowed to run any type of tool (e.g. peer-to-peer) this can be used to limit the amount of concurrent sessions per client to prevent overall session saturation.

SYN-cookie vs Random Early Drop

SYN-cookie is a superior mechanism to counter SYN flood attacks. It is preferred over Random Early Drop for defending against SYN floods. For other traffic, Random Early Drop is the default and only option. Random Early Drop starts randomly dropping packets if the packet rate is between the Activate Rate value and Maximal Rate value. The drop probability increases linearly with the packet rate. If the packet rate exceeds the Maximal Rate value then all excess packets will be dropped.

Use Case

The following examples will provide some ideas on how to implement DoS Protection profiles.

With only one or more web servers to protect, DoS Protection profiles and rules can be very generic. However, it is good practice to already plan for future service additions by making the DoS Protection rules as specific as possible.

Requirements:

- A DoS Protection profile that protects against SYN floods as well as TCP session exhaustion attacks.
- One or more DoS Protection rules that apply the profile to matching traffic.

DoS Protection profile

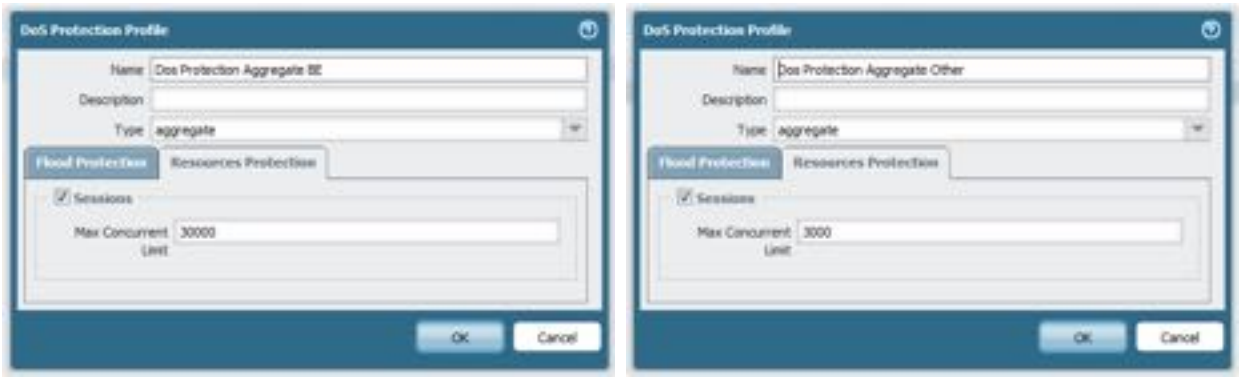
To protect against SYN Floods, it is advised to use the SYN-cookie mechanism over Random Early Drop as it is superior in preventing floods to reach the target. Adjust the rate values based on actual session data from the protected environment.

To protect against session exhaustion, you should balance the advantages and disadvantages of the following techniques and choose the most appropriate one or a combination for your environment:

- Use a classified profile with limited concurrent sessions per source-destination-ip. This will reduce the impact from a limited botnet attack while maintaining availability for regular clients. If you set the session

limit too low, this may affect concurrent clients accessing your web service from behind a NAT device. If you set it too high, it becomes easier for a small botnet attack to exhaust all available sessions.

- Use an aggregate profile in combination with geo-IP location data in the DoS Protection rule base. An aggregate profile in itself is not sufficient to prevent a session exhaustion attack, but combined with geo-IP location data in the DoS Protection rule base, it can reduce the impact of a global DDoS attack. This approach is acceptable if the main audience for your web service resides in a select list of countries.
- Use a classified profile with limited concurrent sessions per source-destination-ip in combination with geo-IP location data. This approach will require some time to set up and tune, but may prove very effective in the end.



DoS Protection rules

Once the necessary DoS Protection profiles have been defined, you can make them active by creating DoS Protection rules. DoS Protection rules define source and destination parameters on which the profile will be applied.

It is best practice to make the rules as specific as possible. For this particular example situation, you would define a rule for each server/service you want to protect. By doing so, the counter values defined in the profile will apply only to the server/service defined in the rule.

The screenshot below shows a DoS rule configuration using an aggregate profile in combination with geo-IP location data.



***Note**: you can use negate to assign a specific DoS Protection profile to any source IP address not from a particular country. This can be a worthwile approach in DDoS defense by limiting sessions for countries that do not typically fall within the general visitor audience for an enterprise's web site.*

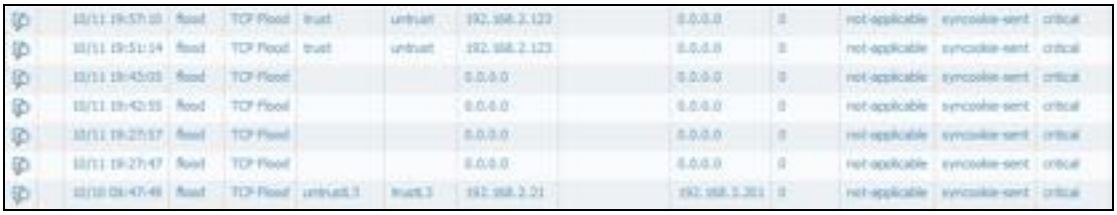
Observations

Flood Protection

When enabled and triggered, Flood Protection alerts are sent to the Threat Log. Log entries will display different values for source and destination zones and addresses depending on the type of DoS Flood Protection that was configured. Destination port will always show 0.

- When set to aggregate, then both source and destination address fields will display 0.0.0.0 and no zone information is displayed.

- When set to classified, source and destination address fields will display the actual addresses if they are part of the classification criteria.
    - E.g. source-ip will only list actual source IP addresses
    - E.g. source-dest-ip will list actual source and destination IP addresses



Resources Protection

When enabled, Resources Protection will limit the number of concurrent sessions that match the DoS Protection rule. No log entries will be created when the maximum number is reached. There is a session timeout that will keep a session in the state table for 30 seconds by default after the session has been ended with FIN/RST. The result is that the actual number of concurrent active sessions as perceived from the client side may not always match the maximum number configured in the DoS Protection profile. See Appendix D: Slow HTTP Test Output for the output of a slowhttptest attack for further clarification.

Zone Protection

Zone Protection profiles provide the possibility to configure additional protection mechanism that can be applied to specific network zones.

The following protection mechanisms are available:

- **Flood Protection–**Prevents flooding attacks in the same way as the Flood Protection feature in DoS Protection profiles.
- **Reconnaissance Protection–**Detects and blocks port scans and host sweeps.
- **Packet Based Attack Protection–**Provides protection against specific IP level attacks.

Kill-Chain

	Delivery	Exploitation	Installation	Command and Control	Actions on Objectives
DoS/Zone Protection		Prevent L3/L4 evasions			Prevent DoS attacks and lateral movement

Zone Protection can prevent L3/L4 evasions through its Packet Based Attack Protection functionality and can as such play an important role in disrupting the kill-chain at the Exploitation stage. It can also prevent network/host scans which are typically used during the Reconnaissance stage.

Recommendations

Zone protection profiles can only be applied to entire zones. Therefore it is important to investigate any possible issues that may arise when applying zone protection profiles.

Flood Protection

Configure Flood Protection settings based on the number of packets you want to allow to each service behind the firewall. Settings apply to all traffic that enters the network through any interface in the zone on which the Zone Protection Profile is active, but a separate counter is maintained for each source IP/destination IP/destination port tuple.

Reconnaissance Protection

In general it should be safe to enable this feature on external zones. For internal zones however, you should make sure settings will not negatively affect any monitoring tools that often use the same scanning techniques to determine if servers and services are up and running.

Packet Based Attack Protection

In general it should be safe to enable this feature on external zones. For internal zones however, you should make sure settings will not negatively affect network communications from other devices that may rely on some of these techniques for proper operation.

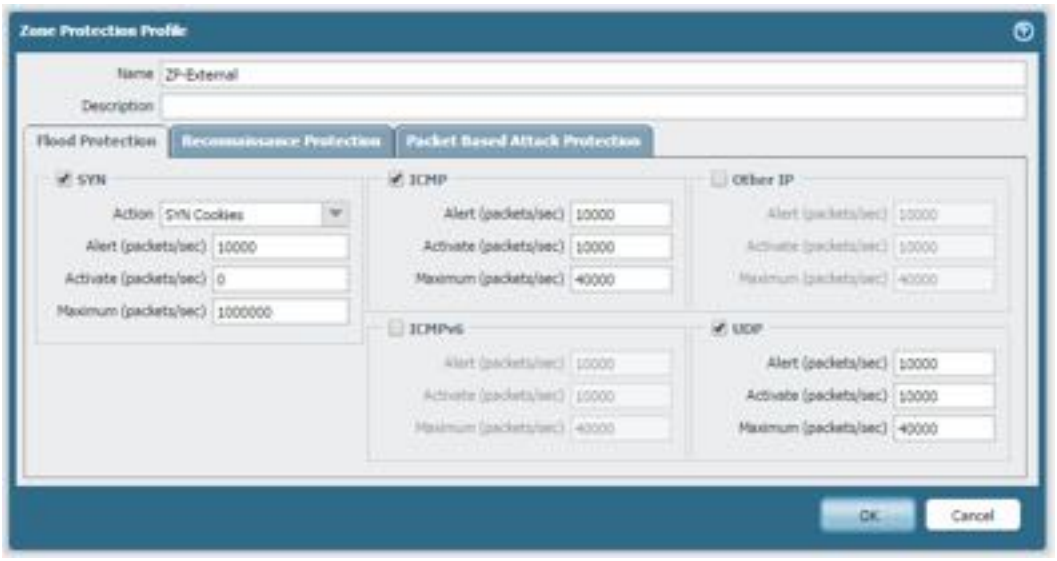
One specific example is the use of ICMP Ping ID 0, where other devices may use this type of packet to check availability of hosts they need to communicate with. Blue Coat proxy devices have been known to do this.

L4 evasions

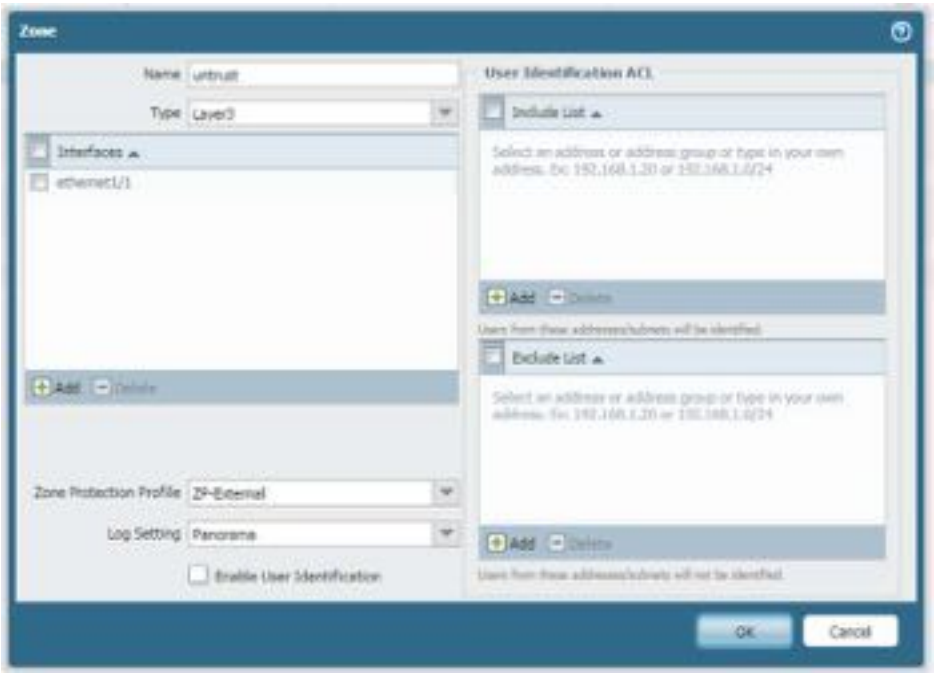
Packet Based Attack Protection can play an important role in securing your network from L4 evasions. For more information on how to correctly configure a firewall to protect against evasion techniques, please refer to Appendix E: Evasions.

Use Case

In our use case, we will enable a Zone Protection profile for the 'untrust' zone, i.e. the zone where public traffic is reaching the firewall.



Zone Protection profile



Zone settings

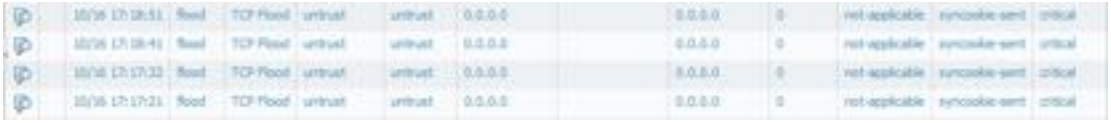
Observations

Flood Protection

Flood Protection is identical to the same feature available in DoS Protection profiles. However, since it can be applied to a zone only, there are less fine-tuning options available to match specific traffic flows based on service or IP address as you could do with DoS Protection rules.

Flood Protection enabled under Zone Protection is applied to the aggregate traffic seen on a specific zone. It will maintain a single counter for all traffic, regardless of source IP/destination IP/destination port. This is similar to Flood Protection under DoS Protection, when an aggregate profile type is selected.

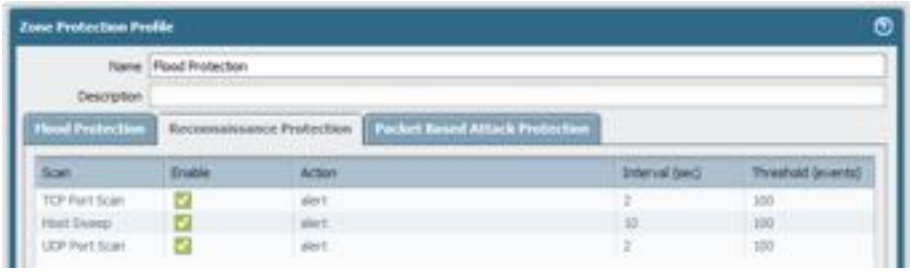
When enabled and triggered, Flood Protection alerts are sent to the Threat Log. Log entries will show the zone name for which the profile was triggered in both source and destination zone fields, while source and destination addresses will show 0.0.0.0. Destination port will also show 0, even if the flood attack was against a specific port.



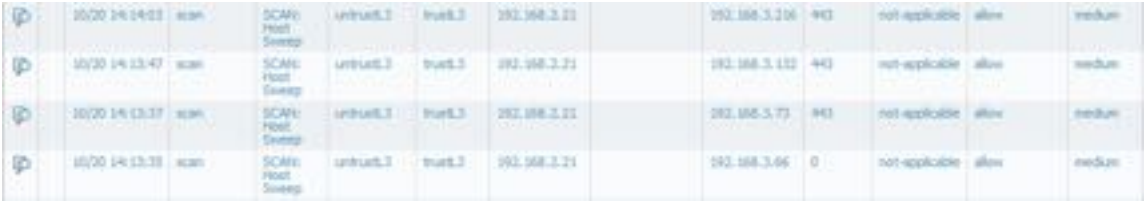
Reconnaissance Protection

Reconnaissance protection can detect and block host sweep and TCP & UDP port scans. It will trigger when the amount of scan packets exceeds the thresholds within the intervals specified.

- TCP and UDP Port scan will trigger when a TCP or UDP port scan is executed against a single host and the scan packet rate exceeds the configured threshold.
- Host Sweep will trigger when a range of hosts is scanned on specific ports either through TCP, UDP or ICMP.



When enabled and triggered, Reconnaissance Protection alerts are sent to the Threat Log.



Packet Based Attack Protection

Packet Based Attack Protection can detect and prevent attacks that try to evade firewall inspection. When enabled, packets that match detection criteria will be dropped and since this type of traffic is considered noise, no log entry will be written to the Threat Log.

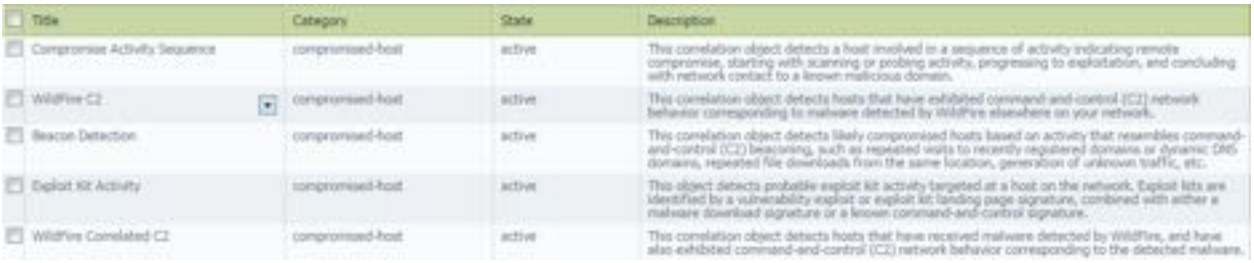
Automated Correlation Engine

The Automated Correlation Engine is an exciting new addition to the alerting capabilities of Palo Alto Networks firewalls and Panorama. It 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 attack on your network. 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.

More information on the Automated Correlation Engine can be found here:
https://www.paloaltonetworks.com/documentation/70/pan-os/newfeaturesguide/management-features/automated-correlation-engine.html

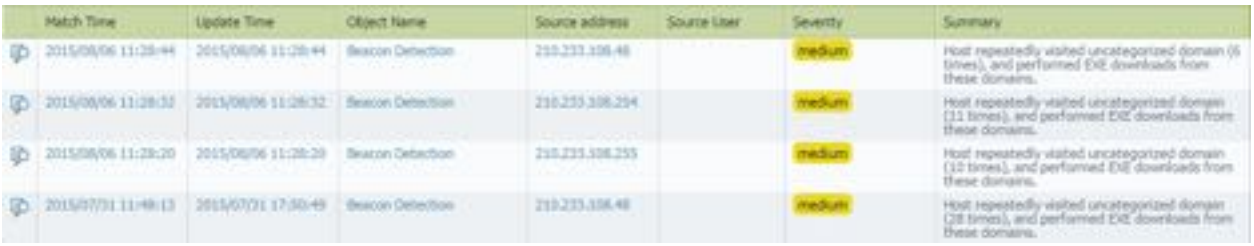
Recommendations

The Automated Correlation Engine uses Correlation Objects which are by default all enabled. Correlation Objects can be viewed and enabled/disabled via Monitor > Automated Correlation Engine > Correlation Objects.



Observations

Correlated events can be viewed under Monitor > Automated Correlation Engine > Correlated Events.



Behavioral Botnet Report

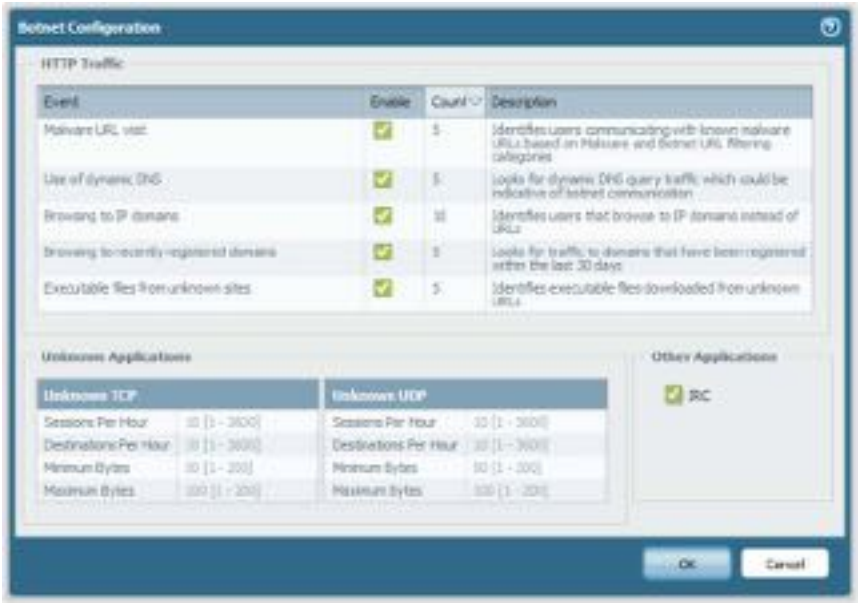
The Palo Alto Networks Behavioral Botnet Report automates the process of tracking and correlating the behaviors that indicate the presence of a bot. This intelligence looks for a variety of characteristics, which are briefly summarized below:

- **Unknown TCP/UDP** – Botnet traffic is regularly encrypted and unknown. Since Palo Alto Networks identifies all traffic tracking unknown TCP and UDP traffic can be a perfect starting point for finding bot-infected machines. The report allows staff to track unknown traffic by sessions, destinations and bytes.
- **Presence of Dynamic DNS** – Malware will often use dynamic DNS in order to make botnet communications more difficult to track. By bouncing traffic between multiple infected hosts with an ever-changing list of IP addresses, it can become very difficult to track the path of the bot and its true source and destination.
- **Activity on Known Malware Sites** – As part of the URL filtering solution, Palo Alto Networks constantly tracks sites that have hosted malware whether intentionally or unintentionally. Palo Alto Networks can track if a user is repeatedly visiting one of these sites and attempting to download files.
- **Visiting Recently Registered Domains** – Botnets are constantly moving around in order to avoid detection and to recover as servers are discovered or disabled. As a result, botnets will often have to use new domains to support the command and control infrastructure. A user repeatedly visiting a newly registered domain will certainly not be conclusive, but may help to provide corroborating evidence of an infection.
- **Browsing to IP domains Instead of URL** – In a similar vein, bots will often use hard-coded IP addresses or known IP ranges in order to communicate as opposed to users which typically prefer to use URLs. As with tracking newly registered domains, tracking connections using IP domains can sometimes indicate the presence of a bot at work as opposed to a human.
- **Executable files from unknown sites** – Identifies executable files downloaded from unknown URLs.
- **IRC traffic** – IRC traffic is one of the most well-known communication methods for botnets, and provides an additional strong piece of correlating data for finding a bot.

The Behavioral Botnet Report takes all of the factors above and automatically correlates them to find hosts that are likely infected with a bot. When run, the report provides specific directory user names of the users or machines that are likely infected along with what behaviors contributed to the analysis. Each user is also provided a score based on how many of the factors listed above were correlated, allowing staff to focus on the devices that are the most likely to be infected.

Recommendations

The behavioral botnet report is enabled by default. It will run on a daily basis and generate a report on the detected behaviors for the previous day. It may be required to further tune the default detection threshold values based on the initial report findings.



Observations

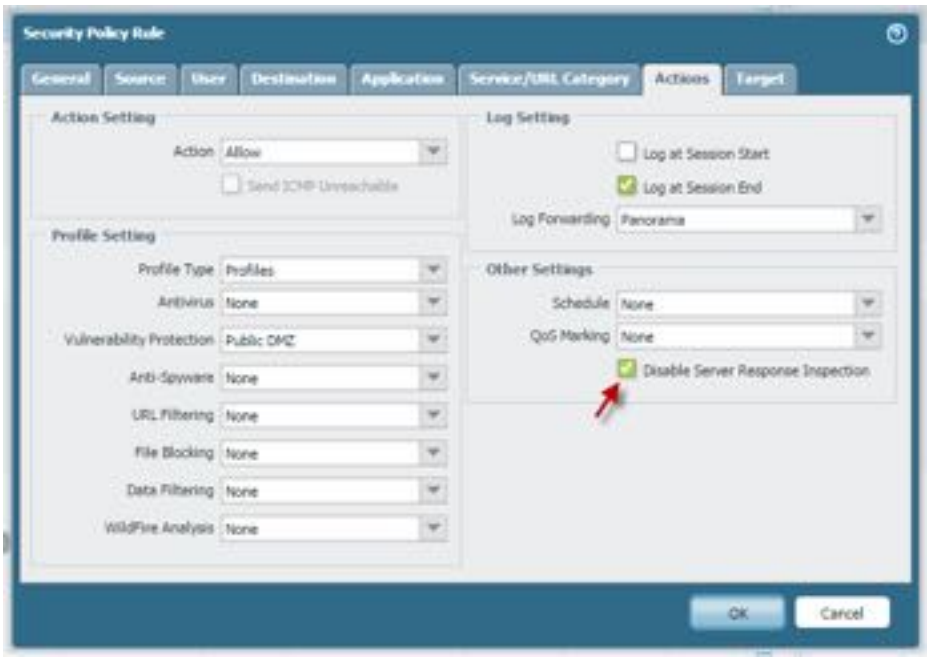
When enabled, the botnet report is run every night on the log data that was collected for the previous day. The report will display hosts together with a confidence score and suspicious behavior details.

	Confidence	Source Address	Source User	Description
1	4	210.233.108.254		Repeatedly (10) downloads executables from the same unknown URL 63.251.135.17/cbad-554.0000 . Repeatedly (30) downloads executables from the same unknown URL vfvyqhbqcg.hk/progs/vdgcmwt/mjailti.php . Repeatedly visited (16) the same malicious URL fp.outerinfo.com/dispatcher.php . Repeatedly visited (22) the same URL 81.95.150.42/pr/gt.php
2	4	210.233.108.255		Downloads executable(s) from unknown URL 81.95.146.206/wincg.exe . Repeatedly (40) downloads executables from the same unknown URL vfvyqhbqcg.hk/progs/kvolhr/tmf/pmm.php . Repeatedly visited (12) the same malicious URL fp.outerinfo.com/dispatcher.php . Repeatedly visited (64) the same URL 66.246.246.236/sync0/upd/u.txt
3	3	210.233.108.48		Repeatedly (21) downloads executables from the same unknown URL wfhcckyedu.hk/progs/fxbkuqw/ddmmhpkf.php . Visited malware URL fp.outerinfo.com/dispatcher.php . Repeatedly visited (24) the same URL 75.126.64.8/me/cfg.bin
4	2	10.0.101.2	labs\bob	Repeatedly visited (11) the same URL 212.113.82.167/dcs5wetukbedgcm6fgwnoan0_8f9w/dcs.gif?&dcsdat=1174051402119&dcssip=www.standaard.be&dcsuri=/&dcsfil=en&js=standaard&met=&btnG=Google+Search&WT.tz=1&WT.bh=GB&WT.cd=16&WT.sr=1400x1050&WT.jo=Yes&WT.ti=De%20Stand...
5	2	10.0.101.1	labs\alice	Repeatedly visited (39) the same URL 208.85.146.37/g/22325/?penns

Appendix A: Performance

Threat Prevention performance will not be impacted by the number of Threat Prevention profiles or the number of signatures that are enabled thanks to the single pass parallel processing architecture (SP3) unique to the Palo Alto Networks firewall. Any session content is inspected in parallel by the Antivirus, Anti-Spyware and Vulnerability Protection features.

One notable configuration setting that will have a positive impact on performance is DSRI or Disable Server Response Inspection. With DSRI turned on, server response traffic is not inspected, which will increase the throughput capacity. Obviously, enabling this feature is only recommended for trusted servers.



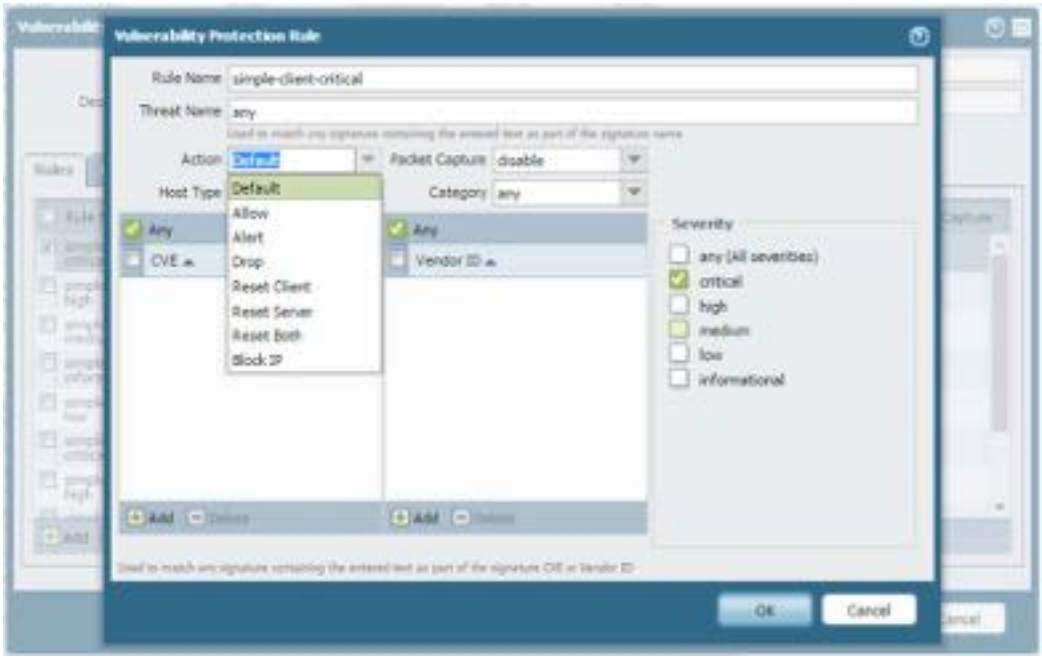
Appendix B: Default Action for Threat Prevention Signatures

Every threat signature present in the Anti-Spyware and Vulnerability Protection profiles has a default action assigned to it. This default action is the action that will be performed when 'default' is selected as the action in threat prevention rules that define the security profile.

Upon the initial installation of a Threat Prevention update, the default action for each signature is set by the Palo Alto Networks Threat Prevention research team to the most appropriate response at any moment based on a wider set of criteria. The default action can change over time when new threat prevention updates are released. The default action can also be modified on a per-signature basis when creating a custom profile.

Below are the available actions that can be applied to individual Anti-Spyware and Vulnerability Protection events.

- Default
- Allow
- Alert
- Drop
- Reset-both
- Reset-client
- Reset-server
- Block-IP



For the Antivirus feature, the default action is not set on a per-signature basis, but configured and applied per protocol.

Below are the available actions that can be applied to Antivirus events.

- Allow
- Alert
- Drop
- Reset-client
- Reset-server
- Reset-both

Appendix C: Guidelines for the Vulnerability Protection Analysis Phase

The goal of this phase will be to determine the correct vulnerability protection profile settings for each of the protected segments and hosts. Multiple profile configurations may be needed for different segments and hosts in your network.

The following steps will help you determine the correct profile settings for a given location or host.

1. Configure an 'alert-only' protection profile.
2. Configure the necessary firewall rules for hosts and segments you want to protect.
3. Apply the 'alert-only' protection profile to each rule.
4. Monitor the threat logs for a representative period of time (e.g. 1 week, 1 month).
5. Investigate any potential false positives.
6. Use the gathered analysis information to build and fine-tune a block-enabled protection profile.

Configure an 'alert-only' protection profile

The goal here is to create a 'alert-only' protection profile that has 'alert' defined as the action for each signature. This will prevent the accidental blocking of legitimate traffic during the analysis phase. You can accomplish this with one profile rule that includes all signatures. The disadvantage of this approach is that the threat log will not display the actual action applied when in blocking mode using the default action for each signature.

Alternatively you can monitor a segment or host via a port of the firewall in TAP-mode and use a protection profile that applies the 'default' action for each signature. This will also prevent the accidental blocking of legitimate traffic, but the advantage here is that the threat log displays the actual action that would be taken for each signature when in blocking mode.

Configure the necessary firewall rules for hosts and segments you want to protect.

During the analysis phase, it is important to define and match the firewall security policy rules as close as possible to what the final security policy will look like. This will help you map vulnerability protection events to specific segments, hosts and traffic flows in your network and facilitate the analysis phase.

Apply the 'alert-only' protection profile to each rule.

Initially, you will have the 'alert-only' profile applied to each rule. Over time, it may be required to create additional profiles that are fine-tuned for specific network segments or hosts.

Monitor the threat logs for a representative period of time.

The goal here is to acquire a representative baseline set of threat log events that can be used as the starting point for further fine-tuning of the vulnerability protection profile(s). It is important to define the monitor period in line with regular business and infrastructure management processes that may only occur at specific times, e.g. once a month.

Investigate any potential false positives.

This step is crucial in defining the correct vulnerability protection profile to be deployed in blocking mode. Special attention should be given to those events that are more generic in nature and have a blocking action as their default action.

Depending on the severity and type of threat detected, the source and destination as well as the recurrence of an event, this may or may not be an easy task at hand. Especially for those events that are generic in nature and

where custom applications are involved, this may require some detailed analysis. The possibility to take a packet capture of these events can be of great help here.

A classic example is a custom web application that relies on a back-end database system. If the web application is not well coded, certain injection attack signatures may in fact trigger on what is typical traffic for the application. The same applies to generic cross-site scripting signatures.

Another example is the 'probing' that some network infrastructure monitoring tools perform against components in the network. While these can also trigger certain signatures, the severity for this type of event is typically low or informational. Nevertheless, it is important to verify that there is an exception defined in case a block action should be required.

Use the gathered analysis information to build and fine-tune a block-enabled protection profile.

Once the analysis results are consistent, they can be used to create vulnerability protection profiles that provide adequate protection tailored to the requirements for the different segments or hosts in the network. In most cases, fine-tuning will consist of adding exceptions to the general vulnerability protection rules, either in order to avoid a false positive or to modify the default action for specific signatures.

Appendix D: Slow HTTP Test Output

A DoS Protection profile enabled with Resources Protection Session Limit set to 300 concurrent sessions.

Slowhttptest tool used to open 1000 connections at a rate of 200 new connections per second.

```
root@attacker:~/slowhttptest/bin# ./slowhttptest -c 1000 -B -g -o my_server_stats -i 110 -r 200 -
s 8192 -t FAKEVERB -u http://192.168.2.20 -x 10 -p 3
Thu Oct 20 17:50:15 2011:
Using:
slow section:                      BODY
number of connections:             1000
URL:                               http://192.168.2.20/
verb:                              FAKEVERB
Content-Length header value:       8192
follow up data max size:           22
interval between follow up data:   110 seconds
connections per seconds:           200
probe connection timeout:          3 seconds
test duration:                     240 seconds
Thu Oct 20 17:50:15 2011:slow HTTP test status on 0th second:
inititalizing:       0
pending:             1
connected:           0
error:               0
closed:              0
service available:   YES
Thu Oct 20 17:50:20 2011:slow HTTP test status on 5th second:
inititalizing:       0
pending:             677
connected:           226
error:               0
closed:              0
service available:   NO
Thu Oct 20 17:50:25 2011:slow HTTP test status on 10th second:
inititalizing:       0
pending:             662
connected:           338
error:               0
closed:              0
service available:   NO
Thu Oct 20 17:50:31 2011:slow HTTP test status on 15th second:
inititalizing:       0
pending:             629
```

```
connected:           334
error:               0
closed:              37
service available:   NO
Thu Oct 20 17:50:35 2011:slow HTTP test status on 20th second:
inititalizing:       0
pending:             629
connected:           221
error:               0
closed:              150
service available:   NO
Thu Oct 20 17:50:41 2011:slow HTTP test status on 25th second:
inititalizing:       0
pending:             519
connected:           294
error:               0
closed:              187
service available:   NO
Thu Oct 20 17:50:45 2011:slow HTTP test status on 30th second:
inititalizing:       0
pending:             473
connected:           228
error:               0
closed:              299
service available:   NO
Thu Oct 20 17:50:51 2011:slow HTTP test status on 35th second:
inititalizing:       0
pending:             473
connected:           190
error:               0
closed:              337
service available:   NO
Thu Oct 20 17:50:55 2011:slow HTTP test status on 40th second:
inititalizing:       0
pending:             473
connected:           78
error:               0
closed:              449
service available:   YES
Thu Oct 20 17:51:01 2011:slow HTTP test status on 45th second:
inititalizing:       0
pending:             473
```

```
connected:           40
error:               0
closed:              487
service available:   YES
Thu Oct 20 17:51:05 2011:slow HTTP test status on 50th second:
inititalizing:       0
pending:             193
connected:           284
error:               0
closed:              523
service available:   YES
Thu Oct 20 17:51:10 2011:slow HTTP test status on 55th second:
inititalizing:       0
pending:             183
connected:           292
error:               0
closed:              525
service available:   NO
Thu Oct 20 17:51:15 2011:slow HTTP test status on 60th second:
inititalizing:       0
pending:             183
connected:           147
error:               0
closed:              670
service available:   YES
Thu Oct 20 17:51:21 2011:slow HTTP test status on 65th second:
inititalizing:       0
pending:             183
connected:           142
error:               0
closed:              675
service available:   YES
Thu Oct 20 17:51:25 2011:slow HTTP test status on 70th second:
inititalizing:       0
pending:             183
connected:           8
error:               0
closed:              809
service available:   YES
Thu Oct 20 17:51:31 2011:slow HTTP test status on 75th second:
inititalizing:       0
pending:             183
```

```
connected:           0
error:               0
closed:              817
service available:   YES
Thu Oct 20 17:51:36 2011:slow HTTP test status on 80th second:
inititalizing:       0
pending:             183
connected:           0
error:               0
closed:              817
service available:   YES
Thu Oct 20 17:51:41 2011:slow HTTP test status on 85th second:
inititalizing:       0
pending:             183
connected:           0
error:               0
closed:              817
service available:   YES
Thu Oct 20 17:51:46 2011:slow HTTP test status on 90th second:
inititalizing:       0
pending:             183
connected:           0
error:               0
closed:              817
service available:   YES
Thu Oct 20 17:51:51 2011:slow HTTP test status on 95th second:
inititalizing:       0
pending:             183
connected:           0
error:               0
closed:              817
service available:   YES
Thu Oct 20 17:51:56 2011:slow HTTP test status on 100th second:
inititalizing:       0
pending:             0
connected:           183
error:               0
closed:              817
service available:   YES
Thu Oct 20 17:52:00 2011:slow HTTP test status on 105th second:
inititalizing:       0
pending:             0
```

```
connected:           183
error:               0
closed:              817
service available:   NO
Thu Oct 20 17:52:05 2011:slow HTTP test status on 110th second:
inititalizing:       0
pending:             0
connected:           61
error:               0
closed:              939
service available:   YES
Thu Oct 20 17:52:11 2011:slow HTTP test status on 115th second:
inititalizing:       0
pending:             0
connected:           33
error:               0
closed:              967
service available:   YES
Thu Oct 20 17:52:12 2011:
Using:
slow section:                      BODY
number of connections:             1000
URL:                               http://192.168.2.20/
verb:                              FAKEVERB
Content-Length header value:       8192
follow up data max size:           22
interval between follow up data:   110 seconds
connections per seconds:           200
probe connection timeout:          3 seconds
test duration:                     240 seconds
Thu Oct 20 17:52:12 2011:Test ended on 116th second
status:                            No open connections left
root@attacker:~/slowhttptest/bin#
```

Appendix E: Evasions

Evasions are a way to bypass a firewall's inspection functionalities by applying techniques such as fragmentation or overlapping TCP segments. In addition to L4 evasion techniques, L7 evasions also exist and are even more popular due to their simplicity. SSL encryption for example is a very simple yet effective technique to bypass inspection mechanisms where no SSL decryption is performed.

The following topic from the product documentation provides best practices for securing your network from layer-4 and layer-7 evasions.

https://www.paloaltonetworks.com/documentation/70/pan-os/pan-os/threat-prevention/best-practices-for-securing-your-network-from-layer-4-and-layer-7-evasions.html#16594

Revision History

Date	Version	Revision	Comment
August 11, 2015	3.0	A	Adapted content to firmware 7.0 Added attack kill-chain and zero-trust principles and additional information for each threat prevention functionality.
March 4, 2014	2.0	A	Adapted content to firmware 6.0
July 8, 2013	1.2	A	Adapted content to firmware 5.0
July 24, 2012	1.1	B	URL Filtering use case p14: inserted a note that explains and provides alternative solution for the warning messages during commit on the security policy rule for the first requirement of this use case. Modified the rule so destination zone is now untrustL3 instead of any.
June 22th, 2012	1.1	A	First release of this document.

What's New in PAN-OS 6.0

With the release of PAN-OS 6.0, we continue to strengthen our leadership through innovation with more than 60 new features to prevent advanced threats, secure virtualized environments, and protect mobile devices. Here is a summary of the new features in PAN-OS and Panorama 6.0.

WildFire Features

- **Additional File Type Support**—As part of the WildFire subscription, the following advanced file types are now supported: Microsoft Office .doc,.xls, and .ppt; Portable Document Format (PDF); Java Applet (jar and class); and Android Application Package (APK). **NOTE:** The WF-500 does not support APK file analysis.

- **Expanded Sandbox Operating Systems—**Microsoft Windows 7 32/bit has been added to the WildFire environment. When a file is analyzed by WlidFire, it will be run in both 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.

- **WildFire Report Incorrect Verdict 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 Analysis Report**—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:
    - Export the full report to a PDF.
    - Download the file sample that was analyzed.
    - View all processes or filter by an individual process.
    - View the analysis results for each virtual machine environment in which the file was analyzed.
    - Re-submit the file sample to Palo Alto Networks for reevaluation if you think the file verdict (benign/malware) is incorrect.

- **WildFire Logs on the Firewall**—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 WildFire logs on the firewall.

- **WildFire Reporting**—The WildFire detailed report is now integrated into the firewall showing session details and the WildFire detailed report, which was previously hosted on the WildFire cloud or WildFire appliance. In addition, Panorama no longer requires that all managed firewalls forward files to the same WildFire system as long as Panorama and the managed firewalls are running 6.0 or later.

- **WildFire Submissions Log Forwarding**—You can now configure the firewall to automatically forward WildFire Submissions logs independently of the threat log forwarding configuration.

Content Inspection Features

- **DNS Sinkholing**—DNS 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. 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 (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.

- **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 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. 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 as well as to improve accuracy and malware detection abilities within PAN-DB URL filtering, DNS-based command-and-control signatures, and WildFire. Passive DNS monitoring is a configurable feature that is disabled by default.

© 2014 Palo Alto Networks (PAN-OS 6.0)

- **URL Filtering Search Engine Cached Site Support**—URL filtering policies will be applied when end-users attempt to view the Google and Internet Archive cached copies of websites.

- **URL Filtering Translation Site Filtering Support**—URL filtering policies will 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 who are searching the Internet using one of the top three search providers—Google, Bing, or Yahoo—from viewing the search results unless the strictest safe search option is set in their browsers for these search engines. If the strictest safe search option is not set in the browser, users will see a block page instructing them on how to set the option for the given search provider.

- **Extended Packet Capture**—Enables the firewall to capture from 1 to 50 packets for anti- spyware and vulnerability protection profiles for rules and exceptions defined in the profile. The extended capture provides more context when analyzing threat logs.

Application Identification Features

- **Support for Hardware Security Modules**—Certificate signing functions for SSL forward proxy, SSL inbound inspection, and the master key storage functions can now be offloaded to a dedicated Hardware Security Module (HSM) for enhanced key management security. HSM support is generally required when FIPS 140-2 Level 3 protection for CA keys is required.
    - HSMs supported: SafeNet Luna SA and Thales Nshield Connect.
    - Platforms supported: PA-3000 Series, PA-4000 Series, PA-5000 Series firewalls, VM-Series, and on the M-100 appliance.

- **TLS 1.2 Decryption**—The ability to decrypt inbound sessions and forward proxy sessions that negotiate with TLS 1.2 is now supported.

- **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.

User Identification Features

- **User-ID Integration with Syslog**—The Syslog Listener will listen for syslog messages from non-standard user authentication services (Proxies, NAC, Wireless Controllers) so that the User-ID 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.
    - Syslog Listener natively supports BlueCoat Proxy, Citrix Access Gateway, Aerohive AP, Cisco ASA, Juniper SA Net Connect, and the Juniper Infranet Controller.

- **Terminal Services API**—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 logon and logoff by including port mappings in addition to IP addresses.

- **Increased User-ID Active Users Limit**—The User-ID active user limit has been increased on the high-end firewall platforms, based on the memory capacity of the individual platforms. The following table summarizes the User-ID active limits on all Palo Alto Networks next-generation firewall platforms:

Firewall Platform	User-ID Active User Limit
PA-5060	256,000
PA-5050 and PA-5020	128,000
PA-4000 Series, PA-3000 Series,PA-2000 Series, PA-500, and PA-200	64,000

© 2014 Palo Alto Networks (PAN-OS 6.0)

Networking Features

- **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 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. **Note:** Decryption port mirroring is available on the PA-5000 Series and PA-3000 Series platforms only.

- **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.

- **OSPF Graceful Restart**—Reduces the Active/Passive HA failover times by allowing 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 network outage that is experienced when OSPF adjacencies are reestablished.

- **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.

- **Increase Jumbo Frame Size**—The maximum transmission unit (MTU) size has been increased to provide compatibility with equipment from other vendors. 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.

- **IPv6 Neighbor Discovery Table Capacity Increase**—In previous versions of PAN-OS, 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.

- **Consolidation of Timers Used in a High Availability (HA) Setup**—To reduce the complexity in configuring HA timers used to detect a firewall failure and trigger a failover, three profiles have been added:
  - Recommended profile is for typical failover timer settings
  - Aggressive profile is for faster failover timer settings
  - Advanced profile allows you to customize the timer values to suit your network requirements.

  The profiles auto-populate the optimum HA timer values for the specific firewall platform to enable a more rapid HA deployment

- **Enhanced Use for Address Objects**—A Layer 3 interface on the firewall can now use an address object in lieu of an IP address which allows you to separate the object from its value/IP address. For Panorama this new feature improves scalability by allowing you to create a template that references 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. Panorama templates support address objects in the following locations:
  - IP addresses for Layer 3 interfaces
  - Virtual address of the Virtual IP address in an active/active HA set up
  - Service routes
  - NAT policy for source translation
  - GlobalProtect portal and gateway IP addresses
  - GlobalProtect satellite and site-to-site IP addresses
  - IKE Gateway local address
  - Multicast static and candidate RP addresses
  - BGP peer group local addresses
  - Hardware Security Module client IP addresses

Virtualization Features

- **Virtual Machine (VM) Monitoring Agent**—This feature eliminates manual scripting from the previous release and allows you to dynamically retrieve virtual machine IP address changes in your virtualized environment without making configuration changes on the firewall.

  The User-ID Windows agent and the firewall (agentless User-ID) can now be configured to proactively monitor the following VM Sources: VMware ESX(i) server (4.1 and 5.0) and VMware VCenter. The Windows User-ID agent supports up to 100 sources and the firewall supports up to 10 sources.

- **Dynamic Address Groups**—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. 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.

- **VM-Series on Citrix SDX**—The VM-Series firewall is now supported on the Citrix SDX hardware platform running Citrix XenServer. Deploying the VM- Series firewall (one or more instances) on the SDX server provides the ability to consolidate the NetScaler VPX and the VM-Series firewall on the same physical platform. This addresses consolidated application delivery controller and security needs for multi-tenant cloud deployments (business units, application owners, service providers) or Citrix XenApp XenDesktop deployments. Supported Citrix platforms are the 11500 or 17500 Series running Citrix XenServer version 6.0.2 or later.

- **VM-Series for VMware NSX—NOTE: The VM-Series for VMware NSX will be available in late Q1 CY2014.** The Palo Alto Networks and VMware joint solution addresses challenges associated with applying network security to software defined networks. With this new offering, customers will be able to safely enable intra-server virtual machine communications. NSX, VMware's Networking and Security platform, automates the process of deploying and provisioning the VM-Series firewall as a service (also called Security Virtual Machine) on ESXi servers. VM to VM traffic is automatically steered to the VM-Series without requiring any manual virtual network configuration. VM context is also shared between NSX and Panorama, to keep track of virtual machine provisioning and changes.

VMware/VM-Series Integration Components and Requirements	
VMware NSX	VMware vCenter 5.5.0 VMware NSX Manager 6.0 VMware ESXI vSphere 5.5
VM-Series	VM-Series for VMware NSX (model available in late Q1 CY2014) in PAN-OS/Panorama 6.0
Minimum system requirements	• Two vCPUs. One will be used 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. • 5GB of memory. Any additional memory will be used by the management plane only. • 40GB of virtual disk space

GlobalProtect Features

- **Agent Deployment Customization**—All GlobalProtect agent customization settings can now be set in the Windows Registry allowing you to deploy 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. Additional configuration options include:
  - **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.
  - **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.
  - **Allow user to continue if portal server certificate is invalid**—Disabling this option prevents users from continuing if there is a warning screen that might indicate a man-in-the-middle (MITM) attack.

- **Agent Update Control**—The client configurations delivered by GlobalProtect now have two additional options (disable and manual) for controlling GlobalProtect agent upgrades:
    - By default, the portal will prompt the agent to upgrade whenever a new version is available.
    - Transparent: enables automatic upgrades that do not require user interaction.
    - Disable: (new) prevents agent upgrades.
    - Manual: (new) allows end users to initiate agent upgrades.

- **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 HIP profiles you create, thereby enabling you to use DLP compliance as criteria for your security policies.

- **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
    - **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.
    - **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 accelerates the authentication process when the portal and the gateway require different credentials (either differentOTPs 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.

- **Certificate Authentication Enforcement**—Enhancements have been made to how client certificate authentication is enforced in various scenarios as follows:
    - 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 pre- deployed to the end clients before their initial portal connection.
    - 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.
    - 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.

- **GlobalProtect Mobile Security Manager and the GP-100 Appliance**—Provides, management, visibility, and automated configuration deployment of iOS and Android-based mobile devices—either company provisioned or employee owned. 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. Additional device management action you can execute include: locking the device, sounding an alarm to help locate the device, or even wiping a device that has been compromised. GlobalProtect gateways can retrieve 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. For more information, please review the GlobalProtect datasheet.

© 2014 Palo Alto Networks (PAN-OS 6.0)

Management Features

- **Commit Improvement**—The commit operation in PAN-OS and Panorama 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 the use of a CDN for delivery of software and content updates to devices running PAN-OS/Panorama 6.0 or later. 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.

  In addition, a new option has been added to strengthen all secure communication between firewalls/Panorama and the Palo Alto Networks update servers. When this feature is 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.

- **Rule Numbering**—Each rule in the Policies tab 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.

  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.

- **Reporting Enhancements**—Reporting enhancements in this release include support for creating group activity reports, complementing user activity reports and the ability to disable predefined reports. Because Panorama does not have the user to group mapping information, you cannot generate group activity reports on Panorama.

- **CLI Find Command**—The new CLI find command helps you find a command when you don't know where to start looking in the hierarchy. The command—which is available in all CLI modes—has two forms. You can either use find command alone to display the entire command hierarchy in the current command mode. Or, you can use find command with the keyword argument to locate all commands that have the specified keyword.

- **Support for Syslog over TCP and SSL**—PAN-OS 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).

- **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.
  The HOST-RESOURCES-MIB includes objects for monitoring the following:
  - 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.
  - Packet buffers on the data plane.

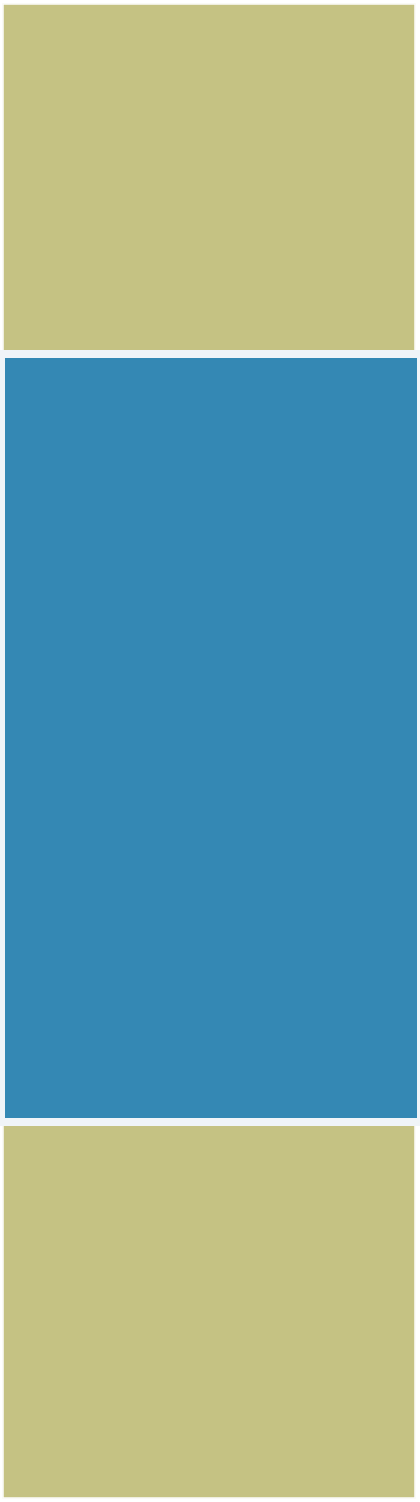
  The Enterprise MIBs include objects for monitoring the following:
  - GlobalProtect gateway utilization in the PAN- Common MIB- to monitor the total number of current active tunnels on the firewall device and the percentage on current tunnel utilization that is based on the number of active tunnels and the maximum number of tunnels allowed.
  - 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.
- **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.

- **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.

Panorama Features

- **Scheduling Dynamic Updates from Panorama**— Dynamic updates for Applications and Threats, WildFire, Antivirus, and URL Database can now be scheduled. 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.
- **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.
- **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.

© 2014 Palo Alto Networks (PAN-OS 6.0)





Palo Alto Networks

WildFire™ 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 describes the administrative tasks required to use and maintain the Palo Alto Networks WildFire feature. Topics covered include licensing information, configuring firewalls to forward files for inspection, viewing reports, and how to configure and manage the WF-500 appliance.

- 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 the support programs, or to manage your account or devices, refer to https://support.paloaltonetworks.com.
- For the latest release notes, go to the software downloads page at https://support.paloaltonetworks.com/Updates/SoftwareUpdates.

To provide feedback on the documentation, please write to us at: documentation@paloaltonetworks.com.

Palo Alto Networks, Inc.
www.paloaltonetworks.com
© 2014–2015 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: November 18, 2015

Table of Contents

WildFire Overview 7

About WildFire 8

WildFire Deployments 11
- WildFire Public Cloud 11
- WildFire Private Cloud 11
- WildFire Hybrid Cloud 12

WildFire Concepts 13
- Virtual Sandbox 13
- Verdicts 13
- File Type Analysis 14
- Email Link Analysis 14
- Signatures 15
- Alerts 15
- Logs and Reporting 16

WildFire Subscription Requirements 19

Best Practices for Keeping Signatures up to Date 20

Configure the WF-500 Appliance 21

About the WF-500 Appliance 22

Configure the WF-500 Appliance 23

Set Up the VM Interface 29
- Virtual Machine Interface Overview 29
- Configure the VM Interface on the WF-500 Appliance 30
- Connect the Firewall to the WF-500 Appliance VM Interface 32

Set Up WF-500 Content Updates 34
- Install WF-500 Content Updates Directly from the Update Server 34
- Install WF-500 Content Updates from an SCP-Enabled Server 35

Enable Signatures and URL Generation 37

Upgrade a WF-500 Appliance 39

Submit Files for WildFire Analysis 43

Forward Files for WildFire Analysis 44

Verify WildFire Submissions 48
- Test a Sample Malware File 48
- Verify File Forwarding 48

Manually Upload Files to the WildFire Portal 53

Submit Malware or Reports from the WF-500 Appliance 54

Firewall File Forwarding Capacity by Platform 55

Monitor WildFire Activity 57

Use the Firewall to Monitor WildFire Activity 58
- Configure WildFire Submissions Log Settings 58
- Monitor WildFire Submissions and Analysis Reports 60
- Set Up Alerts for Malware Logs 62

Use the WildFire Portal to Monitor WildFire Activity 65
- Configure WildFire Portal Settings 65
- Add WildFire Portal Users 66
- View Reports on the WildFire Portal 67

WildFire Analysis Reports—Close Up 69

Use the WildFire API 73

About WildFire Subscriptions and API Keys 74

WildFire API File Submission Methods 75
- Submit a File to the WildFire Cloud Using the Submit File Method 75
- Submit a File to WildFire Using the Submit URL Method 75

Query for a WildFire PDF or XML Report 77

Use the API to Retrieve a Sample Malware Test File 78

Use the API to Retrieve a Sample or PCAP 79

Use the WildFire API on a WF-500 Appliance 81
- Generate API Keys on the WF-500 Appliance 81
- Manage API Keys on the WF-500 Appliance 82
- Use the WildFire API on a WF-500 Appliance 83

Use the WF-500 Appliance CLI 85

WF-500 Appliance Software CLI Concepts 86
- WF-500 Appliance Software CLI Structure 86
- WF-500 Appliance Software CLI Command Conventions 86
- WF-500 Appliance CLI Command Messages 87
- Command Option Symbols 87
- Privilege Levels 89

WildFire CLI Command Modes 90
- Configuration Mode 90
- Operational Mode 94

Access the CLI 95
- Establish a Direct Console Connection 95
- Establish an SSH Connection 95

Use the CLI 96
- Access Operational and Configuration Modes 96
- Display WF-500 Appliance Software CLI Command Options 96
- Restrict Command Output 97
- Set the Output Format for Configuration Commands 98

Configuration Mode Command Reference 99
- set deviceconfig setting wildfire 99
- set deviceconfig system update-schedule 101
- set deviceconfig system vm-interface 102

Operational Mode Command Reference. 104
create wildfire api-key 104
delete wildfire api-key 105
delete wildfire-metadata 106
edit wildfire api-key 107
load wildfire api-key 108
request system raid 109
request system wildfire-vm-image 110
request wf-content 111
save wildfire api-key 112
set wildfire portal-admin 113
show system raid 114
show wildfire 115
test wildfire registration 119

WildFire Example 121



WildFire Overview

WildFire™ provides detection and prevention of zero-day malware using a combination of malware sandboxing, signature-based detection, and blocking of malware. WildFire extends the capabilities of Palo Alto Networks next-generation firewalls to identify and block targeted and unknown malware.

- About WildFire
- WildFire Deployments
- WildFire Concepts
- WildFire Subscription Requirements

About WildFire

Modern malware is at the heart of many of today's most sophisticated network attacks and is increasingly customized to avoid traditional security solutions. Palo Alto Networks has developed an integrated approach that addresses the full malware life cycle, which includes preventing infections, identifying zero-day malware (undiscovered malware), or targeted malware (malware targeting a specific industry or corporation), as well as pinpointing and disrupting active infections.

The Palo Alto Networks WildFire engine exposes zero-day and targeted malware through direct observation in a virtual environment within the WildFire system. The WildFire feature also makes extensive use of the Palo Alto Networks App-ID technology by identifying file transfers within all applications, not just email attachments or browser-based file downloads.

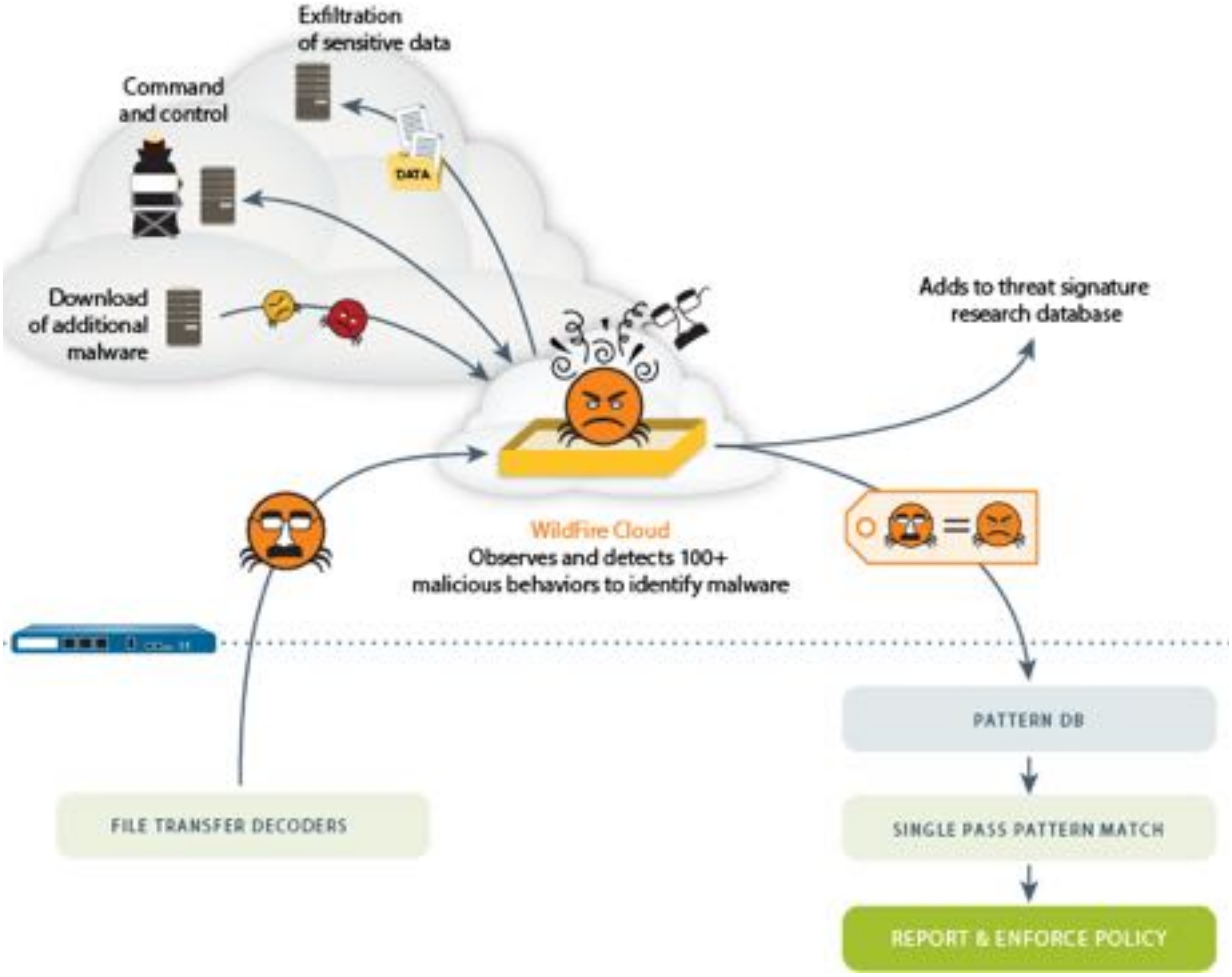
For information on Palo Alto Networks WildFire privacy policy, refer to https://live.paloaltonetworks.com/docs/DOC-2880.

High-Level WildFire Decision Workflow illustrates the basic WildFire workflow, Detailed WildFire Decision Flow describes the entire WildFire lifecycle from the time a user downloads a malicious file to the point where WildFire generates a signature to be used by Palo Alto Networks firewalls to protect against future exposure to the malware.

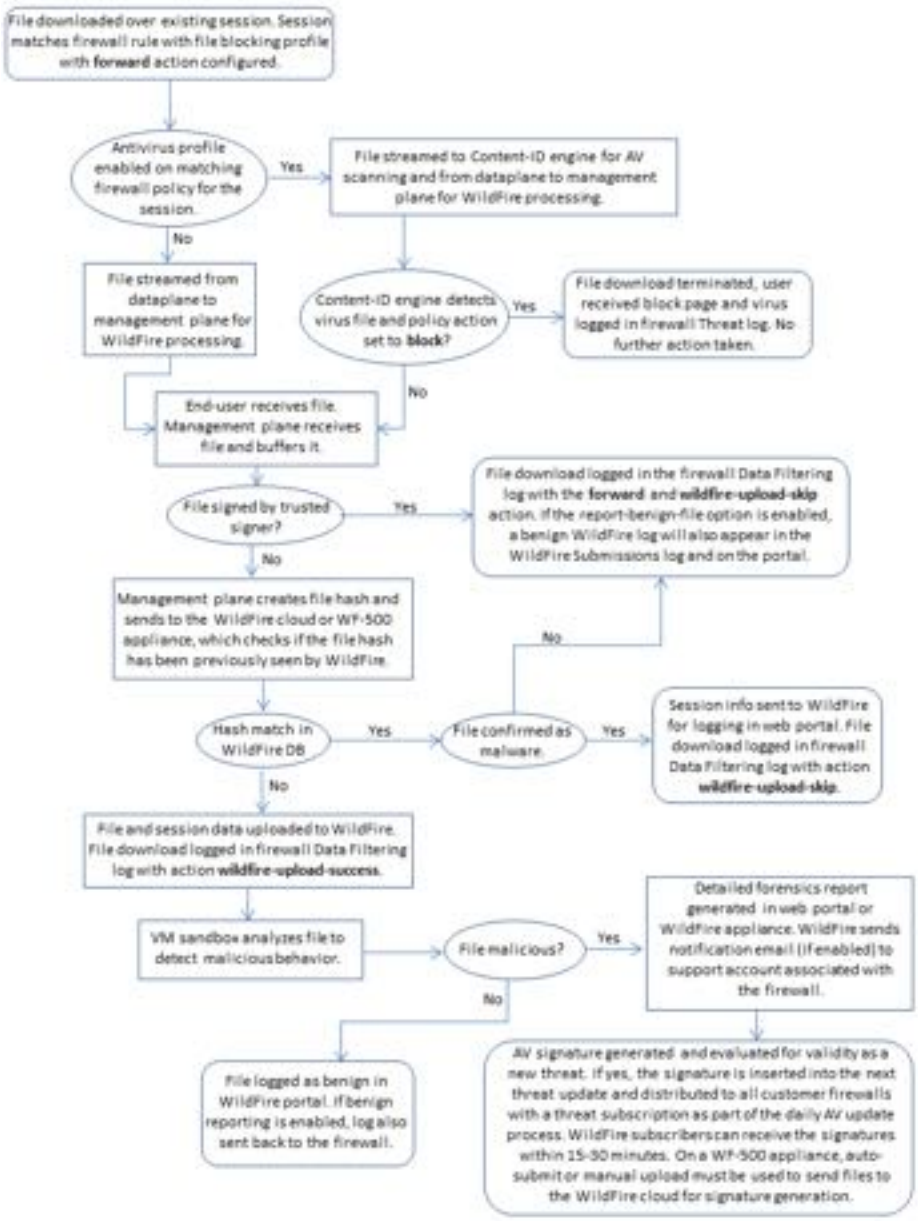


The the High-Level WildFire Decision Workflow describes the workflow for a file download. The analysis of an HTTP/HTTPS link contained in an email is very similar, but there are minor differences. For details, see Email Link Analysis.

High-Level WildFire Decision Workflow



Detailed WildFire Decision Flow



WildFire Deployments

You can set up a Palo Alto Networks firewall to submit unknown files for WildFire analysis using the following types of WildFire deployments:

- WildFire Public Cloud
- WildFire Private Cloud
- WildFire Hybrid Cloud

WildFire Public Cloud

In a WildFire public cloud deployment, a Palo Alto Networks firewall forwards files to the hosted WildFire environment that Palo Alto Networks owns and maintains. As WildFire detects new malware, it generates new signatures within 15-30 minutes. Firewalls equipped with a WildFire subscription can receive the new signatures within 15 minutes; firewalls with a Threat Prevention subscription will receive the new signatures in the next antivirus signature update within 24-48 hours.

The available WildFire cloud servers are wildfire.paloaltonetworks.com for the WildFire cloud server hosted in the United States and wildfire.paloaltonetworks.jp for the WildFire cloud hosted in Japan. You might want your firewalls to use the Japan server if you do not want benign files forwarded to the U.S. cloud servers. However, the Japan cloud does forward malicious files to the U.S. cloud servers so that a signature can be generated and distributed to WildFire users, in order to identify and block the threat. If your firewalls are located in the Japan region, you will also see faster response time for sample submissions and report generation.

Multiple virtual machines run on the WildFire cloud to represent a variety of operating systems and applications that are then used to analyze the behavior of samples files forwarded by Palo Alto Networks firewalls. The WildFire public cloud also supports multi-version analysis—application-specific files, such as Microsoft Office documents, PDFs, and media files, are analyzed by the WildFire public cloud across several versions of the application in order to identify malware intended to uniquely target specific versions of client applications.

WildFire Private Cloud

In a Palo Alto Networks private cloud deployment, Palo Alto Networks firewalls forward files to a WF-500 appliance on your corporate network that is being used to host a private cloud analysis location. The WildFire private cloud hosted on a local WF-500 appliance can receive and analyze files from up to 100 Palo Alto Networks firewalls.

With the local sandboxing of malware provided by the WildFire private cloud, benign or grayware files never leave your network. By default, the private cloud also does not send malware outside of your network; however, you can choose to automatically forward malware to the WildFire public cloud. The WildFire public cloud re-analyzes the sample, generates a signature if it is malware, and distributes the signature to global Palo Alto Networks firewalls with threat prevention or WildFire subscriptions. Alternatively, Manually Upload Files to the WildFire Portal or Use the WildFire API to submit files discovered in the private cloud for further analysis and global signature distribution.

If you do not want the WildFire private cloud to forward even malicious samples outside of your network, it is recommended that you enable the appliance to forward the malware report (and not the sample itself) to the WildFire public cloud. WildFire reports provide statistical information that helps Palo Alto Networks assess the pervasiveness and propagation of the malware. For more details, see Submit Malware or Reports from the WF-500 Appliance.

Enable Signatures and URL Generation on the WF-500 appliance to locally generate signatures for malware identified in the private cloud. Signatures generated by the WF-500 appliance are distributed to connected firewalls so that the firewalls can effectively block the malware the next time it is detected.

WildFire Hybrid Cloud

In a WildFire hybrid cloud deployment, a single firewall can forward certain samples the WildFire public cloud and other samples to a WildFire private cloud hosted by a WF-500 appliance. Configure settings on the firewall to forward files to a WildFire analysis location (either public-cloud or private-cloud) based on file type, application, and transmission direction of the file (upload or download). A WildFire hybrid cloud deployment allows the flexibility to analyze private documents locally and inside your network, while the WildFire public cloud analyzes files sourced from the Internet. For example, forward Payment Card Industry (PCI) and Protected Health Information (PHI) data exclusively to the WildFire private cloud for analysis, while forwarding Portable Executables (PEs) to the WildFire public cloud for analysis. In a WildFire hybrid cloud deployment, offloading files to the public cloud for analysis allows you benefit from a prompt verdict for files that have been previously processed in the WildFire public cloud, and also frees up the WF-500 appliance capacity to process sensitive content. Additionally, you can forward certain file types to the WildFire public cloud that are not currently supported for WF-500 appliance analysis, such as Android Application Package (APK) files.

To set up hybrid cloud forwarding, see Forward Files for WildFire Analysis.

WildFire Concepts

- Virtual Sandbox
- Verdicts
- File Type Analysis
- Email Link Analysis
- Signatures
- Alerts
- Logs and Reporting

Virtual Sandbox

WildFire executes suspect files it receives in a virtual environment and observes files behavior for signs of malicious activities, such as changes to browser security settings, injection of code into other processes, modification of files in the Windows system folder, or domains that the sample attempted to access. The WildFire public cloud also analyzes files across application versions in order to identify malware intended to uniquely target specific versions of client applications (the WildFire private cloud does not support multi-version analysis, and does not analyze application-specific files are analyzed across several versions of the application). When the WildFire engine completes file analysis, it generates a detailed forensics report that summarizes the observed behaviors and assigns a verdict of malware, benign, or grayware to the file. Similarly, WildFire will extracts links in emails messages and visits the links to determine if the corresponding web page hosts any exploits. If WildFire detects malicious behavior, it generates a report and submits the URL to PAN-DB and categorizes the URL as malware.

WildFire includes sandbox support for the following operating system environments:

- Microsoft Windows XP 32-bit
- Microsoft Windows 7 32-bit (Supported as an option for WF-500 appliance only)
- Microsoft Windows 7 64-bit

Verdicts

WildFire delivers verdicts to identify samples it analyzes as safe, malicious, or unwanted (grayware is considered obtrusive but not malicious):

- **Benign**—The sample is safe and does not exhibit malicious behavior.
- **Grayware**—The sample does not pose a direct security threat, but might display otherwise obtrusive behavior. Grayware typically includes adware, spyware, and Browser Helper Objects (BHOs).

- **Malware**—The sample is malicious in nature or intent and poses a security threat. Malware can include viruses, worms, Trojans, Remote Access Tools (RATs), rootkits, and botnets. For files identified as malware, WildFire generates and distributes a signature to prevent against future exposure to the threat.

File Type Analysis

A Palo Alto Networks firewall can be configured with a WildFire analysis profile in order to forward samples for WildFire analysis based on file type. If a user downloads a file sample over a session that matches the security rule to which the WildFire analysis profile is attached, the firewall performs a file hash check with WildFire to determine if WildFire has previously analyzed the file. If the file is unknown, the firewall forwards it to WildFire.

The firewall can forward files for WildFire analysis based on the following file types:

- **apk**—Android Application Package (APK) files. APK files are not supported for WildFire private cloud analysis using a WF-500 appliance.
- **email-link**—HTTP/HTTPS email links contained in SMTP and POP3 email messages.
- **flash**—Adobe Flash applets and Flash content embedded in web pages.
- **jar**—Java applets (JAR/class files types).
- **ms-office**—Microsoft Office files, including documents (DOC, DOCX, RTF), workbooks (XLS, XLSX), and PowerPoint (PPT, PPTX) presentations, and Office Open XML (OOXML) 2007+ documents.
- **pe**—Portable Executable (PE) files. PEs include executable files, object code, DLLs, and FON (fonts). A subscription is not required to forward PE files for WildFire analysis, but is required for all other supported file types.
- **pdf**—Portable Document Format (PDF) files.

For details on enabling the firewall to forward samples to WildFire based on file type, see Forward Files for WildFire Analysis.

Email Link Analysis

A Palo Alto Networks firewall can extract HTTP/HTTPS links contained in SMTP and POP3 email messages and forward the links to the WildFire public or private clouds for analysis. Enable forwarding of unknown links contained in emails by configuring a WildFire analysis profile with the **email-link** file type. The firewall only extracts links and associated session information (sender, recipient, and subject) from the email messages that traverse the firewall; it does not receive, store, forward, or view the email message.

After receiving a link from a firewall, WildFire visits the link to determine if the corresponding web page hosts any exploits. If WildFire determines that the page itself is benign or grayware, it will not generate a log. However, if it detects malicious behavior on the page, the firewall returns a malicious verdict and:

- Generates a detailed analysis report and logs the report to the **WildFire Submissions** log on the firewall. The log includes the email header information—email sender, recipient, and subject—so that you can identify the message and delete it from the mail server, or mitigate the threat if the email has already been delivered or opened.
- Adds the URL to PAN-DB and categorizes the URL as malware.

The firewall forwards email links to WildFire in batches of 100 email links or every two minutes (depending on which limit is hit first). Each batch upload to WildFire counts as one upload toward the upload per-minute capacity for the given firewall platform (Firewall File Forwarding Capacity by Platform).

If a link included in an email corresponds to a file download instead of a URL, the firewall forwards the file to WildFire for analysis only if corresponding file type is enabled for WildFire analysis.

For details on enabling the firewall to forward links included in emails to a WildFire cloud, see Forward Files for WildFire Analysis.

Signatures

The key benefits of the Palo Alto Networks WildFire feature are that it can discover zero-day malware in web traffic (HTTP/HTTPS), email protocols (SMTP, IMAP, and POP), and FTP traffic and can quickly generate signatures to protect against future infections from the malware it discovers. WildFire will automatically generate a signature based on the malware payload of the sample and tests it for accuracy and safety. Because malware evolves rapidly, the signatures that WildFire generates will address multiple variants of the malware. As WildFire detects new malware, it generates new signatures within 15-30 minutes. Firewalls equipped with a WildFire subscription can receive the new signatures within 15 minutes. If you do not have a WildFire subscription, signatures are made available within 24-48 hours as part of the antivirus update for firewalls equipped with a Threat Prevention subscription.

As soon as the firewall downloads and installs the new signature, the firewall drops any files that contain that malware (or a variant of the malware). Information gathered by WildFire during the analysis of malware is used to fortify other Threat Prevention features, such as adding malware URLs to PAN-DB, generating DNS signatures, antivirus, and anti-spyware signatures. Palo Alto Networks also develops signatures for command and control traffic, enabling immediate disruption in the communications of any malware inside the network. For details on signatures and the benefits of having a WildFire subscription, see WildFire Subscription Requirements.

Alerts

Enable the firewall to provide instant notifications when it detects malware on your network. Set Up Alerts for Malware Logs to be delivered as email alerts, syslog messages, or SNMP traps. Alert notifications enable you to quickly identify the user who downloaded the malware and eradicate the malware before it causes extensive damage or propagates to other users. In addition, every signature that WildFire generates is automatically propagated to all Palo Alto Networks firewalls protected with a Threat Prevention and/or WildFire subscription, which provides automatic protection from malware discovered on networks all over the world.

Logs and Reporting

You can Monitor WildFire Activity using a firewall submitting samples for analysis, the WildFire portal, or the WildFire API.

For each sample WildFire analyzes, WildFire generates a WildFire analysis report that categorizes the sample as malware, grayware, or benign and details sample information and behavior.

WildFire analysis reports can be accessed on the firewall that submitted the sample and the WildFire cloud (public or private) that analyzed the sample:

View Reports for Samples Submitted to the WildFire Public Cloud

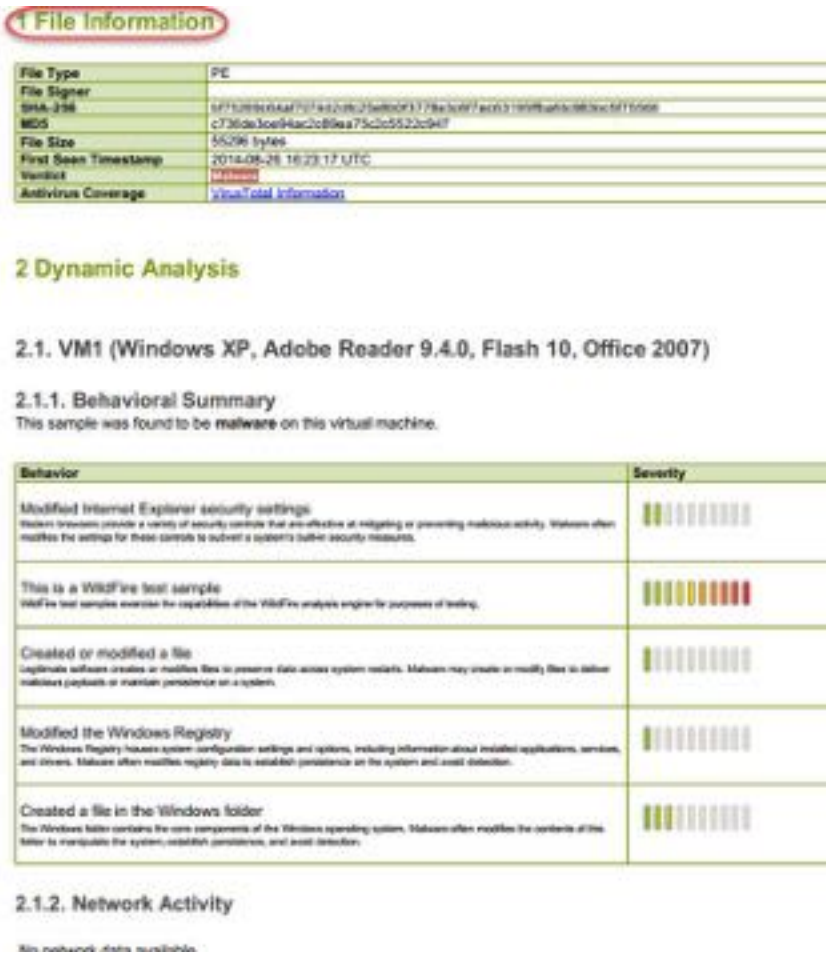
- Use the firewall—All samples submitted by a firewall for WildFire analysis are logged as **WildFire Submissions** entries (**Monitor > WildFire Submissions**). For each WildFire submission entry you can open a detailed log view to view the WildFire analysis report for the sample or to download the report as a PDF.
- Use the WildFire portal—Monitor WildFire activity, including the WildFire analysis report for each sample, which can also be downloaded as a PDF. In a WildFire public cloud deployment, details are available on the WildFire portal for samples submitted to the public cloud by connected firewalls and samples that are manually uploaded to the portal.

View Reports for Samples Submitted to the WildFire Private Cloud

- Use the firewall—All samples submitted by a firewall for WildFire analysis are logged as WildFire Submissions entries (Monitor > WildFire Submissions). For each WildFire submission entry you can open a detailed log view to view the WildFire analysis report for the sample or to download the report as a PDF.
- Use the WildFire portal—(Only when cloud intelligence is enabled on the WF-500 appliance) Monitor WildFire activity, including the WildFire analysis report for each sample, which can also be downloaded as a PDF. In a WildFire private cloud deployment, the WildFire portal provides details for samples that are manually uploaded to the portal and samples submitted by a WF-500 appliance with cloud intelligence enabled.
- Use the WildFire API—Retrieve WildFire analysis reports from a WF-500 appliance.

WildFire analysis reports show detailed behavioral information about the sample, information on the targeted user, the application that delivered the file, and all URLs involved in the delivery or phone-home activity of the file. For details on report fields, see WildFire Analysis Reports—Close Up.

The following screen capture shows part of a WildFire analysis report for a file submitted to the WildFire cloud, followed by a second screen capture showing a WildFire analysis report for an email link.



1 File Information

File Type	PE
File Signer	
SHA-256	
MD5	c736de3ce94ac2c89ea75c2c5522c947
File Size	55296 bytes
First Seen Timestamp	2014-06-26 16:23:17 UTC
Verdict	Malware
Antivirus Coverage	VirusTotal Information

2 Dynamic Analysis

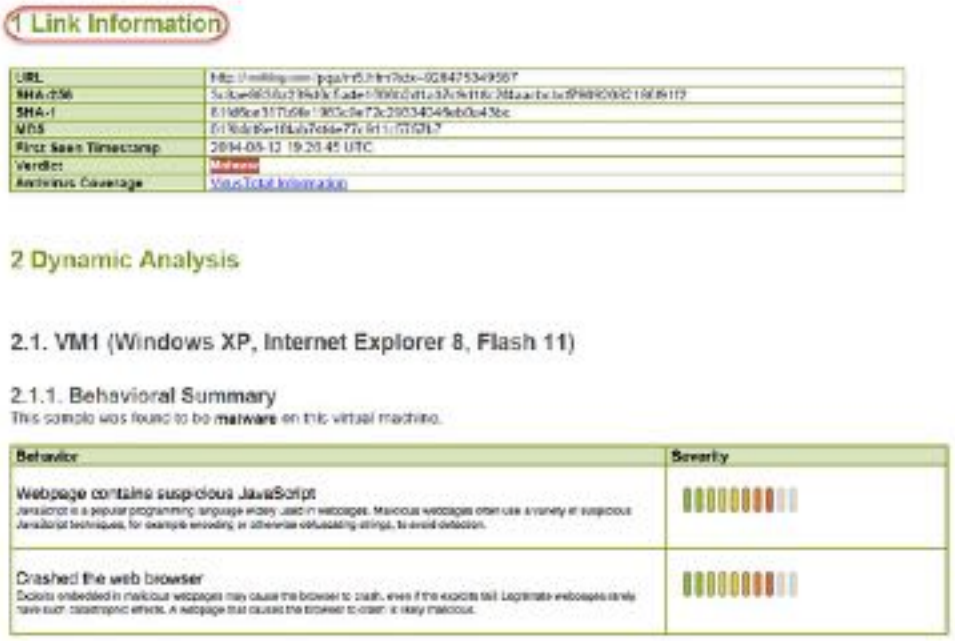
2.1. VM1 (Windows XP, Adobe Reader 9.4.0, Flash 10, Office 2007)

2.1.1. Behavioral Summary

This sample was found to be **malware** on this virtual machine.

Behavior	Severity
Modified Internet Explorer security settings	
This is a WildFire test sample	
Created or modified a file	
Modified the Windows Registry	
Created a file in the Windows folder	

2.1.2. Network Activity



1 Link Information

URL	
SHA-256	
SHA-1	
MD5	
First Seen Timestamp	2014-05-12 19:20:45 UTC
Verdict	Malware
Antivirus Coverage	VirusTotal Information

2 Dynamic Analysis

2.1. VM1 (Windows XP, Internet Explorer 8, Flash 11)

2.1.1. Behavioral Summary

This sample was found to be **malware** on this virtual machine.

Behavior	Severity
Webpage contains suspicious JavaScript JavaScript is a popular programming language widely used in webpages. Malicious webpages often use unusual or suspicious JavaScript techniques, for example encoding or otherwise obfuscating strings, to avoid detection.	
Crashed the web browser Exploits embedded in malicious webpages may cause the browser to crash, even if the exploits fail. Legitimate webpages rarely have such catastrophic effects. A webpage that caused the browser to crash is likely malicious.	

WildFire Subscription Requirements

WildFire provides detection and prevention of zero-day malware using a combination of malware sandboxing and signature-based detection of malware. No subscription is required to use WildFire to sandbox files sent from Palo Alto Networks firewalls to the WildFire cloud.

For the firewall to perform detection and blocking of known malware discovered by WildFire, the firewall requires a Threat Prevention and/or WildFire subscription. The Threat Prevention subscription enables the firewall to receive daily antivirus signature updates, which provide coverage for all malware samples that WildFire discovers globally. The Threat Prevention subscription also provides access to weekly content updates that include new vulnerability protection and anti-spyware signatures.

To enable a WF-500 appliance for local analysis, you only need to install a support license. This enables the appliance to communicate with the Palo Alto Networks update server to download the operating system images and daily content updates. The content updates support the ability to generate signatures on the local WF-500 appliance and equip the appliance with the most up-to-date threat information for accurate malware detection and to improve the ability of the appliance to differentiate the malicious from the benign.

To receive the full benefits of the WildFire service, each firewall connected to the WildFire public cloud or private cloud must have a WildFire subscription, which includes the following:

- **WildFire Dynamic Updates**—Provides new malware signatures on a sub-hourly basis, configurable through **Device > Dynamic Updates**. Within 15-30 minutes after WildFire identifies a malicious sample, WildFire generates a new malware signature and distributes it through the WildFire dynamic updates, which the firewall can poll every 15, 30, or 60 minutes. You can configure the firewall to take specific actions on malware signatures separate from the regular antivirus signature actions in the antivirus profile. The WildFire signatures delivered in the dynamic update include signatures generated for malware detected in files submitted to WildFire by all Palo Alto Networks WildFire customers, not just the file samples that your firewalls send to WildFire.

- **WildFire Advanced File Type Support**—Allows the firewall to forward samples for WildFire analysis based on file type (see File Type Analysis) and to configure the firewall to extract and forward links included in email messages for WildFire Email Link Analysis.

- **WildFire API**—Provides access to the WildFire API, which enables direct programmatic access to the WildFire service on the Palo Alto Networks WildFire cloud or a WF-500 appliance. Use the WildFire API to submit files for analysis and to retrieve the subsequent WildFire analysis reports. The WildFire API supports up to 1,000 file submissions and up to 10,000 queries a day.

Only firewalls with a valid WildFire subscription can forward files to a WF-500 appliance for private cloud analysis.

Best Practices for Keeping Signatures up to Date

This section describes the best practices for keeping a firewall with Threat Prevention and WildFire subscriptions up-to-date with the latest protection. For a streamlined workflow, use Panorama to push dynamic update schedules to managed firewalls using Panorama templates. This ensures consistency across all firewalls and simplifies management of update schedules.

These guidelines provide two schedule options: the minimum recommended schedule and a more aggressive schedule. Choosing the more aggressive approach causes the device to perform download and install updates much more frequently, some of which can be very large (over 100MB for antivirus updates).



You can stagger automatic updates by using the **Threshold (Hours)** field to specify how long after a release to wait before performing a content update.

- **Antivirus**—New antivirus content updates are released by Palo Alto Networks on a daily basis. To get the latest content, schedule these updates daily at minimum. For a more aggressive schedule, schedule them hourly.
- **Applications and Threats**—New App-ID, vulnerability protection, and anti-spyware signatures are released by Palo Alto Networks as weekly content updates (normally on Tuesdays). To receive the latest content, schedule the updates at least weekly. For a more aggressive schedule to ensure that the firewall receives the latest content soon after the release (including occasional off-schedule emergency content releases), schedule the firewall to download/install daily.
- **WildFire**—New WildFire antivirus signatures are published every 15 minutes. Depending on when WildFire discovers new malware within the release cycle, coverage is provided in the form of a WildFire signature 15-30 minutes after it is discovered. To get the latest WildFire signatures, schedule these updates every hour or half-hour. For a more aggressive schedule, configure the firewall to check for updates every 15 minutes.
- **WF-Private**—If the WF-500 appliance is configured to generate signatures and URL categories (antivirus signatures, DNS signatures, and URL entries for PAN-DB), configure the firewall to download/install the updates using the WF-Private dynamic update. After the appliance receives a malicious sample, it will generate a signature within five minutes in most cases. When configuring the firewall to retrieve these updates, set the schedule to download and install every hour or half-hour. For a more aggressive schedule (recommended), configure the firewall to download and install the updates every 5 minutes. If you configure your firewalls to retrieve WF-Private updates, it is highly recommended that the firewalls also download content updates from Palo Alto Networks (Antivirus, Applications/Threats, and WildFire) to ensure that firewalls have the latest protection. This is important due to the fact that when the local storage for WF-Private updates on the appliance is full, new signatures/URL categorizations will overwrite existing ones, beginning with the oldest ones first. For details on local signatures generation, see Enable Signatures and URL Generation.



Configure the WF-500 Appliance

This topic describes how to configure a WF-500 appliance in order to host a WildFire private cloud to analyze files on your network. The following topics describe readying the WF-500 appliance to receive files for analysis, how to manage the appliance, and how to enable the appliance to locally generate threat signatures and URL categories.

- About the WF-500 Appliance
- Configure the WF-500 Appliance
- Set Up the VM Interface
- Set Up WF-500 Content Updates
- Enable Signatures and URL Generation
- Upgrade a WF-500 Appliance

About the WF-500 Appliance

The WF-500 appliance provides an on-premises WildFire™ private cloud, enabling you to analyze suspicious files in a sandbox environment without requiring the firewall to sends files out of network. To use the WF-500 appliance to host a WildFire private cloud, configure the firewall to submit samples to the WF-500 appliance for analysis. The WF-500 appliance sandboxes all files locally and analyzes them for malicious behaviors using the same engine used by the WildFire public cloud. Within minutes, the private cloud returns analysis results to the firewall **WildFire Submissions** logs.

The WF-500 appliance has a cloud intelligence feature that you can enable (it is disabled by default) to send confirmed malware to the public cloud for signature generation. You can also configure the cloud intelligence feature to send only reports on malware, allowing Palo Alto Networks to gather statistics on malware. It is recommended that you configure the appliance to send malware samples to the WildFire public cloud, so that signatures are generated and distributed to protect global users. If you do not want to automatically send detected malware to the WildFire public cloud for signature generation, you can manually upload to the WildFire Portal.

You can also configure the WF-500 appliance to generate signatures locally and the connected firewalls can retrieve the updates directly from the appliance. For information on configuring local signature generation and to learn about the types of content updates that the appliance can provide, see Enable Signatures and URL Generation.

Configure up to 100 Palo Alto Networks firewalls with valid WildFire subscriptions to forward to a single WF-500 appliance.

The WF-500 appliance has two interfaces:

- **MGT**—Receives all files forwarded from the firewalls and returns logs detailing the results back to the firewalls. See Configure the WF-500 Appliance.
- **Virtual Machine Interface (VM interface)**—Provides network access for the WildFire sandbox systems to enable sample files to communicate with the Internet, which allows WildFire to better analyze the behavior of the sample. When the VM interface is configured, WildFire can observe malicious behaviors that the malware would not normally perform without network access, such as phone-home activity. However, to prevent malware from entering your network from the sandbox, configure the VM interface on an isolated network with an Internet connection. You can also enable the Tor option to hide the public IP address used by your company from malicious sites that are accessed by the sample. For more information on the VM interface, see Set Up the VM Interface.

Configure the WF-500 Appliance

Configure the WF-500 Appliance

This section describes the steps required to integrate a WF-500 appliance into a network and perform basic setup.

Configure the WF-500 Appliance	
Before you begin: ❑ Rack mount and cable the WF-500 appliance. Refer to the WildFire Appliance Hardware Reference Guide. ❑ Obtain the information required to configure network connectivity on the MGT port and the VM interface from your network administrator (IP address, subnet mask, gateway, hostname, DNS server). All communication between the firewalls and the appliance occurs over the MGT port, including file submissions, WildFire log delivery, and appliance administration. Therefore, ensure that the firewalls have connectivity to the MGT port on the appliance. In addition, the appliance must be able to connect to the `updates.paloaltonetworks.com` site to retrieve its operating system software updates. ❑ Have a computer ready with either a console cable or Ethernet cable to connect to the device for the initial configuration.	
Step 1 Connect the management computer to the appliance using the MGT or Console port and power on the appliance.	1. Connect to the console port or the MGT port. Both are located on the back of the appliance. • **Console Port**—This is a 9-pin male serial connector. Use the following settings on the console application: 9600-8-N-1. Connect the provided cable to the serial port on the management computer or USB-To-Serial converter. • **MGT Port**—This is an Ethernet RJ-45 port. By default, the MGT port IP address is 192.168.1.1. The interface on your management computer must be on the same subnet as the MGT port. For example, set the IP address on the management computer to 192.168.1.5. 2. Power on the appliance. The appliance will power on as soon as you connect power to the first power supply and a warning beep will sound until you connect the second power supply. If the appliance is already plugged in and is in the shutdown state, use the power button on the front of the appliance to power on.

Configure the WF-500 Appliance (Continued)	

Step 2 Register the WF-500 appliance.	1. Obtain the serial number from the S/N tag on the appliance, or run the following command and refer to the `serial` field:
	`admin@WF-500> **show system info**`
	2. From a browser, navigate to the Palo Alto Networks Support site.
	3. Register the device as follows:
	• If this is the first Palo Alto Networks device that you are registering and you do not yet have a login, click **Register** on the right side of the page. To register, provide an email address and the serial number of the device. When prompted, set up a username and password for access to the Palo Alto Networks support community.
	• For existing accounts, log in and then click **My Devices**. Scroll down to the **Register Device** section at the bottom of the screen and enter the serial number of the device, the city and postal code, and then click **Register Device**.
	4. To confirm WildFire registration on the WF-500 appliance, log in to the appliance with an SSH client or by using the Console port. Enter a username/password of admin/admin and enter the following command on the appliance:
	`admin@WF-500> **test wildfire registration**`
	The following output indicates that the appliance is registered with one of the Palo Alto Networks WildFire cloud servers.
	`Test wildfire`
	`wildfire registration: successful`
	`download server list: successful`
	`select the best server:`
	`**cs-s1.wildfire.paloaltonetworks.com**`

Step 3 Reset the admin password.	1. Set a new password by running the command: `admin@WF-500# **set password**` 2. Type the old password, press enter and then enter and confirm the new password. There is no need to commit the configuration because this is an operational command. 3. Type `exit` to log out and then log back in to confirm that the new password is set.

Configure the WF-500 Appliance (Continued)		

Step 4	Configure the management interface settings.
	
	This example uses the following values:
	• IPv4 address - 10.10.0.5/22
	• Subnet Mask - 255.255.252.0
	• Default Gateway - 10.10.0.1
	• Hostname - wildfire-corp1
	• DNS Server - 10.0.0.246
		1. Log in to the appliance with an SSH client or by using the Console port and enter configuration mode:
		`admin@WF-500> configure`
		2. Set the IP information:
		`admin@WF-500# set deviceconfig system ip-address 10.10.0.5 netmask 255.255.252.0 default-gateway 10.10.0.1 dns-setting servers primary 10.0.0.246`
		Configure a secondary DNS server by replacing primary with secondary in the above command, excluding the other IP parameters. For example:
		`admin@WF-500# set deviceconfig system dns-setting servers secondary 10.0.0.247`
		3. Set the hostname (wildfire-corp1 in this example):
		`admin@WF-500# set deviceconfig system hostname wildfire-corp1`
		4. Commit the configuration to activate the new management (MGT) port configuration:
		`admin@WF-500# commit`
		5. Connect the MGT interface port to a network switch.
		6. Put the management PC back on your corporate network, or whatever network is required to access the appliance on the management network.
		7. From your management computer, use an SSH client to connect to the new IP address or hostname assigned to the MGT port on the appliance. In this example, the IP address is 10.10.0.5.

Step 5	Activate the appliance with the WildFire authorization code that you received from Palo Alto Networks.
	
	Though it will function without an auth-code, the WF-500 appliance cannot retrieve software or content updates without a valid auth-code.
		1. Change to operational mode:
		`admin@WF-500# exit`
		2. Fetch and install the WildFire license:
		`admin@WF-500> request license fetch auth-code <auth-code>`
		3. Verify the license:
		`admin@WF-500> request support check`
		Information about the support site and the support contract date is displayed. Confirm that the date displayed is valid.

Configure the WF-500 Appliance (Continued)

Step 6 Manually set the current date, time, and timezone, or synchronize the WF-500 appliance local clock with a Network Time Protocol (NTP) server.

Manually set the WF-500 appliance date, time, and timezone:

1. Set the date and time:

```
admin@WF-500> set clock date <YY/MM/DD> time <hh:mm:ss>
```

2. Enter configuration mode:

```
admin@WF-500> configure
```

3. Set the local time zone:

```
admin@WF-500# set deviceconfig system timezone <timezone>
```



The time stamp that will appear on the WildFire detailed report will use the time zone set on the appliance. If administrators in various regions will view reports, consider setting the time zone to UTC.

Synchronize the WF-500 appliance local clock with a Network Time Protocol (NTP) server:

1. Enter configuration mode:

```
admin@WF-500> configure
```

2. Enter the IP address of the NTP server you want to use to synchronize the WF-500 local clock:

```
admin@WF-500# set deviceconfig system ntp-servers primary-ntp-server ntp-server-address <server IP address>
```

3. (Optional) Enter the IP address of a secondary NTP server:

```
admin@WF-500# set deviceconfig system ntp-servers secondary-ntp-server ntp-server-address <server IP address>
```



The WF-500 appliance does not prioritize the primary or secondary NTP server; it synchronizes with either server.

4. Set up NTP authentication:
   - Disable NTP authentication:

     ```
     admin@WF-500# set deviceconfig system ntp-servers primary-ntp-server authentication-type none
     ```

   - Enable symmetric key exchange (shared secrets) to authenticate the NTP server time updates:

     ```
     admin@WF-500# set deviceconfig system ntp-servers primary-ntp-server authentication-type symmetric-key
     ```

     Continue to enter the **key-ID** (1 - 65534), choose the **algorithm** to use in NTP authentication (**MD5** or **SHA1**), and then enter and confirm the authentication algorithm **authentication-key**.

   - Use autokey (public key cryptography) to authenticate the NTP server time updates:

     ```
     admin@WF-500# set deviceconfig system ntp-servers primary-ntp-server authentication-type autokey
     ```

Configure the WF-500 Appliance (Continued)	

Step 7 Choose the virtual machine image for the appliance to use to analyze files.
 	
The image should be based on the attributes that most accurately represent the software installed on your end user computers. Each virtual image contains different versions of operating systems and software, such as Windows XP or Windows 7 32-bit or 64-bit and specific versions of Adobe Reader, and Flash. Although you configure the appliance to use one virtual machine image configuration, the appliance uses multiple instances of the image to improve performance.
	• To view a list of available virtual machines to determine which one best represents your environment:
	`admin@WF-500> show wildfire vm-images`
	• View the current virtual machine image by running the following command and refer to the `Selected VM` field:
	`admin@WF-500> show wildfire status`
	• Select the image that the appliance will use for analysis:
	`admin@WF-500# set deviceconfig setting wildfire active-vm <vm-image-number>`
	For example, to use vm-1:
	`admin@WF-500# set deviceconfig setting wildfire active-vm vm-1`

Continue to configure the following optional settings for the WF-500 appliance:	
• Enable the WF-500 appliance cloud intelligence feature to automatically submit malware samples discovered in the WildFire private cloud to the WildFire public cloud. The WildFire public cloud re-analyzes the sample and generates a signature if the sample is malware. The signature is added to WildFire signature updates to be distributed to global users.	Submit Malware to the WildFire Public Cloud.
• If you do not want to forward malware samples outside of the WildFire private cloud, instead submit WildFire analysis reports for malware to the WildFire public cloud. If the WF-500 appliance is not enabled to submit malware to the WildFire public cloud, it is a best practice to enable malware analysis report submissions to contribute to and improve WildFire threat intelligence.	Submit Analysis Reports to the WildFire Public Cloud.

• Configure additional user accounts for managing the WF-500 appliance.
 	
You can assign two role types: superuser and superreader. Superuser is equivalent to the admin account, and superreader only has read access.
	In this example, you will create a superreader account for the user bsimpson:
	1. Enter configuration mode:
	`admin@WF-500> configure`
	2. Create the user account:
	`admin@WF-500# set mgt-config users bsimpson <password>`
	3. Enter and confirm a new password.
	4. Assign the superreader role:
	`admin@WF-500# set mgt-config users bsimpson permissions role-based superreader yes`

Configure the WF-500 Appliance (Continued)	

• Configure RADIUS authentication for administrator access.	1. Create a RADIUS profile using the following options:
	`admin@WF-500# set shared server-profile radius <profile-name>`
	(Configure the RADIUS server and other attributes.)
	2. Create an authentication profile:
	`admin@WF-500# set shared authentication-profile <profile-name> method radius server-profile <server-profile-name>`
	3. Assign the profile to a local admin account:
	`admin@WF-500# set mgt-config users username authentication-profile authentication-profile-name>`

• Set a password for the WildFire portal admin account.  The WildFire portal admin can use the portal to manually submit samples for WildFire public cloud analysis, and to view analysis reports for these samples (see Manually Upload Files to the WildFire Portal). If the WF-500 appliance is enabled to automatically Submit Malware to the WildFire Public Cloud, analysis reports for these samples can also be accessed using the WildFire portal.	Only the password can be changed for the WildFire portal admin account, and additional admin accounts cannot be created. The default username and password is admin/admin. 1. To change the WildFire portal admin account password: `admin@WF-500> set wildfire portal-admin password` 2. Press enter and type and confirm the new password.
Next Steps...	• Set Up the VM Interface to enable the WF-500 appliance to observe malicious behaviors where the file being analyzed seeks network access. • Use the WildFire API to retrieve WildFire analysis reports for samples analyzed in a WildFire private cloud. • If you have enabled cloud intelligence to Submit Malware to the WildFire Public Cloud for further analysis, use the WildFire portal to view reports for submitted malware.

Set Up the VM Interface

The virtual machine interface (vm-interface) provides external network connectivity from the sandbox virtual machines in the WF-500 appliance to enable observation of malicious behaviors in which the file being analyzed seeks network access. The following sections describe the VM interface and the steps required for configuring it. You can optionally enable the Tor feature with the VM interface, which will mask any malicious traffic sent from the WF-500 appliance through the VM interface, so the malware sites that the traffic may be sent to cannot detect your public-facing IP address.

This section also describes the steps required to connect the VM interface to a dedicated port on a Palo Alto Networks firewall to enable Internet connectivity.

- Virtual Machine Interface Overview
- Configure the VM Interface on the WF-500 Appliance
- Connect the Firewall to the WF-500 Appliance VM Interface

Virtual Machine Interface Overview

The VM interface (labeled **1** on the back of the appliance) is used by WildFire to improve malware detection capabilities. The interface allows a file sample running on the WildFire virtual machines to communicate with the Internet and enables WildFire to better analyze the behavior of the sample file to determine if it exhibits characteristics of malware.

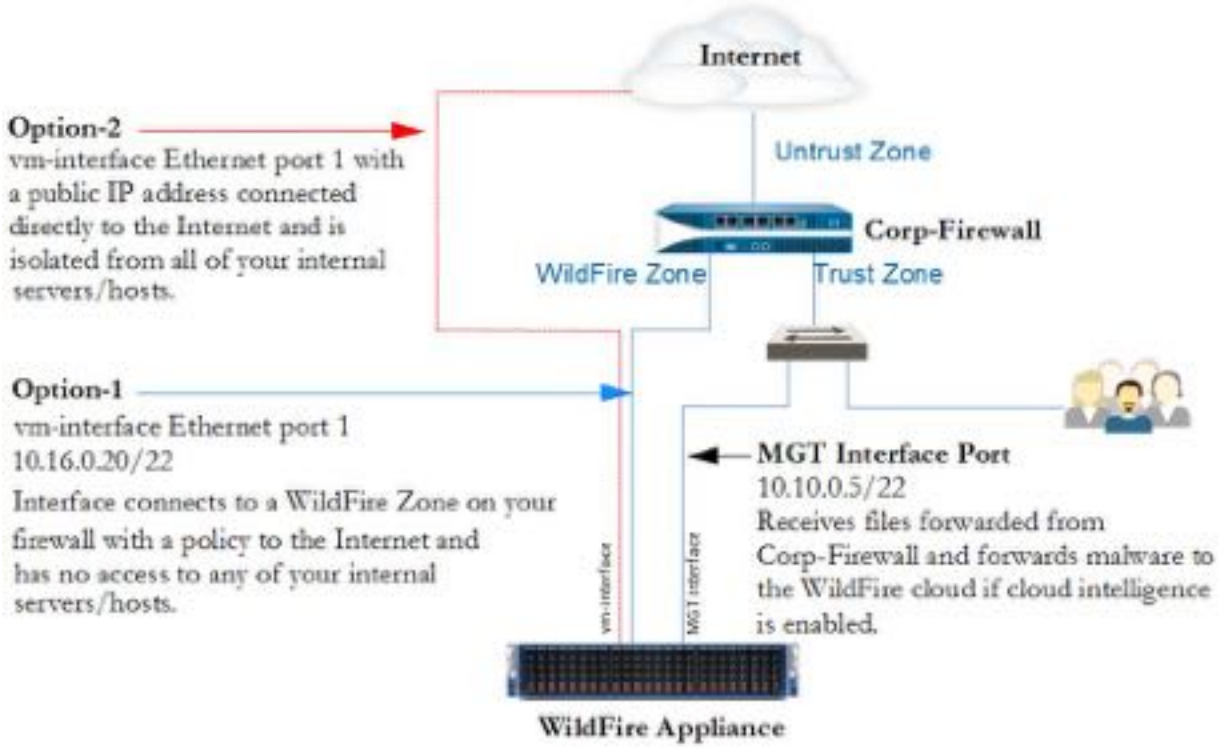


While it is recommended that you enable the VM interface, it is very important that you do not connect the interface to a network that allows access to any of your servers/hosts because malware that runs in the WildFire virtual machines could potentially use this interface to propagate itself.

This connection can be a dedicated DSL line or a network connection that only allows direct access from the VM interface to the Internet and restricts any access to internal servers/client hosts.

The following illustration shows two options for connecting the VM interface to the network.

Virtual Machine Interface Example



- Option-1 (recommended)—Connect the VM interface to an interface in a dedicated zone on a firewall that has a policy that only allows access to the Internet. This is important because malware that runs in the WildFire virtual machines can potentially use this interface to propagate itself. This is the recommended option because the firewall logs will provide visibility into any traffic that is generated by the VM interface.
- Option-2—Use a dedicated Internet provider connection, such as a DSL, to connect the VM interface to the Internet. Ensure that there is no access from this connection to internal servers/hosts. Although this is a simple solution, traffic generated by the malware out the VM interface will not be logged unless you place a firewall or a traffic monitoring tool between the WF-500 appliance and the DSL connection.

Configure the VM Interface on the WF-500 Appliance

This section describes the steps required to configure the VM interface on the WF-500 appliance using the Option 1 configuration detailed in the Virtual Machine Interface Example. After configuring the VM interface using this option, you must also configure an interface on a Palo Alto Networks firewall through which traffic from the VM interface is routed as described in Connect the Firewall to the WF-500 Appliance VM Interface.

By default, the VM interface has the following settings:

- IP Address: 192.168.2.1
- Netmask: 255.255.255.0
- Default Gateway: 192.168.2.254

- DNS: 192.168.2.254

If you plan on enabling this interface, configure it with the appropriate settings for your network. If you do not plan on using this interface, leave the default settings. Note that this interface must have network values configured or a commit failure will occur.

Configure the VM Interface		

Step 1	Set the IP information for the VM interface on the WF-500 appliance. The following settings are used in this example:
	• IPv4 address - 10.16.0.20/22
	• Subnet Mask - 255.255.252.0
	• Default Gateway - 10.16.0.1
	• DNS Server - 10.0.0.246
	The VM interface cannot be on the same network as the management interface (MGT).
		1. Enter configuration mode:
		`admin@WF-500> configure`
		2. Set the IP information for the VM interface:
		`admin@WF-500# set deviceconfig system vm-interface ip-address 10.16.0.20 netmask 255.255.252.0 default-gateway 10.16.0.1 dns-server 10.0.0.246`
		You can only configure one DNS server on the VM interface. As a best practice, use the DNS server from your ISP or an open DNS service.

Step 2	Enable the VM interface.	1. Enable the VM interface:
		`admin@WF-500# set deviceconfig setting wildfire vm-network-enable yes`
		2. Commit the configuration:
		`admin@WF-500# commit`

Step 3	Test connectivity of the VM interface.	Ping a system and specify the VM interface as the source. For example, if the VM interface IP address is 10.16.0.20, run the following command where *ip-or-hostname* is the IP or hostname of a server/network that has ping enabled:
		`admin@WF-500> ping source 10.16.0.20 host ip-or-hostname`
		For example:
		`admin@WF-500> ping source 10.16.0.20 host 10.16.0.1`

Step 4	(Optional) Enable the Tor network. When this option is enabled, any malicious traffic that the malware generates to the Internet is sent to the Tor network. The Tor network will mask your public facing IP address, so the owners of the malicious site cannot determine the source of the traffic.	Enable the Tor network: 1. `admin@WF-500# set deviceconfig setting wildfire vm-network-use-tor` 2. Commit the configuration: `admin@WF-500# commit`
Next Step...		Connect the Firewall to the WF-500 Appliance VM Interface.

Connect the Firewall to the WF-500 Appliance VM Interface

The following example workflow describes how to connect the VM interface to a port on a Palo Alto Networks firewall. Before connecting the VM interface to the firewall, the firewall must already have an Untrust zone connected to the Internet. In this example, you configure a new zone named wf-vm-zone that will contain the interface used to connect the VM interface on the appliance to the firewall. The policy associated with the wf-vm-zone will only allow communication from the VM interface to the Untrust zone.

Configure the Firewall to Control Traffic for the WF-500 Appliance VM Interface		
Step 1	Configure the interface on the firewall that the VM interface will connect to and set the virtual router. The wf-vm-zone should only contain the interface (ethernet1/3 in this example) used to connect the VM interface on the appliance to the firewall. This is done to avoid having any traffic generated by the malware from reaching other networks.	1. From the web interface on the firewall, select **Network > Interfaces** and then select an interface, for example **Ethernet1/3**. 2. In the **Interface Type** drop-down, select **Layer3**. 3. On the **Config** tab, from the **Security Zone** drop-down box, select **New Zone**. 4. In the Zone dialog **Name** field, enter wf-vm-zone and click **OK**. 5. In the **Virtual Router** drop-down box, select **default**. 6. 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 10.16.0.0/22. 7. To save the interface configuration, click **OK**.
Step 2	Create a security policy on the firewall to allow access from the VM interface to the Internet and block all incoming traffic. In this example, the policy name is WildFire VM Interface. Because you will not create a security policy from the Untrust zone to the wf-vm-interface zone, all inbound traffic is blocked by default.	1. Select **Policies > Security** and click **Add** 2. In the **General** tab, enter a **Name**. 3. In the **Source** tab, set the **Source Zone** to wf-vm-zone. 4. In the **Destination** tab, set the **Destination Zone** to **Untrust**. 5. In the **Application** and **Service/ URL Category** tabs, leave the default as **Any**. 6. In the **Actions** tab, set the **Action Setting** to **Allow**. 7. Under **Log Setting**, select the **Log at Session End** check box. If there are concerns that someone might inadvertently add other interfaces to the wf-vm-zone, clone the WildFire VM Interface security policy and then in the **Action** tab for the cloned rule, select **Deny**. Make sure this new security policy is listed below the WildFire VM interface policy. This will override the implicit intra-zone allow rule that allows communications between interfaces in the same zone and will deny/block all intra-zone communication.
Step 3	Connect the cables.	Physically connect the VM interface on the WF-500 appliance to the port you configured on the firewall (Ethernet 1/3 in this example) using a straight through RJ-45 cable. The VM interface is labeled **1** on the back of the appliance.

Configure the Firewall to Control Traffic for the WF-500 Appliance VM Interface (Continued)	

Step 4 Verify that the VM interface is transmitting and receiving traffic.	1. View the VM interface settings:
	`admin@WF-500> show interface vm-interface`
	2. Verify that received/transmitted counters are incrementing. You can run the following command to generate ping traffic from the VM interface to an external device:
	`admin@WF-500> ping source vm-interface-ip host <gateway-ip>`
	For example:
	`admin@WF-500> ping source 10.16.0.20 host 10.16.0.1`

Set Up WF-500 Content Updates

Configure daily content updates for the WF-500 appliance. WF-500 content updates provide the appliance with threat intelligence to facilitate accurate malware detection, improve appliance capability to differentiate malicious samples from benign samples, and ensure that the appliance has the most recent information needed to generate signatures.

- Install WF-500 Content Updates Directly from the Update Server
- Install WF-500 Content Updates from an SCP-Enabled Server

Install WF-500 Content Updates Directly from the Update Server

Install Threat Intelligence Content Updates Directly from the Update Server

Step 1 Verify connectivity from the appliance to the update server and identify the content update to install.

1. Log in to the WF-500 appliance and run the following command to display the current content version:

   ```
   admin@wf-500> show system info | match wf-content-version
   ```

2. Confirm that the appliance can communicate with the Palo Alto Networks Update Server and view available updates:

   ```
   admin@wf-500> request wf-content upgrade check
   ```

   The command queries the Palo Alto Networks Update Server and provides information about available updates and identifies the version that is currently installed on the appliance.

   ```
   Version Size  Released on            Downloaded Installed
   ---------------------------------------------------------
   2-253   57MB  2014/09/20 20:00:08 PDT no          no
   2-39    44MB  2014/02/12 14:04:27 PST yes         current
   ```

   If the appliance cannot connect to the update server, you will need to allow connectivity from the appliance to the Palo Alto Networks Update Server, or download and install the update using SCP as described in Install WF-500 Content Updates from an SCP-Enabled Server.

Step 2 Download and install the latest content update.

1. Download the latest content update:

   ```
   admin@wf-500> request wf-content upgrade download latest
   ```

2. View the status of the download:

   ```
   admin@wf-500> show jobs all
   ```

   You can run `show jobs pending` to view pending jobs. The following output shows that the download (job id 5) has finished downloading (Status FIN):

   ```
   Enqueued            ID  Type    Status Result Completed
   ------------------------------------------------------------
   2014/04/22 03:42:20  5  Downld   FIN     OK   03:42:23
   ```

3. After the download is complete, install the update:

   ```
   admin@wf-500> request wf-content upgrade install version latest
   ```

   Run the `show jobs all` command again to monitor the status of the install.

Install Threat Intelligence Content Updates Directly from the Update Server (Continued)		

Step 3	Verify the content update.	Run the following command and refer to the `wf-content-version` field:
		`admin@wf-500> show system info`
		The following shows an example output with content update version 2-253 installed:
		`admin@wf-500> show system info`
		`hostname: wf-500`
		`ip-address: 10.5.164.245`
		`netmask: 255.255.255.0`
		`default-gateway: 10.5.164.1`
		`mac-address: 00:25:90:c3:ed:56`
		`vm-interface-ip-address: 192.168.2.2`
		`vm-interface-netmask: 255.255.255.0`
		`vm-interface-default-gateway: 192.168.2.1`
		`vm-interface-dns-server: 192.168.2.1`
		`time: Mon Apr 21 09:59:07 2014`
		`uptime: 17 days, 23:19:16`
		`family: m`
		`model: WF-500`
		`serial: abcd3333`
		`sw-version: 6.1.0`
		`wf-content-version: 2-253`
		`wfm-release-date: 2014/08/20 20:00:08`
		`logdb-version: 6.1.2`
		`platform-family: m`

Step 4	(Optional) Schedule content updates to be installed on a daily or weekly basis.	1. Schedule the appliance to download and install content updates:
		`admin@WF-500# set deviceconfig system update-schedule wf-content recurring [daily | weekly] action [download-and-install | download-only]`
		For example, to download and install updates daily at 8:00 am:
		`admin@WF-500# set deviceconfig system update-schedule wf-content recurring daily action download-and-install at 08:00`
		2. Commit the configuration
		`admin@WF-500# commit`

Install WF-500 Content Updates from an SCP-Enabled Server

The following procedure describes how to install threat intelligence content updates on a WF-500 appliance that does not have direct connectivity to the Palo Alto Networks Update Server. You will need a Secure Copy (SCP)-enabled server to temporarily store the content update.

Install Threat Intelligence Content Updates from an SCP-Enabled Server		
Step 1	Retrieve the content update file from the update server.	1. Log in to the Palo Alto Networks Support site and click Dynamic Updates. 2. In the WF-500 Appliance section, locate the latest WF-500 appliance content update and download it. 3. Copy the content update file to an SCP-enabled server and note the file name and directory path.

Install Threat Intelligence Content Updates from an SCP-Enabled Server (Continued)	

Step 2 Install the content update on the WF-500 appliance.	1. Log in to the WF-500 appliance and download the content update file from the SCP server:
	`admin@WF-500> scp import wf-content from username@host:path`
	For example:
	`admin@WF-500> scp import wf-content from bart@10.10.10.5:c:/updates/panup-all-wfmeta-2-253.tgz`
	If your SCP server is running on a non-standard port or if you need to specify the source IP, you can also define those options in the **scp import** command.
	2. Install the update:
	`admin@WF-500> request wf-content upgrade install file panup-all-wfmeta-2-253.tgz`
	View status of the install:
	`admin@WF-500> show jobs all`

Step 3 Verify the content update.	Verify the content version:
	`admin@wf-500> show system info | match wf-content-version`
	The following output now shows version 2-253:
	`wf-content-version: 2-253`

Enable Signatures and URL Generation

The WF-500 appliance can generate signatures locally based on the samples received from connected firewalls and the WildFire API, as an alternative to sending malware to the public cloud for signature generation. The appliance can generate the following types of signatures for the firewalls to use to block malware and any associated command and control traffic:

- **Antivirus signatures**—Detect and block malicious files. WildFire adds these signatures to WildFire and Antivirus content updates.
- **DNS signatures**—Detect and block callback domains for command and control traffic associated with malware. WildFire adds these signatures to WildFire and Antivirus content updates.
- **URL categories**—Categorizes callback domains as malware and updates the URL category in PAN-DB.

Configure the firewalls to receive the signatures generated by the WF-500 appliance as frequently as every five minutes. You can also send the malware sample to the WildFire public cloud, in order to enable the signature to be distributed globally through Palo Alto Networks content releases.



If you're planning to configure the WF-500 appliance as a part of a WildFire Hybrid Cloud deployment, you should also enable the firewall to receive the latest signatures distributed by the WildFire public cloud.

The following steps describe how to enable the WF-500 appliance to generate signatures and URL categories and how to distribute the signatures and categories to the firewalls in your network.

Enable the WF-500 Appliance to Generate and Distribute Signature and URL Categories	
Before you begin, make sure that you have Set Up WF-500 Content Updates for the WF-500 appliance to receive the latest threat intelligence from Palo Alto Networks.	

Step 1 Enable signature and URL category generation.	1. Log in to the appliance and type `configure` to enter configuration mode.
	2. Enable all threat prevention options:
	`admin@WF-500# set deviceconfig setting wildfire signature-generation av yes dns yes url yes`
	3. Commit the configuration:
	`admin@WF-500# commit`

Enable the WF-500 Appliance to Generate and Distribute Signature and URL Categories	
Step 2 Create a schedule for firewalls connected to the appliance to receive the signatures and URL categories generated by the appliance. It is a best practice to configure your firewalls to retrieve content updates from both the WildFire public cloud and WF-500 appliance. This ensures that your firewalls receive signatures based on threats detected world wide, in addition to the signatures generated by the local appliance.	For multiple firewalls managed by Panorama: Launch Panorama and select **Panorama > Device Deployment > Dynamic Updates**, click **Schedules**, and **Add** scheduled content updates for managed devices. For details on using Panorama to set up managed firewalls to receive signatures and URL categories from a WF-500 appliance, see Schedule Content Updates to Devices Using Panorama. For a single firewall: 1. Log in to the firewall web interface and select **Device > Dynamic Updates**. For firewalls configured to forward files to a WF-500 appliance (in either a WildFire private cloud or hybrid cloud deployment), the WF-Private section is displayed. 2. Set the **Schedule** for the firewall to download and install content updates from the WF-500 appliance.  For details on setting up and managing firewall content updates, see Manage Content Updates.

Upgrade a WF-500 Appliance

This section describes how to upgrade the WF-500 appliance operating system. First download and install the WF-500 appliance guest VM image for that release. The VM images can be as large as 4GB, so you must download them from the Palo Alto Networks update servers and host them on an SCP-enabled server that you provide. Use the SCP client on the appliance to download the software and VM images from the SCP-enabled server prior to upgrading the appliance.

The appliance can only use one environment at a time to analyze samples, so after upgrading the appliance, review the list of available VM images and then choose the image that best fits your environment. In the case of Windows 7, if your environment has a mix of Windows 7 32-bit and Windows 7 64-bit systems, it is recommended that you choose the Windows 7 64-bit image, so WildFire will analyze both 32-bit and 64-bit PE files. Although you configure the appliance to use one virtual machine image configuration, to improve the appliance uses multiple instances of the image to perform file analyses.

The VM files can be as large as 4GB, so ensure that your Secure Copy (SCP) enabled server software supports file transfers over 4GB and verify that there is enough free space to temporarily store the files.

Upgrade the WF-500 Appliance		

Step 1	Download the 7.0 software version to the WF-500 appliance.
	You cannot skip any major release versions when upgrading the WF-500 appliance. For example, if you want to upgrade from version 6.0 to 7.0, you must first install version 6.1.
		Log in to the WF-500 appliance and download the 7.0.0 software version:
		`admin@WF-500> request system software download version 7.0.0`
		To check the status of the download, use the following command:
		`admin@WF-500> show jobs all`

Step 2	Download the 7.0 VM image to an SCP-enabled server.	1. On the Palo Alto Networks Support site, click Software Updates. 2. In the WF-500 Guest VM Images section, download the 7.0.0 guest VM image—WFWinXpGf_m-1.0.0-c6.xpgf—to your local system:  3. Move the file to your SCP-enabled server and note the file name and directory path.

Upgrade the WF-500 Appliance (Continued)		

Step 3	Download the VM image to the WF-500 appliance.	Download the VM image file to the WF-500 appliance from the SCP-enabled server by running the following operation command on the appliance:
		`admin@WF-500> scp import wildfire-vm-image from <server>:/<path>/WFWinXpGf_m-1.0.0-c6.xpgf`
		For example:
		`admin@WF-500> scp import wildfire-vm-image from bart@10.43.15.41:c:/root/WF-Elink-Image/WFWinXpGf_m-1.0.0-c6xpgf`
		The SCP path following the IP address or hostname varies depending on the SCP software that you are using. For Windows, the path is `c:/folder/filename` or `//folder/filename`; for Unix/Mac systems, the path is `/folder/filename` or `//folder/filename`.

Step 4	Install the VM image on the WF-500 appliance.	Run the following command to install the VM image to the WF-500 appliance: `admin@WF-500> request system wildfire-vm-image upgrade install file WFWinXpGf__m-1.0.0_c6.xpgf`
Step 5	Install the 7.0 software version on the WF-500 appliance.	Install the WF-500 appliance operating system image (downloaded in Step 1) by running the following command: `admin@WF-500> request system software install version 7.0.0`

Step 6	Restart the WF-500 appliance and confirm that the installation was successful.	1. Confirm that the upgrade is complete. Run the following command and look for the job type `Install` and status `FIN`:
		`admin@WF-500> show jobs all`
		`Enqueued ID Type Status Result Completed`
		`2015/05/15 10:38:48 2 Downld FIN OK 10:39:08`
		2. Restart the appliance:
		`admin@WF-500> request restart system`
		3. Verify that the `sw-version` field shows 7.0:
		`admin@WF-500> show system info | match sw-version`

Upgrade the WF-500 Appliance (Continued)		

Step 7	(Optional) Enable the Windows 7 64-bit sandbox environment.	1. View the active virtual machine image by running the following command and refer to the `Selected VM field`:
		`admin@WF-500> show wildfire status`
		2. View a list of available virtual machines images:
		`admin@WF-500> show wildfire vm-images`
		The following output shows that `vm-5` is the Windows 7 64-bit image:
		`vm-5`
		`Windows 7 64bit, Adobe Reader 11, Flash 11, Office 2010. Support PE, PDF, Office 2010 and earlier`
		3. Select the image to be used for analysis:
		`admin@WF-500# set deviceconfig setting wildfire active-vm <vm-image-number>`
		For example, to use vm-5, run the following command:
		`admin@WF-500# set deviceconfig setting wildfire active-vm vm-5`
		4. `Commit` the configuration:
		`admin@WF-500# commit`

Step 8	Upgrade the firewalls connected to the appliance.	Upgrade firewalls to PAN-OS 7.0.

Upgrade the WF-500 Appliance (Continued)	
Step 9 Enable firewalls to continue to forward files to the WF-500 appliance for analysis. This step is required to ensure WildFire settings migrated during the upgrade are configured correctly. During the upgrade to PAN-OS 7.0, file forwarding settings are migrated from the File Blocking profile to the WildFire Analysis profile. The new WildFire Analysis profile enables the firewall to support a WildFire hybrid cloud deployment, where files can be forwarded to both the public cloud and private cloud for analysis.	1. Define the traffic to be forwarded for WildFire analysis: a. Select **Objects > Security Profiles > WildFire Analysis** and for each profile rule, update the Analysis setting to **private-cloud**. b. (Optional) To use the WildFire public cloud to analyze certain files (for example, APK files which are not supported for WF-500 appliance analysis), add or modify a WildFire Analysis profile rule to use the **Analysis** location **public-cloud**. Traffic matched to rules with **public-cloud** analysis location configured will be forwarded to the WildFire public cloud for analysis. 2. Select **Device > Setup > WildFire** and configure the WildFire locations to which the firewall will forward files for analysis: a. In the **WildFire Private Cloud** field, enter the IP address or FQDN of a WF-500 appliance. b. If the firewall was connected to a WF-500 appliance before the upgrade, the **WildFire Public Cloud** field is automatically populated with the IP address or FQDN of the WF-500 appliance. Populate this field depending on whether you want to analyze files using only the WF-500 appliance, or both the WF-500 appliance and the WildFire public cloud: – To continue forwarding files only to the WF-500 appliance, delete the WildFire Public Cloud entry, leaving the field blank. – To enable the firewall to forward files to both the WF-500 appliance and the WildFire Public Cloud hosted in the United States, enter wildfire.paloaltonetworks.com. To use the Japan cloud, enter wildfire.paloaltonetworks.jp.
Next Step...	Forward Files for WildFire Analysis.



Submit Files for WildFire Analysis

The following topics describe how to submit files for WildFire™ analysis. You can set up Palo Alto Networks firewalls to automatically forward unknown files to the WildFire public cloud or a WildFire private cloud, and you can also manually submit files for analysis using the WildFire portal. Samples submitted for WildFIre analysis receive a verdict of benign, grayware, or malware, and a detailed analysis report is generated for each sample.

- Forward Files for WildFire Analysis
- Verify WildFire Submissions
- Manually Upload Files to the WildFire Portal
- Submit Malware or Reports from the WF-500 Appliance
- Firewall File Forwarding Capacity by Platform

Forward Files for WildFire Analysis

Configure Palo Alto Networks firewalls to forward unknown files or email links for analysis. Use the **WildFire Analysis** profile to define files to forward to the WildFire cloud (use the public cloud or a private cloud), and then attach the profile to a security rule to trigger inspection for zero-day malware.

Specify traffic to be forwarded for analysis based on the application in use, the file type detected, links contained in email messages, or the transmission direction of the sample (upload, download, or both). For example, you can set up the firewall to forward Portable Executables (PEs) or any files that users attempt to download during a web-browsing session.

If you are using a WF-500 appliance to host a WildFire private cloud, you can extend WildFire analysis resources to a WildFire Hybrid Cloud, by configuring the firewall to continue to forward sensitive files to your WildFire private cloud for local analysis, and forward less sensitive or unsupported file types to the WildFire public cloud.

Configure a Firewall to Forward Files and Email Links to WildFire

Before you begin:

- ❑ If another firewall resides between the firewall you are configuring to forward files and the WildFire cloud or WF-500 appliance, make sure that the firewall in the middle allows the following ports:
  - The WildFire public cloud uses port 443 for registration and file submissions.
  - The WF-500 appliance uses port 443 for registration and 10443 for file submissions.
- ❑ Verify that the firewall has valid Threat Prevention and WildFire subscriptions (**Device > Licenses**).
- ❑ Verify that content updates are scheduled and up-to-date. Select **Device > Dynamic Updates** and **Check Now** to ensure that the firewall has the most recent Antivirus, Applications and Threats, and WildFire updates.
- ❑ (PA-7000 Series Firewalls Only) To enable a PA-7000 Series firewall to forward files and email links for WildFire analysis, you must first configure a data port on an NPC as a Log Card interface.

Configure a Firewall to Forward Files and Email Links to WildFire

Step 1 Configure WildFire settings.

WildFire settings including defining the WildFire public and private cloud analysis locations and the option to enable reporting for grayware or benign files.



Panorama Only:

If Panorama detects a **WildFire Submissions** log entry with incomplete fields, Panorama can connect to WildFire to gather information about the sample and populate the log entry details. Select **Panorama > Setup > WildFire** and enter a **WildFire Server** for Panorama to communicate with to gather sample details (by default, Panorama will use the WildFire public cloud).

1. Select **Device > Setup > WildFire** and edit the General Settings edit icon.
2. Enter the WildFire public cloud and WildFire private cloud that you want to use for WildFire analysis:
   - To forward files to the **WildFire Public Cloud** hosted in the United States, enter `wildfire.paloaltonetworks.com`.

     To forward files to the WildFire cloud hosted in Japan, enter `wildfire.paloaltonetworks.jp`. If you are in the Japan region, you might experience faster response times for sample submissions and report generation when using the Japan cloud. You can also use the Japan cloud if you do not want grayware or benign files forwarded to the United States cloud servers; however, if a file sent to the Japan cloud is determined to be malicious, it will continue to be forwarded to the United States servers for analysis and signature generation.
   - To forward files to a **WildFire Private Cloud**, enter the IP address or FQDN of the WF-500 appliance.
   - Leave either field empty if you do not plan to use that cloud for file analysis.
3. (Optional) Modify the **File Size Limits** for files forwarded from the firewall. For example, if you set PDF to 5MB, any PDF larger than 5MB will not be forwarded.
4. (Optional) Enable reporting for benign files and grayware:
   - Select **Report Benign Files** to allow logging for files that receive a WildFire verdict of benign.
   - Select **Report Grayware Files** to allow logging for files that receive a WildFire verdict of grayware.

   To view logs for files that receive grayware and benign verdicts, select **Monitor > WildFire Submissions**.
5. (Optional) Define what session information is recorded in WildFire analysis reports:
   a. Edit the Session Information Settings.
   b. By default, all session information is displayed in WildFire analysis reports. Clear the check boxes to remove the corresponding fields from WildFire analysis reports and click **OK** to save the settings.

Configure a Firewall to Forward Files and Email Links to WildFire

Step 2 Define traffic to be forwarded for WildFire analysis.

If you have a WF-500 appliance set up, you can use both the private cloud and the public cloud in a hybrid cloud deployment. Analyze sensitive files locally on your network, while sending all other unknown files to the WildFire public cloud for comprehensive analysis and prompt verdict returns.

1. Select **Objects > Security Profiles > WildFire Analysis**, **Add** a new WildFire analysis profile, and give the profile a descriptive **Name**.
2. **Add** a profile rule to define traffic to be forwarded for analysis and give the rule a descriptive **Name**, such as local-PDF-analysis.
3. Define for the profile rule to match to unknown traffic and to forward samples for analysis based on:
   - **Applications**—Forward files for analysis based on the application in use.
   - **File Types**—Forward files for analysis based on file types, including links contained in email messages. For example, select **PDF** to forward unknown PDFs detected by the firewall for analysis.
   - **Direction**—Forward files for analysis based the transmission direction of the file (upload, download, or both). For example, select **both** to forward all unknown PDFs for analysis, regardless of the transmission direction.
4. Set the **Analysis** location to which files matched to the rule will be forwarded.
   - Select **public-cloud** to forward files matched to the rule to the WildFire public cloud for analysis.
   - Select **private-cloud** to forward files matched to the rule to the WildFire private cloud for analysis.

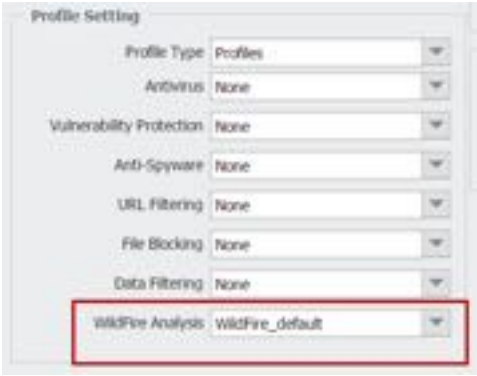
   For example, to analyze PDFs that could contain sensitive or proprietary information without sending these documents out of your network, set the **Analysis** location for the rule local-PDF-analysis to **private-cloud**.



In a hybrid cloud deployment, files that match to both **private-cloud** and **public-cloud** rules are forwarded only to the private cloud as a cautionary measure.

5. (Optional) Continue to add rules to the WildFire analysis profile as needed. For example, you could add a second rule to the profile to forward Android application package (APK), Portable Executable (PE), and Flash files to the WildFire public cloud for analysis.
6. Click **OK** to save the WildFire analysis profile.



Configure a Firewall to Forward Files and Email Links to WildFire	
Step 3 Attach the WildFire analysis profile to a security policy rule. Traffic allowed by the security policy rule is evaluated against the attached WildFire analysis profile; the firewalls forwards traffic matched to the profile for WildFire analysis.	1. Select **Policies > Security** and **Add** or modify a policy rule. 2. Click the **Actions** tab within the policy rule. 3. In the Profile Settings section, select **Profiles** as the **Profile Type** and select a **WildFire Analysis** profile to attach to the policy rule 
Step 4 (Optional) Enable the firewall to forward decrypted traffic for WildFire analysis. Traffic that is decrypted by the firewall is evaluated against security policy and, if matched to the WildFire analysis profile attached to a policy rule, it can be forwarded to and analyzed by WildFire before it is re-encrypted. Only a superuser can enable this option.	To forward decrypted traffic for WildFire analysis, the firewall must first be enabled to perform decryption. On a single firewall: 4. Select **Device > Setup > Content-ID**. 5. Edit the URL Filtering options and enable **Allow Forwarding of Decrypted Content**. 6. Click **OK** to save the changes. On a firewall with virtual systems configured: Select **Device > Virtual Systems**, click the virtual system you want to modify, and select the **Allow Forwarding of Decrypted Content** check box.
Step 5 Commit the configuration.	Click **Commit** to apply the settings.
Next Steps...	• Verify WildFire Submissions to confirm that the firewall is successfully forwarding files for WildFire analysis. • (WF-500 Appliance Only) Submit Malware or Reports from the WF-500 Appliance. Enable this feature to automatically forward malware identified in your WildFire private cloud to the WildFire public cloud. The WildFire public cloud re-analyzes the sample and generates a signature if the sample is malware. The signature is distributed to global users through Wildfire signature updates. • Monitor WildFire Activity to assess alerts and details reported for malware.

Verify WildFire Submissions

Test your WildFire setup using malware test samples, and also verify that the firewall is correctly forwarding files for WildFire analysis.

▲ Test a Sample Malware File

▲ Verify File Forwarding

Test a Sample Malware File

Palo Alto Networks provides a sample malware file that you can use to test a WildFire configuration. Take the following steps to download the malware sample file, verify that the file is forwarded for WildFire analysis, and view the analysis results.

Use a Sample Malware File to Test the WildFire Configuration	
Step 1	Download the malware test file: https://wildfire.paloaltonetworks.com/publicapi/test/pe. The test file is named wildfire-test-pe-file.exe and each test file has a unique SHA-256 hash value. You can also Use the API to Retrieve a Sample Malware Test File.
Step 2	On the firewall web interface, select **Monitor > WildFire Submissions** to confirm that the file was forwarded for analysis. It might take about five minutes for analysis results to be displayed for the file on the **WildFire Submissions** page. The verdict for the test file will always display as malware.

Verify File Forwarding

After the firewall is set up to Forward Files for WildFire Analysis, use the following options to verify the connection between the firewall and the WildFire public or private cloud, and to monitor file forwarding.

Several of the options to verify that a firewall is forwarding samples for WildFire analysis are CLI commands; for details on getting started with and using the CLI, refer to the PAN-OS CLI Quick Start Guide.

Verify File Forwarding	

❑ Verify that the firewall is communicating with a WildFire server(s).	Use the `test wildfire registration` command to verify that the firewall is connected to a WildFire private cloud, the WildFire public cloud, or both.
	The following example output is for a firewall in a WildFire Private Cloud deployment:
	`admin@VM-FW> test wildfire registration`
	`This test may take a few minutes to finish.`
	`Do you want to continue? (y or n)`
	`Test wildfire Public Cloud`
	`Testing cloud server wildfire.paloaltonetworks.com ...`
	`wildfire registration: failed`
	`Test wildfire Private Cloud`
	`Testing cloud server X.X.X.X ...`
	`wildfire registration: successful`
	`download server list: successful`
	`select the best server: X.X.X.X:XXXXX`
	The example output confirms that the firewall is connected to the WildFire private cloud, and is not connected to the WildFire public cloud (public cloud registration fails).
	If the firewall is configured in a WildFire Hybrid Cloud deployment, check that the firewall is successfully registered with and connected to both the WildFire public cloud and a WildFire private cloud.

Verify File Forwarding	

❑ Verify the status of the firewall connection to the WildFire public and/or private cloud, including the total number of files forwarded by the firewall for analysis.	Use the `show wildfire status` command to:
	• Check the status of the WildFire public and/or private cloud to which the firewall is connected. The status `Idle` indicates that the WildFire cloud (public or private) is ready to receive files for analysis.
	• Confirm the configured size limits for files forwarded by the firewall (**Device > Setup > WildFire**).
	• Monitor file forwarding, including how the total count of files forwarded by the firewall for WildFire analysis. If the firewall is in a WildFire hybrid cloud deployment, the number of files forwarded to the WildFire public cloud and the WildFire private cloud are also displayed.
	The following example shows the `show wildfire status` output for a firewall in a WildFire private cloud deployment:
	`admin@VM-FW> show wildfire status`
	
	`Connection info:`
	`  Signature verification:  enable`
	`  Server selection:  enable`
	`  File cache:  enable`
	`WildFire Public Cloud:`
	`  Server address:  wildfire.paloaltonetworks.com`
	`  Status:  Disabled due to configuration`
	`  Best server:`
	`  Device registered:  no`
	`  Through a proxy:  no`
	`  Valid wildfire license:  yes`
	`  Service route IP address:  X.X.X.X`
	`WildFire Private Cloud:`
	`  Server address:  X.X.X.X`
	`  Status:  Idle`
	`  Best server:  X.X.X.X:XXXXX`
	`  Device registered:  yes`
	`  Through a proxy:  no`
	`  Valid wildfire license:  yes`
	`  Service route IP address:  X.X.X.X`
	`File size limit info:`
	`  pe  9 MB`
	`  apk  49 MB`
	`  pdf  1000 KB`
	`  ms-office  9500 KB`
	`  jar  9 MB`
	`  flash  10 MB`
	`  MacOSX  1 MB`
	`Forwarding info:`
	`  file idle time out (second):  90`
	`  total concurrent files:  0`
	`  Public Cloud:`
	`  total file forwarded:  0`
	`  file forwarded in last minute:  0`
	`  concurrent files:  0`
	`  Private Cloud:`
	`  total file forwarded:  0`
	`  file forwarded in last minute:  0`
	`  concurrent files:  0`
	To view forwarding information for only the WildFire public cloud or WildFire private cloud, use the following commands:
	• `show wildfire status channel public`
	• `show wildfire status channel private`

Verify File Forwarding	
❑ View samples forwarded by the firewall according to file type (including email links). Use this option to confirm that email links are being forwarded for WildFire analysis, since only email links that receive a malware verdict are logged as **WildFire Submissions** entries on the firewall, even if logging for benign and grayware samples is enabled. This is due to the sheer number of WildFire Submissions entries that would be logged for benign email links.	Use the `show wildfire statistics` command to confirm the file types being forwarded to the WildFire public or private cloud: • The command displays the output of a working firewall and shows counters for each file type that the firewall forwards for WildFire analysis. If a counter field shows 0, the firewall is not forwarding that file type. • Confirm that email links are being forwarded for analysis by checking that the following counters do not show zero: – `FWD_CNT_APPENDED_BATCH`—Indicates the number of email links added to a batch waiting for upload to WildFire. – `FWD_CNT_LOCAL_FILE`— Indicates the total number of email links uploaded to WildFire.
❑ Verify that a specific sample was forwarded by the firewall and check that status of that sample. This option can be helpful when troubleshooting to: • Confirm that samples that have not yet received a WildFire verdict were correctly forwarded by the firewall. Because **WildFire Submissions** are logged on the firewall only when WildFire analysis is complete and the sample has received a WildFire verdict, use this option to verify the firewall forwarded a sample that is currently undergoing WildFire analysis. • Track the status for a single file or email link that was allowed according to your security policy, matched to a WildFire Analysis profile, and then forwarded for WildFire analysis. • Check that a firewall in a WildFire Hybrid Cloud deployment is forwarding the correct file types and email links to either the WildFire public cloud or a WildFire private cloud.	Execute the following CLI commands on the firewall to view samples the firewall has forwarded WildFire analysis: • View all samples forwarded by the firewall by with the CLI command `debug wildfire upload-log`. • View only samples forwarded to the WildFire public cloud with the CLI command `debug wildfire upload-log channel public`. • View only samples forwarded to the WildFire private cloud with the CLI command `debug wildfire upload-log channel private`. The following example shows the output for the three commands listed above when issued on a firewall in a WildFire public cloud deployment: 

```
user@firewall> debug wildfire upload-log
+ channel    WildFire channel (Public/Private)
  |          Pipe through a command
  <Enter>    Finish input

user@firewall> debug wildfire upload-log channel private

Private Cloud upload logs:


user@firewall> debug wildfire upload-log channel public

Public Cloud upload logs:

    log: 0, filename: support-login.swf
    processed 353590 seconds ago, action: skipped - remote benign dup
    vsys_id:  1, session_id: 169651, transaction_id: 261
    file_len:  91536, flag: 0x81c, file type: flash
    threat id: 52145, user_id: 1238, app_id: 872
    from XX.XX.XX.XX/XXXX to XX.XXX.XXX.XXX/XXX
    SHA256: 6b2f1a23407ab2db9a17ccdf66beace6dad7d2489c65ba90dbdf02305b3d4efd

    log: 1, filename: GEM_D_because_12.03.2014_300x250.swf
    processed 611505 seconds ago, action: skipped - remote benign dup
    vsys_id:  1, session_id: 259049, transaction_id: 260
    file_len:  39206, flag: 0x81c, file type: flash
    threat id: 52145, user_id: 20583, app_id: 872
    from XX.XX.XX.XX/XXXXX to XXX.XX.XXX.XXX/XX
    SHA256: ca52d1b7a7521a14237c1531edb109627fee084804a300d907b57322b1efd4e7
```

Verify File Forwarding	
❑ Monitor samples successfully submitted for WildFire analysis.	Using the firewall web interface, select **Monitor > Logs > WildFire Submissions**. All files forwarded by a firewall to the WildFire public or private cloud for analysis are logged on the WildFire Submissions page. • Check the WildFire verdict for a sample: By default, only samples that receive malware verdicts are displayed as **WildFire Submissions** entries. To enable logging for benign and/or grayware samples, select **Device > Setup > WildFire > Report Benign Files/ Report Grayware Files**.  Enable logging for benign files as a quick troubleshooting step to verify that the firewall is forwarding files. Check the **WildFire Submissions** logs to verify that files are being submitted for analysis and receiving WildFire verdicts (in this case, a benign verdict). • Confirm the analysis location for a sample: The **WildFire Cloud** column displays the location to which the file was forwarded and where it was analyzed (public cloud or private cloud). This is useful when deploying a WildFire Hybrid Cloud.

Manually Upload Files to the WildFire Portal

All Palo Alto Networks customers with a support account can use the Palo Alto Networks WildFire portal to manually submit files for WildFire analysis.

Upload Samples the WildFire Portal		
Step 1	Manually upload files or URLs from your network to the WildFire portal for analysis.	1. Log in to the WildFire Portal. If your firewall is forwarding to the WildFire portal in Japan, use https://wildfire.paloaltonetworks.jp. 2. Click **Upload Sample** on the menu bar then click **Add files**. 3. **Open** the file for which you want to receive a WildFire verdict and analysis report. The file name will appear below the **Add files** icon. 4. Click the **Start** icon to the right of the file, or click the **Start upload** button if multiple files are waiting for upload. If the file(s) upload successfully, Success will appear next to each file.  5. Close the **Uploaded File Information** pop-up.
Step 2	View the analysis results for the file. WildFire takes approximately five minutes to complete a file analysis. Because a manual upload is not associated with a specific firewall, manual uploads will appear separately from your registered firewalls and will not show session information in the reports.	1. Refresh the portal page from your browser. 2. Click **Manual** under the source column to view the results of manual sample upload. 3. The report page will show a list of all files that have been uploaded to your account. Find the file you uploaded and click the detail icon to the left of the date field. The portal displays a full report of the file analysis detailing the observed file behavior. If WildFire identifies the file as malware, it generates a signature, which is then distributed to all Palo Alto Networks firewalls configured with a WildFire or Threat Prevention subscription.

Submit Malware or Reports from the WF-500 Appliance

Enable the WF-500 appliance cloud intelligence feature to automatically submit malware samples discovered in the WildFire private cloud to the WildFire public cloud. The WildFire public cloud further analyzes the malware and generates a signature to identify the sample. The signature is then added to WildFire signature updates, and distributed to global users to prevent future exposure to the threat. If you do not want to forward malware samples outside of your network, you can instead choose to submit only WildFire reports for the malware discovered on your network to contribute to WildFire statistics and threat intelligence.

Enable the WF-500 Appliance to Submit Malware or Reports to the WildFire Public Cloud

Submit Malware to the WildFire Public Cloud

Step 1 Execute the following CLI command from the WF-500 appliance to enable the appliance to automatically submit malware samples to the WildFire public cloud:

```
admin@WF-500# set deviceconfig setting wildfire cloud-intelligence submit-sample yes
```



If the firewall that originally submitted the sample for WildFire private cloud analysis has packet captures (PCAPs) enabled, the PCAPs for the malware will also be forwarded to the WildFire public cloud.

Step 2 Go to the WildFire portal to view analysis reports for malware automatically submitted to the WildFire public cloud.

Submit Analysis Reports to the WildFire Public Cloud

If the WF-500 appliance is enabled to Submit Malware to the WildFire Public Cloud, you do not need to also enable the appliance to submit reports to the public cloud. When malware is submitted to the WildFire public cloud, the public cloud generates a new analysis report for the sample.

If you want the WF-500 appliance to automatically submit malware reports to the WildFire public cloud (and not the malware sample), execute the following CLI command on the WF-500 appliance:

```
admin@WF-500# set deviceconfig setting wildfire cloud-intelligence submit-report yes
```

Verify Cloud Intelligence Settings

Check to confirm that cloud intelligence is enabled to either submit malware or submit reports to the WildFire public cloud by running the following command:

```
admin@WF-500> show wildfire status
```

Refer to the `Submit sample` and `Submit report` fields.

Firewall File Forwarding Capacity by Platform

File forwarding capacity is the maximum rate per minute at which each Palo Alto Networks firewall platform can submit files to the WildFire cloud or a WF-500 appliance for analysis. If the firewall reaches the per-minute limit, it queues any remaining samples.

The Reserved Drive Space column in the following table lists the amount of drive space on the firewall that is reserved for queuing files. If the firewall reaches the drive space limit, it cancels forwarding of new files to WildFire until more space in the queue is available.



The speed at which the firewall can forward files to WildFire also depends on the bandwidth of the upload link to the WildFire systems.

Platform	Maximum Files Per Minute	Reserved Drive Space
VM-100	5	100MB
VM-200	10	200MB
VM-300	20	200MB
PA-200	5	100MB
PA-500	10	200MB
PA-2000 Series	20	200MB
PA-3020	50	200MB
PA-3050/3060	50	500MB
PA-4020	20	200MB
PA-4050/4060	50	500MB
PA-5000 Series	50	500MB
PA-7000 Series	100	1GB

Monitor WildFire Activity

Depending on your WildFire™ deployment—public, private, or hybrid cloud—you can view samples submitted to WildFire and analysis results for each sample using the WildFire portal, by accessing the firewall that submitted the sample (or Panorama, if you are centrally managing multiple firewalls), or by using the WildFire API. If the WF-500 appliance is enabled to Submit Malware to the WildFire Public Cloud, analysis results for the malware submitted to the public cloud can also be viewed on the WildFire portal.

After WildFire has analyzed a sample and delivered a verdict of malware, grayware, or benign, a detailed analysis report is generated for the sample. WildFire analysis reports viewed on the firewall that submitted the sample also include details for the session during which the sample was detected. For samples identified as malware, the WildFire analysis report includes details on existing WildFire signatures that might be related to the newly-identified malware and information on file attributes, behavior, and activity that indicated that the sample was malicious.

See the following topics to use the WildFire portal or a connected firewall to monitor WildFire submissions, to view reports for analyzed samples, and to set up alerts and notifications based on submissions and analysis results:

- Use the Firewall to Monitor WildFire Activity
- Use the WildFire Portal to Monitor WildFire Activity
- WildFire Analysis Reports—Close Up

Use the Firewall to Monitor WildFire Activity

Samples forwarded by the firewall are added as entries to the **WildFire Submissions** logs. A detailed WildFire analysis report is displayed in the expanded view for each WildFire Submissions entry.

- Configure WildFire Submissions Log Settings
- Monitor WildFire Submissions and Analysis Reports
- Set Up Alerts for Malware Logs

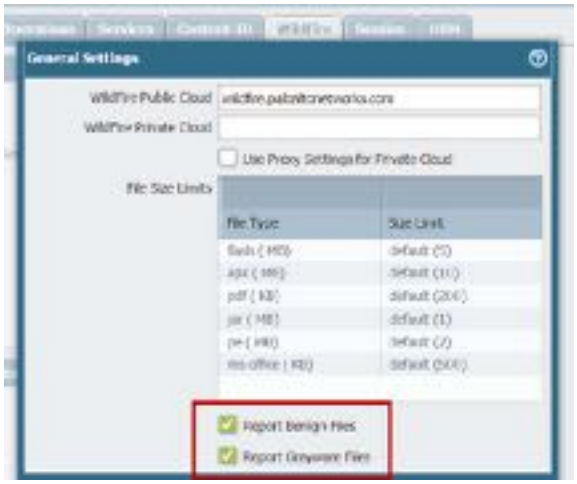
Configure WildFire Submissions Log Settings

Enable the following options for **WildFire Submissions** logs:

- Enable Logging for Benign and Grayware Samples
- Include Email Header Information in WildFire Logs and Reports
- Include User-ID Information in WildFire Logs and Reports

Enable Logging for Benign and Grayware Samples

Logging for benign and grayware samples is disabled by default. Email links that receive benign or grayware verdicts are not logged.

Enable Logging for Benign and Grayware Samples	
Step 1	Select **Device > Setup > WildFire**, edit **General Settings**.
Step 2	Select **Report Benign Files** and/or **Report Grayware Files** and click **OK** to save the settings. 

Include Email Header Information in WildFire Logs and Reports

Use the following steps to include email header information—email sender, recipient(s), and subject—in WildFire logs and reports.

Session information is forwarded to the WildFire cloud along with the sample, and used to generate the WildFire analysis report. Neither the firewall nor the WildFire cloud receive, store, or view actual email contents.



Session information can help you to quickly track down and remediate threats detected in email attachments or links, including how to identify recipients who have downloaded or accessed malicious content.

Include Email Header Information in WildFire Logs and Reports	
Step 1	Select **Device > Setup > WildFire**.
Step 2	Edit the Session Information Settings section and enable one or more of the options (**Email sender**, **Email recipient**, and **Email subject**).
Step 3	Click **OK** to save.

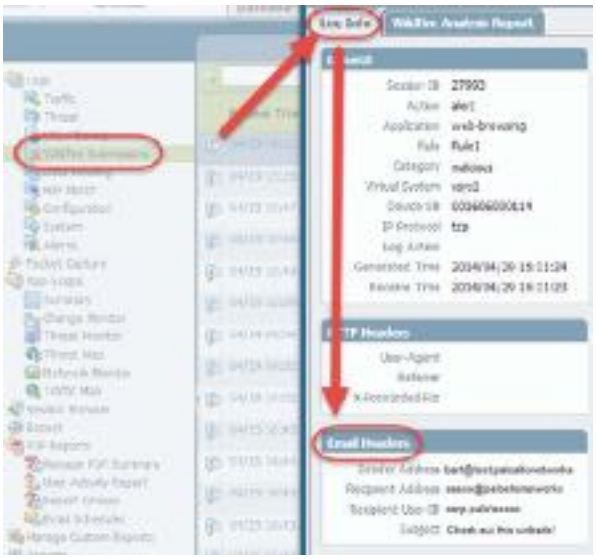
Include User-ID Information in WildFire Logs and Reports

Enable the firewall to match User-ID information with email header information, so that the User-ID for the recipient of a malicious email attachment or link is identified for a WildFire entry.

Include User-ID Information in WildFire Logs and Reports	
Step 1	Select **Device > User Identification > Group Mapping Settings**.
Step 2	Select the desired group mapping profile to modify it.
Step 3	In the **Server Profile** tab in the Mail Domains section, populate the **Domain List** field: • **Mail Attributes**—This field is automatically populated after you fill in the **Domain List** field and click **OK**. The attributes are based on your LDAP server type (Sun/RFC, Active Directory, and Novell). • **Domain List**—Enter the list of email domains in your organization using a comma separated list up to 256 characters.

Include User-ID Information in WildFire Logs and Reports

When email header information is matched to a User-ID, the **Recipient User-ID** field in the **Email Headers** section of the detailed log view will link to a filtered ACC for that user or user group.



Monitor WildFire Submissions and Analysis Reports

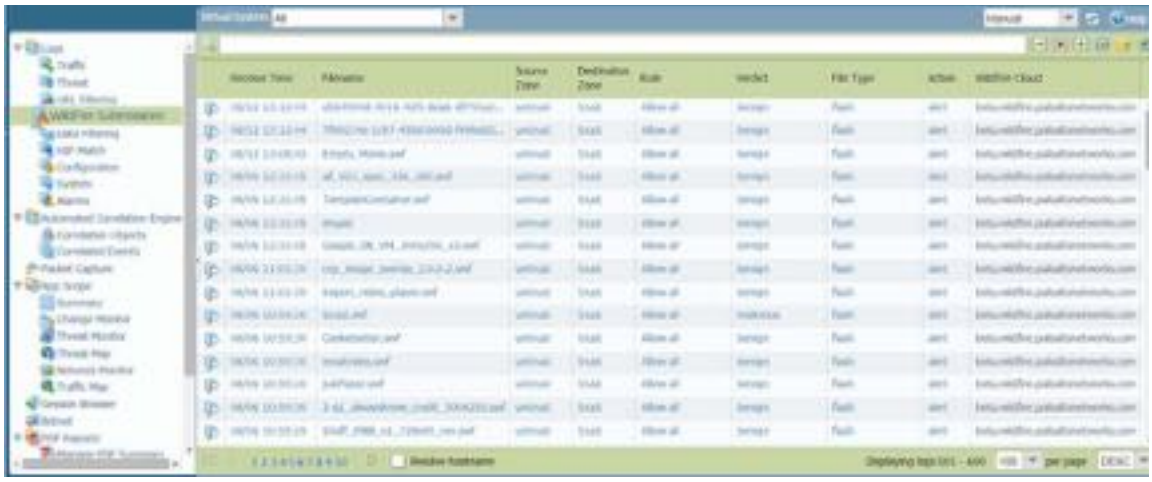
Samples that firewalls submit for WildFire analysis are displayed as entries in the **WildFire Submissions** log on the firewall web interface. For each WildFire entry, you can open an expanded log view which displays log details and the WildFire analysis report for the sample.

Monitor WildFire Submissions and Reports

Step 1	Forward Files for WildFire Analysis.
Step 2	Configure WildFire Submissions Log Settings.

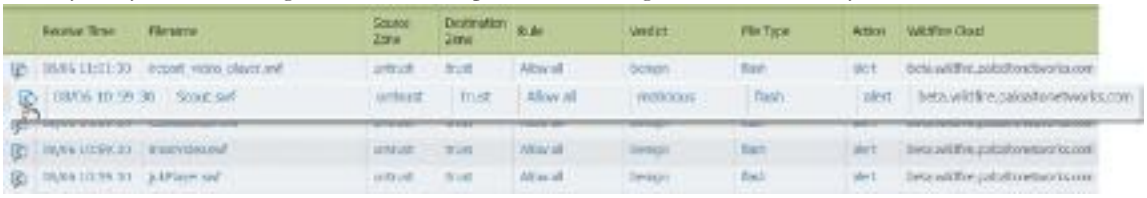
Monitor WildFire Submissions and Reports (Continued)

Step 3 To view samples submitted by a firewall to a WildFire public, private, or hybrid cloud, select **Monitor > WildFire Submissions**. When WildFire analysis of a sample is complete, the results are sent back to the firewall that submitted the sample and are accessible in the WildFire Submissions logs. The Verdict column indicates whether the sample is benign, malicious, or grayware.



Monitor WildFire Submissions and Reports (Continued)

Step 4 For any entry, select the Log Details icon to open a detailed log view for each entry:



The detailed log view displays Log Info and the WildFire Analysis Report for the entry. If the firewall has packet captures (PCAPs) enabled, the sample PCAPs are also displayed.



For all samples, the WildFire analysis report displays file and session details. For malware samples, the WildFire analysis report is extended to include details on the file attributes and behavior that indicated the file was malicious.

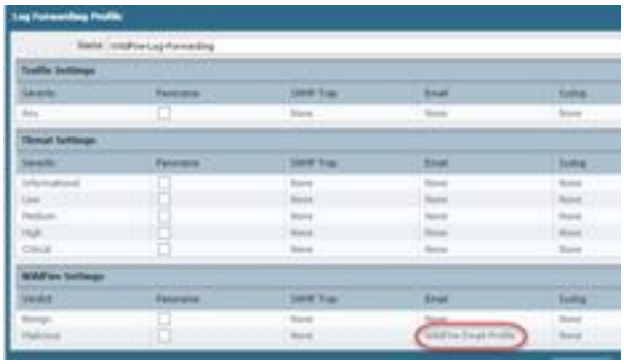


Step 5 (Optional) **Download PDF** of the WildFire Analysis Report.

Set Up Alerts for Malware Logs

You can configure a Palo Alto Networks firewall to send an alert when WildFire identifies a malicious file or email link. You can configure alerts for benign and grayware files as well, but not for benign and grayware email links. This example describes how to configure an email alert; however, you could also configure log forwarding to set up alerts to be delivered as syslog messages, SNMP traps, or Panorama alerts.

Set Up Email Alerts for Malware	
Step 1 Configure an email server profile.	1. Select **Device > Server Profiles > Email**. 2. Click **Add** and then enter a **Name** for the profile. For example, WildFire-Email-Profile. 3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down. 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 (up to four email servers can be added to the profile): • **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. • **Display Name**—The name to show in the From field of the email. • **From**—The email address where notification emails are sent from. • **To**—The email address to which notification emails are sent. • **Additional Recipient(s)**—Enter an email address to send notifications to a second recipient. • **Gateway**—The IP address or host name of the SMTP gateway to use to send the emails. 5. Click **OK** to save the server profile. 6. Click **Commit** to save the changes to the running configuration.
Step 2 Test the email server profile.	1. Select **Monitor > PDF Reports > Email Scheduler**. 2. Click **Add** and select the new email profile from the **Email Profile** drop-down. 3. Click the **Send test** email button and a test email should be sent to the recipients defined in the email profile.

Set Up Email Alerts for Malware (Continued)		
Step 3	Configure a log forwarding profile, to enable WildFire logs to be forwarded to Panorama, an email account, SNMP, and/or a syslog server. In this example you will set up email logs for when a sample is determined to be malicious. You can also enable Benign and Grayware logs to be forwarded, which will produce more activity if you are testing.	1. Select **Objects > Log Forwarding**. 2. Click **Add** and name the profile. For example, **WildFire-Log-Forwarding**. 3. In the **WildFire Settings** section, choose the email profile from the **Email** column for **Malicious** as shown in the screen capture.  To forward logs to Panorama, select the check boxes under the Panorama column for Benign, Grayware, and/or Malicious. For SNMP and Syslog, select the drop-down and choose the appropriate profile or click **New** to configure a new profile. 4. Click **OK** to save the changes.
Step 4	Add the log forwarding profile to a security policy being used for WildFire forwarding (with a WildFire Analysis profile attached). The WildFire Analysis profile defines the traffic that the firewall forwards for WildFire analysis. To set up a WildFire analysis profile and attach it to a security policy rule, see Forward Files for WildFire Analysis.	1. Select **Policies > Security** and click on the policy that is used for WildFire forwarding. 2. In the **Actions** tab **Log Setting** section, select the **Log Forwarding** profile you configured. 3. Click **OK** to save the changes and then **Commit** the configuration.

Use the WildFire Portal to Monitor WildFire Activity

Log in to the Palo Alto Networks WildFire portal 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 subscription or support account. For each device listed, the portal displays statistics for the number of malware samples that have been detected, benign samples that have been analyzed, and the number of pending files that are waiting to be analyzed. Your WildFire portal account displays data for all samples submitted by firewalls on your network that are connected to the WildFire public cloud, as well as data for samples manually submitted to the portal. Additionally, if you have enabled a WF-500 appliance to forward malware to the WildFire public cloud for signature generation and distribution, reports for those malware samples can also be accessed on the portal.

See the following sections for details on using the WildFire portal to monitor WildFire activity:

- Configure WildFire Portal Settings
- Add WildFire Portal Users
- View Reports on the WildFire Portal

Configure WildFire Portal Settings

This section describes the settings that can be customized for a WildFire cloud account, such as time zone and email notifications for each firewall connected to the account. You can also delete firewall logs stored in the cloud.

Customize the WildFire Portal Settings	
Log in to the WildFire cloud and select **Settings** on the menu bar to modify the following time zone, log, and notification settings.	
• Configure the time zone for the WildFire cloud account.	Select a time zone from the **Set Time Zone** drop-down and **Update Time Zone** to save the change.  The time stamp that appears on WildFire analysis reports is based on the time zone configured for the WildFire cloud account.
• Delete WildFire logs hosted on the cloud for specific firewalls.	1. In the **Delete WildFire Reports** drop-down, select a firewall (by serial number) and **Delete Reports** to remove logs for that firewall from WildFire portal. This action does not delete logs stored on the firewall. 2. Click **OK** to proceed with the deletion.

Customize the WildFire Portal Settings	
• Configure email notifications based on WildFire analysis verdicts.	1. In the Configure Alerts section, select **Malware**, **Grayware**, and/or **Benign** check boxes to receive email notifications based on those verdicts: • Select the verdict check boxes in the **All** row to receive verdict notifications for all samples uploaded to the WildFire cloud. • Select the verdict check boxes in the **Manual** row to receive verdict notifications for all samples that are manually uploaded to the WildFire public cloud using the WildFire portal. • Select the verdict check boxes for one or several firewall serial numbers to receive verdict notifications for samples submitted by those firewalls. 2. Select **Update Notification** to enable verdict notifications to be emailed to the email address associated with your support account.

Add WildFire Portal Users

WildFire portal accounts are created by a super user (the registered owner of a Palo Alto Networks device) to give additional users the ability to log in to the WildFire cloud and view device data for which they are granted access by the super user. A WildFire user can be a user associated with an existing Palo Alto Networks account or a user not associated with a Palo Alto Networks support account, to whom you can be allow access to just the WildFire public clouds and a specific set of firewall data.

Add WildFire User Accounts		
Step 1	Select the account for which you want to add users who can access the WildFire portal. WildFire portal users can view data for all firewalls associated with the support account.	1. Log in to the Palo Alto Networks Support site. 2. Under **Manage Account** click on **Users and Accounts**. 3. Select an existing account or sub-account.
Step 2	Add a WildFire user.	1. Click the **Add WildFire User** button. 2. Enter the email address for the user you would like to add. The only restriction when adding a user is that the email address cannot be from a free web-based email account (such as Gmail, Hotmail, and Yahoo). If an email address is entered for a domain that is not supported, a pop-up warning is displayed.

Add WildFire User Accounts (Continued)		
Step 3	Assign firewalls to the new user account and access the WildFire cloud.	Select the firewall(s) by serial number for which you want to grant access and fill out the optional account details. Users with an existing support account will receive an email with a list of the firewalls that are now available for WildFire report viewing. If the user does not have a support account, the portal sends an email with instructions on how to access the portal and how to set a new password. The new user can now log in to the WildFire cloud and view WildFire reports for the firewalls to which they have been granted access. Users can also configure automatic email alerts for these devices in order to receive alerts on files analyzed. They can choose to receive reports on malicious and/or benign files.

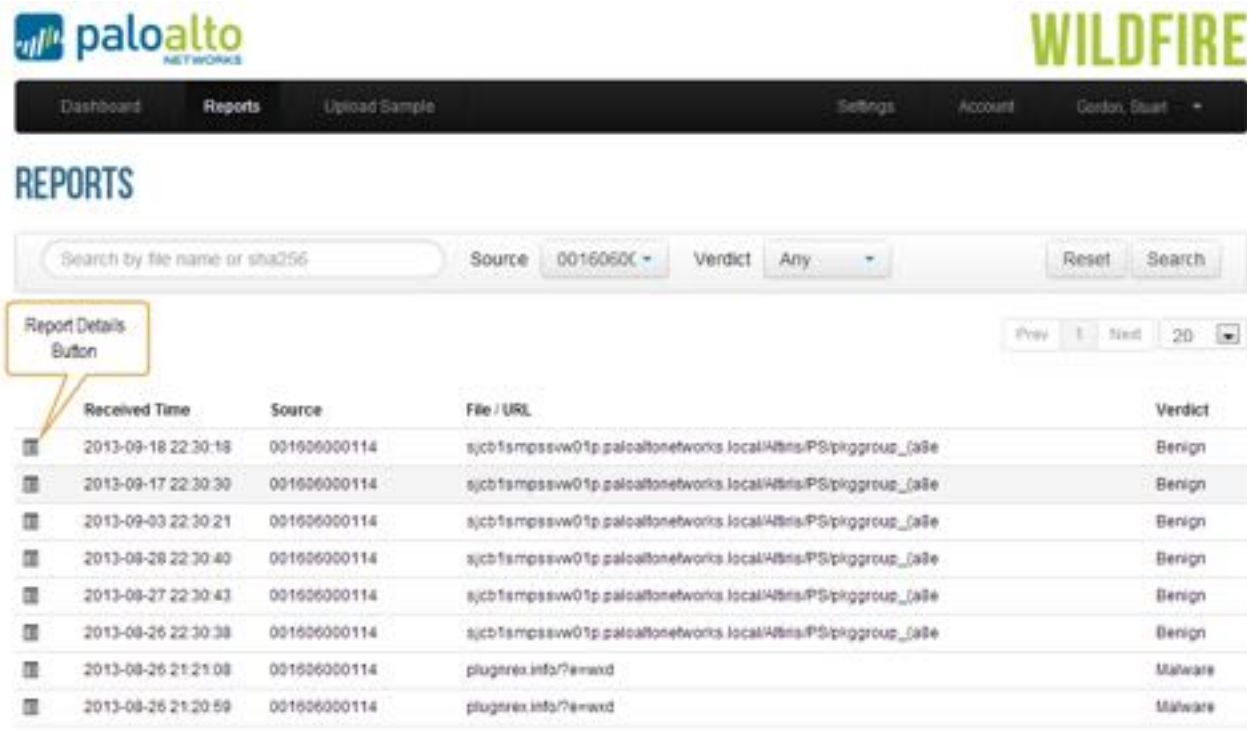
View Reports on the WildFire Portal

The Wildfire portal displays reports for samples that are submitted from firewalls, manually uploaded, or uploaded using the WildFire API. Select **Reports** to display the latest reports for samples analyzed by the WildFire cloud. For each sample listed, the report entry shows the date and time the sample was received by the cloud, the serial number of the firewall that submitted the file, the file name or URL, and the verdict delivered by WildFire (benign, grayware, or malware).

Use the search option to search for reports based on the file name or the sample hash value. You can also narrows the results displayed by viewing only reports for samples submitted by a specific **Source** (view only results submitted manually or by a specific firewall) or for samples that received a specific WildFire **Verdict** (any, benign, malware, grayware, or pending).

To view an individual report from the portal, click the **Reports** icon to the left of the report name. To save the detailed report, click the **Download as PDF** button on the upper right of the report page. For details on WildFire analysis reports, see WildFire Analysis Reports—Close Up.

The following shows a list of sample files submitted by a specific firewall:


paloalto NETWORKS
WILDFIRE
Dashboard
Reports
Upload Sample
Settings
Account
Gordon, Stuart
REPORTS
Search by file name or sha256
Source
Verdict
Any
Reset
Search
Report Details Button
Prev
1
Next
20
Received Time
Source
File / URL
Verdict
2013-09-18 22:30:18
001606000114
Benign
2013-09-17 22:30:30
001606000114
Benign
2013-09-03 22:30:21
001606000114
Benign
2013-08-28 22:30:40
001606000114
Benign
2013-08-27 22:30:43
001606000114
Benign
2013-08-26 22:30:38
001606000114
Benign
2013-08-26 21:21:08
001606000114
Malware
2013-08-26 21:20:59
001606000114
Malware

WildFire Analysis Reports—Close Up

Access WildFire analysis reports on the firewall, the WildFire portal, and the WildFire API.

WildFire analysis reports display detailed sample information, as well as information on the user who was targeted, email header information (if enabled), the application that delivered the file, and all URLs involved in the delivery or phone-home activity of the file. WildFire reports contain some or all of the information described in the following table based on the session information configured on the firewall that forwarded the file and depending on the observed behavior for the file.



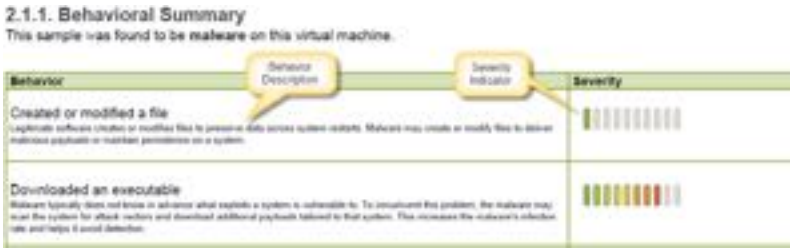
When viewing a WildFire report for a file that was manually uploaded to the WildFire portal or by using the WildFire API, the report will not show session information because the traffic did not traverse the firewall. For example, the report would not show the Attacker/Source and Victim/Destination.

Report Heading	Description

File Information	• **File Type**—Flash, PE, PDF, APK, JAR/Class, or MS Office. This field is named URL for HTTP/HTTPS email link reports and will display the URL that was analyzed.
	• **File Signer**—The entity that signed the file for authenticity purposes.
	• **Hash Value**—A file hash is much like a fingerprint that uniquely identifies a file to ensure that the file has not been modified in any way. The following lists the hash versions that WildFire generates for each file analyzed:
	  • **SHA-1**—Displays the SHA-1 value for the file.
	  • **SHA-256**—Displays the SHA-256 value for the file.
	  • **MD5**—Displays the MD5 information for the file.
	• **File Size**—The size (in bytes) of the file that WildFire analyzed.
	• **First Seen Timestamp**—If the WildFire system has analyzed the file previously, this is the date/time that it was first observed.
	• **Verdict**—Displays the analysis verdict:
	  • **Benign**—The file is safe and does not exhibit malicious behavior.
	  • **Grayware**—The file behaves similarly to malware, but does not pose a direct security threat. Grayware includes executables that display obtrusive behavior, but are not malicious in intent or nature. Examples of grayware include adware, spyware, and Browser Helper Objects (BHOs).
	  • **Malware**—WildFire identified the file as malware and generates a signature to protect against future exposure.
	• **Sample File**—Click the **Download File** link to download the sample file to your local system. Note that you can only download files with the malware verdict, not benign.

Report Heading	Description
Coverage Status	Click the **Virus Total** link to view endpoint antivirus coverage information for samples that have already been identified by other vendors. If the file has never been seen by any of the listed vendors, file not found appears. In addition, when the report is rendered on the firewall, up-to-date information about what signature and URL filtering coverage that Palo Alto Networks currently provides to protect against the threat will also be displayed in this section. Because this information is retrieved dynamically, it will not appear in the PDF report. The following screen capture shows coverage status that appears after rendering the report on the firewall:  Coverage type: wildfire; Signature ID: 5000000; Detail: Virus/Win32.WGeneric.a; Date Released: 2014-05-30 09:31:01; Content Version: 771 The following coverage information is provided for active signatures: • **Coverage Type**—The type of protection provided by Palo Alto Networks (virus, DNS, WildFire, or malware URL). • **Signature ID**—A unique ID number assigned to each signature that Palo Alto Networks provides. • **Detail**—The well-known name of the virus. • **Date Released**—The date that Palo Alto Networks released coverage to protect against the malware. • **Content Version**—The version number for the content release that provides protection against the malware.
Session Information	Contains session information based on the traffic as it traversed the firewall that forwarded the sample. To define the session information that WildFire will include in the reports, select **Device > Setup > WildFire> Session Information Settings**. The following options are available: • Source IP • Source Port • Destination IP • Destination Port • Virtual System (If multi-vsys is configured on the firewall) • Application • User (If User-ID is configured on the firewall) • URL • Filename • Email sender • Email recipient • Email subject

Report Heading	Description
Dynamic Analysis	If a file is low risk and WildFire can easily determine that it is safe, only a static analysis is performed, instead of a dynamic analysis. When a dynamic analysis is performed, this section contains tabs for each virtual environment that the sample was run in when it was analyzed in the WildFire cloud. For example, Virtual Machine 1 tab may have Windows XP, Adobe Reader 9.3.3, and Office 2003 and Virtual Machine 2 may have similar attributes, but with Office 2007. When a file goes through a full dynamic analysis, it is run in each virtual machine and the results of each environment can be viewed by clicking any of the Virtual Machine tabs.  On the WF-500 appliance, only one virtual machine is used for the analysis, which you select based on virtual environment attributes that best match your local environment. For example, if most users have Windows 7 32-bit, that virtual machine would be selected.

Report Heading	Description
Behavior Summary	Each Virtual Machine tab summarizes the behavior of the sample file in the specific environment. Examples include whether the sample created or modified files, started a process, spawned new processes, modified the registry, or installed browser helper objects. The Severity column indicates the severity of each behavior. The severity gauge will show one bar for low severity and additional bars for higher severity levels. This information is also added to the dynamic and static analysis sections.  The following describes the various behaviors that are analyzed: • **Network Activity**—Shows network activity performed by the sample, such as accessing other hosts on the network, DNS queries, and phone-home activity. A link is provided to download the packet capture. • **Host Activity (by process)**—Lists activities performed on the host, such as registry keys that were set, modified, or deleted. • **Process Activity**—Lists files that started a parent process, the process name, and the action the process performed. • **File**—Lists files that started a child processes, the process name, and the action the process performed. • **Mutex**—If the sample file generates other program threads, the mutex name and parent process is logged in this field. • **Activity Timeline**—Provides a play-by-play list of all recorded activity of the sample. This will help in understanding the sequence of events that occurred during the analysis. The activity timeline information is only available in the PDF export of the WildFire reports.
Submit Malware	Use this option to manually submit the sample to Palo Alto Networks. The WildFire cloud will then re-analyze the sample and generate a signatures if it determines that the sample is malicious. This is useful on a WF-500 appliance that does not have signature generation or cloud intelligence enabled, which is used to forward malware from the appliance to the WildFire cloud.
Report an Incorrect Verdict	Click this link to submit 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 malware sample is determined to be safe, the signature for the file is disabled in an upcoming antivirus signature update or if a benign file is determined to be malicious, a new signature is generated. After the investigation is complete, you will receive an email describing the action that was taken.

Use the WildFire API

The WildFire™ API enables you to send file analysis jobs to WildFire and query for report data through a simple XML API interface and is supported on the WildFire public cloud and for the WF-500 appliance. See the following sections for details on using the WildFire API:

- About WildFire Subscriptions and API Keys
- WildFire API File Submission Methods
- Query for a WildFire PDF or XML Report
- Use the API to Retrieve a Sample Malware Test File
- Use the API to Retrieve a Sample or PCAP
- Use the WildFire API on a WF-500 Appliance

About WildFire Subscriptions and API Keys

Access to the WildFire API key is provided if at least one Palo Alto Networks firewall has an active WildFire subscription registered to an account holder in your organization. You can share the same API key within your organization. The API key is displayed in the **My Account** section of the WildFire web portal, along with statistics, such as how many uploads and queries have been performed using the key. The key should be considered secret and should not be shared outside of authorized channels.

The WildFire API 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 REST services.

The API methods are hosted on the WildFire cloud and the HTTPS protocol (not HTTP) is required in order to protect your API key and any other data exchanged with the service.

A WildFire API key allows up to 1000 sample uploads per day and up to 10,000 report queries per day.

To use the WildFire API on a WF-500 appliance, you generate an API key from the appliance and use the IP address or FQDN in the URL used to locate the appliance. All other functions are the same as if you were using the API on the WildFire public cloud. For example, the URL to retrieve a report from the WildFire public cloud is https://**wildfire.paloaltonetworks.com**/publicapi/get/report. The URL to retrieve a report from a a WF-500 appliance with the IP address 10.3.4.50, would be as follows: https://**10.3.4.50**/publicapi/get/report. For an example, see Use the WildFire API on a WF-500 Appliance.

WildFire API File Submission Methods

Use the following methods to submit files to WildFire:

- Submit a File to the WildFire Cloud Using the Submit File Method
- Submit a File to WildFire Using the Submit URL Method

Submit a File to the WildFire Cloud Using the Submit File Method

The WildFire API can be used to submit all supported file types (for a list of supported file types see, File Type Analysis). The file along with your API key is required when submitting to WildFire for analysis. The return code of the submit-file method indicates a success or error condition. If a 200 OK code is returned, the submission is successful and a result is normally available for query within five minutes.

The following table describes the API attributes needed to submit files to the WildFire cloud using the submit file method:

URL	`https://wildfire.paloaltonetworks.com/publicapi/submit/file`	
Method	`POST`	
Parameters	`apikey`	Your WildFire API key
	`file`	The sample file to be analyzed
Return	`200 OK`	Indicates success and a report is returned
	`401 Unauthorized`	API key invalid
	`405 Method Not Allowed`	Method other than POST used
	`413 Request Entity Too Large`	Sample file size over max limit
	`418 Unsupported File Type`	Sample file type is not supported
	`419 Max Request Reached`	Max number of uploads per day exceeded
	`500`	Internal error
	`513`	File upload failed

Submit a File to WildFire Using the Submit URL Method

Use the submit-url method to submit a file for analysis via a URL. This method is identical in interface and functionality to the submit-file method, except that the file parameter is replaced with a url parameter. The url parameter must point to an accessible supported file type. If a 200 OK code is returned, the submission is successful and a result is usually available for query within five minutes.

The following table describes the API attributes needed to submit files to the WildFire cloud using a URL:

URL	`https://wildfire.paloaltonetworks.com/publicapi/submit/url`	
Method	`POST`	
Parameters	`apikey`	Your WildFire API key
	`url`	The URL for the file to be analyzed. The URL must contain the file name, for example http://paloaltonetworks.com/folder1/my-file.pdf.
Return	`200 OK`	Indicates success and a report is returned
	`401 Unauthorized`	API key invalid
	`405 Method Not Allowed`	Method other than POST used
	`413 Request Entity Too Large`	Sample file size over max limit
	`418 Unsupported File Type`	Sample file type is not supported
	`419 Max Request Reached`	Max number of uploads per day exceeded
	`422`	URL download error
	`500`	Internal error

Code Examples for File Submit

The following cURL command demonstrates how to submit a file to WildFire using the submit file method:

```
curl -k -F apikey=yourAPIkey -F  file=@local-file-path
https://wildfire.paloaltonetworks.com/publicapi/submit/file
```

The following shell code example demonstrates a simple script to submit a file to the WildFire API for analysis. The API key is provided as the first parameter and the path to the file is the second parameter:

```
#manual upload sample to WildFire with APIKEY
#Parameter 1: APIKEY
#Parameter 2: location of the file

key=$1
file=$2

/usr/bin/curl -i -k -F apikey=$key -F file=@$file
https://wildfire.paloaltonetworks.com/submit/file
```

The following cURL command demonstrates how to submit a file to WildFire using the submit URL method:

```
curl -k -F apikey=yourAPIkey -F  url=URL
https://wildfire.paloaltonetworks.com/publicapi/submit/url
```

Query for a WildFire PDF or XML Report

Use the get report method to query for an XML or PDF report of analysis results for a particular sample. Use either the MD5 or SHA-256 hash of the sample file as a search query.

The following table describes the API attributes needed to query for reports:

URL	`https://wildfire.paloaltonetworks.com/publicapi/get/report`	
Method	`POST`	
Parameters	`hash`	The MD5 or SHA-256 hash value of the sample
	`apikey`	Your WildFire API key
	`format`	Report format: PDF or XML
Return	`200 OK`	Indicates success and a report is returned
	`401 Unauthorized`	API key invalid
	`404 Not Found`	The report was not found
	`405 Method Not Allowed`	Method other than POST used
	`419`	Request report quota exceeded
	`420`	Insufficient arguments
	`421`	Invalid arguments
	`500`	Internal error

Example API Query for PDF or XML Report

The following cURL command demonstrates a query for a PDF report using the MD5 hash of a sample file:

```
curl -k -F hash=1234556 -F format=pdf -F apikey=yourAPIkey
https://wildfire.paloaltonetworks.com/publicapi/get/report
```



To retrieve the XML version of the report, replace `format=pdf` with `format=xml`. For example:

```
curl -k -F hash=1234556 -F format=xml -F apikey=yourAPIkey
https://wildfire.paloaltonetworks.com/publicapi/get/report
```

Use the API to Retrieve a Sample Malware Test File

The following describes the API syntax to retrieve a sample malware file, which can be used to test end-to-end WildFire sample processing.

For details on the sample file, see Test a Sample Malware File.

To retrieve the file using the API:

```
API : GET https://wildfire.paloaltonetworks.com/publicapi/test/pe
```

This will return a test file and every API call will return a similar file, but with a different SHA256 value.

If there is problem retrieving the file, a 500-Internal Server error is returned.

To retrieve the test file using cURL:

```
curl -k https://wildfire.paloaltonetworks.com/publicapi/test/pe
```

Use the API to Retrieve a Sample or PCAP

▲ Use the API to Retrieve a Sample
▲ Use the API to Retrieve a Packet Capture (PCAP)

Use the API to Retrieve a Sample

Use the get-sample method to retrieve a particular sample. You can use either the MD5 or SHA-256 hash of the sample file as a search query.

URL	`https://wildfire.paloaltonetworks.com/publicapi/get/sample`	
Method	`POST`	
Parameters	`hash`	The MD5 or SHA-256 hash value of the sample
	`apikey`	Your WildFire API key
Return	`200 OK`	Indicates success and a sample is returned
	`401 Unauthorized`	API key invalid
	`403 Forbidden`	Permission Denied
	`404 Not Found`	The sample was not found
	`405 Method Not Allowed`	Method other than POST used
	`419`	`Request sample quota exceeded`
	`420`	Insufficient arguments
	`421`	Invalid arguments
	`500`	Internal error

Example API Query for Get-Sample

The following cURL command demonstrates a query for a sample using the sample's MD5 hash:

```
curl -k -F hash=md5hash -F apikey=yourAPIkey
https://wildfire.paloaltonetworks.com/publicapi/get/sample
```

Use the API to Retrieve a Packet Capture (PCAP)

Use the get-pcap method to query for a PCAP recorded during analysis of a particular sample. Use either the MD5 or SHA-256 hash of the sample file as a search query. You can optionally define the platform of the desired PCAP to specify which PCAP should be returned. If no platform is specified, the method returns a PCAP from a session that yielded a verdict of **Malware**.



Samples uploaded prior to August 2014 are not guaranteed to return a PCAP if no platform parameter is supplied.

The following table describes the available platform parameters:

Platform ID	Description
1	Windows XP, Adobe Reader 9.3.3, Office 2003
2	Windows XP, Adobe Reader 9.4.0, Flash 10, Office 2007
3	Windows XP, Adobe Reader 11, Flash 11, Office 2010
4	Windows 7 32-bit, Adobe Reader 11, Flash 11, Office 2010
5	Windows 7 64bit, Adobe Reader 11, Flash 11, Office 2010.
201	Android 2.3, API 10, avd2.3.1

The following table describes the API attributes needed to query for pcaps:

URL	https://wildfire.paloaltonetworks.com/publicapi/get/pcap	
Method	POST	
Parameters	hash	The MD5 or SHA-256 hash value of the sample
	apikey	Your WildFire API key
	platform*	Target analysis environment
Return	200 OK	Indicates success and a PCAP is returned
	401 Unauthorized	API key invalid
	403 Forbidden	Permission Denied
	404 Not Found	The PCAP was not found
	405 Method Not Allowed	Method other than POST used
	419	Request sample quota exceeded
	420	Insufficient arguments
	421	Invalid arguments
	500	Internal error

* Optional parameter

Example API Query for Get-PCAP

The following cURL command demonstrates a query for a pcap using the sample's MD5 hash:

```
curl -k -F hash=md5hash -F apikey=yourAPIkey -F platform=targetPlatform
https://wildfire.paloaltonetworks.com/publicapi/get/pcap
```

Use the WildFire API on a WF-500 Appliance

To use the WildFire XML API on a WF-500 appliance, you must first generate an API key on the appliance and then use the API key from the host computer that performs the API functions. The URL used to locate the appliance is based on the IP address or FQDN of the appliance. After the keys are generated and you have the URL used to locate the appliance, you can then perform all of the same API functions supported on the WildFire cloud.

The following topics describe how to manage API keys on the appliance and provide an example on using the WildFire API to submit file samples to the appliance.

- Generate API Keys on the WF-500 Appliance
- Manage API Keys on the WF-500 Appliance
- Use the WildFire API on a WF-500 Appliance

Generate API Keys on the WF-500 Appliance

Generate an API Key	
Step 1 Generate a new API key on the WF-500 appliance. The appliance supports up to 100 API keys. As a best practice, leave out the key-value option in this step and the firewall will generate a key automatically. If you manually enter a key, the key-value must be 64 alpha characters (a-z) or numbers (0-9) that you randomly choose.	1. Log in to the WF-500 appliance CLI. 2. Generate the API key using one of the following methods: • Generate a key automatically: `admin@WF-500>` **`create wildfire api-key name`** ***`key-name`*** For example, to create a key with the name `my-api-key`: `admin@WF-500>` **`create wildfire api-key name my-api-key`** • To generate a key manually (where *key-value* is a 64-bit key): `admin@WF-500>` **`create wildfire api-key name my-api-key key`** *key-value* For example: `admin@WF-500>` **`create wildfire api-key name my-api-key key 0377785F3F1A3D2DC6BCF2342730700747FBF4A23BD69F455F142494BC43D4A1`**
Step 2 View the API keys that you generated.	View all API keys: `admin@WF-500>` **`show wildfire api-key all`** This command also shows the date the key was generated and the last time the key was used. In this example, the appliance generated the following key with the name *my-api-key*: **`0377785F3F1A3D2DC6BCF2342730700747FBF4A23BD69F455F142494BC43D4A1`**

Manage API Keys on the WF-500 Appliance

This section describes some useful commands that you can use to manage WildFire API keys on the appliance and describes how to export and import the keys. For example, you may want to export all of your keys for backup purposes or to make it easier to access the keys from the systems that will use the API to perform various functions on the appliance.

Manage API Keys

Disable or enable an API key:

```
admin@WF-500> edit wildfire api-key status [disable | enable] key api-key
```

For example, to disable the API key used in this example:

```
admin@WF-500> edit wildfire api-key status disable key
0377785F3F1A3D2DC6BCF2342730700747FBF4A23BD69F455F142494BC43D4A1
```



In the above command, you can type the first few unique digits of the key and then hit tab to fill in the remaining digits.

Delete an API key:

```
admin@WF-500> delete wildfire api-key key api-key
```

For example:

```
admin@WF-500> delete wildfire api-key key
377785F3F1A3D2DC6BCF2342730700747FBF4A23BD69F455F142494BC43D4A1
```

Use the following commands to import or export API keys from the appliance using Secure Copy (SCP):

- Save all API keys to a file to prepare the keys for export:

  ```
  admin@WF-500# save wildfire api-key to filename
  ```

  For example:

  ```
  admin@WF-500> save wildfire api-key to my-api-keys
  ```

  To SCP the API key file to an SCP-enabled server:

  ```
  admin@WF-500> scp export wildfire-api-keys to username@host:path
  ```

  For example:

  ```
  admin@WF-500> scp export wildfire-api-keys to bart@10.10.10.5:c:/scp/
  ```

  You can also import the key from an SCP-enabled server:

  ```
  admin@WF-500> scp import wildfire-api-keys from bart@10.10.10.5:c:/scp/my-api-keys
  ```

- After importing API keys, you must load the keys:

  ```
  admin@WF-500# load wildfire api-key mode [merge | replace] from my-api-keys
  ```

  If you leave out the `mode` option, the default behavior will merge the new keys. To replace all API keys on the appliance, use the `replace` option. For example, to replace all API keys, enter the command:

  ```
  admin@WF-500# load wildfire api-key mode replace from my-api-keys
  ```

- Confirm that the keys were loaded:

  ```
  admin@WF-500> show wildfire api-keys all
  ```

Use the WildFire API on a WF-500 Appliance

The following workflow describes how to use the WildFire API to submit a sample file to a WF-500 appliance for analysis. After understanding the basic concepts illustrated in this workflow, you can then use any of the API functions that are available on the WildFire cloud. The functions are the same, but in the case of the WF-500 appliance, you will use the API key generated on the appliance and the URL of the appliance.



This workflow requires a host computer that has the cURL command line tool installed. You will then send files from the host computer to the WF-500 appliance using the URL syntax.

Use the WildFire API to Submit a File Sample

Step 1 Generate a WildFire API key for the host computer that will perform API functions on the WF-500 appliance. For details, see Generate API Keys on the WF-500 Appliance.

1. Access the CLI on the WF-500 appliance and generate an API key:

   ```
   admin@WF-500> create wildfire api-key name my-api-key
   ```

2. View the API keys:

   ```
   admin@WF-500> show wildfire api-key all
   ```

3. Make sure the key status is Enabled and then highlight and copy the key. The following screen capture shows an example API key named *my-api-key*.



Use the WildFire API to Submit a File Sample (Continued)

Step 2 Using the new API key that you generated, submit a sample file to the WF-500 appliance.

1. Place a sample file in a folder that can be accessed from the host computer that has the cURL command line tool installed and note the path of the sample file.
2. Submit the file using cURL:

   ```
   curl -k -F apikey=your-API-key -F file=@local-file-path --remote-name
   https://WF-appliance-IP/publicapi/submit/file
   ```

   The syntax will vary based on the host that you are using. The following examples shows the syntax using a Linux host and a Windows host.

   From a Linux host:

   ```
   curl -k -F apikey=87C142CB01CA5BEBE06E226A25C0A473B34050B617073E21E8F1A6BCB8C5C387 -F
   file=@test-wf-api.docx --remote-name https://10.3.4.99/publicapi/submit/file
   ```

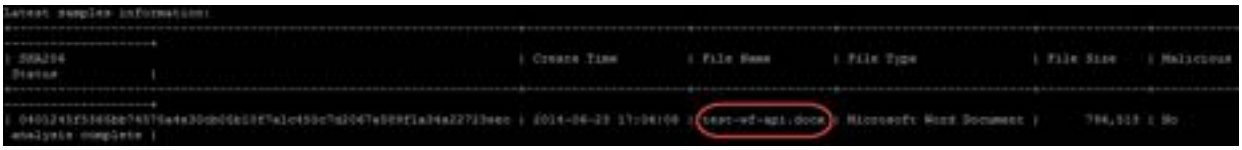
   From a Windows host (The only difference is the file path following the @ symbol):

   ```
   curl -k -F apikey=87C142CB01CA5BEBE06E226A25C0A473B34050B617073E21E8F1A6BCB8C5C387 -F
   file=@c://scp/test-wf-api.docx --remote-name https://10.3.4.99/publicapi/submit/file
   ```

3. Verify that the API successfully submitted the file to the WF-500 appliance. To view a list of recent samples submitted to the appliance:

   ```
   admin@WF-500> show wildfire latest samples
   ```

   The following screen capture shows that the sample file `test-wf-api.docx` successfully submitted to the appliance:



If the sample file does not appear on the appliance, verify connectivity between the host computer and the appliance and confirm that the folder/file path is correct. You can also run `show wildfire status` (status should show `Idle`) and `show wildfire statistics` to verify that the appliance is ready to analyze files. For more information on troubleshooting, refer to the Palo Alto Networks WildFire Administrator's Guide.

Use the WF-500 Appliance CLI

This section describes the CLI commands that are specific to the WildFire™ appliance software. All other commands, such as configuring interfaces, committing the configuration, and setting system information are identical to PAN-OS and are also shown in the hierarchy. For information on the PAN-OS commands, refer to the PAN-OS 7.0 CLI Quick Start.

- WF-500 Appliance Software CLI Concepts
- WildFire CLI Command Modes
- Access the CLI
- Use the CLI
- Configuration Mode Command Reference
- Operational Mode Command Reference

WF-500 Appliance Software CLI Concepts

This section introduces and describes how to use the WF-500 appliance software command line interface (CLI):

- WF-500 Appliance Software CLI Structure
- WF-500 Appliance Software CLI Command Conventions
- WF-500 Appliance CLI Command Messages
- Command Option Symbols
- Privilege Levels

WF-500 Appliance Software CLI Structure

The WF-500 appliance software CLI is used to manage the appliance. The CLI is the only interface to the appliance. Use it to view status and configuration information and modify the appliance configuration. Access the WF-500 appliance software CLI over SSH or by direct console access using the console port.

The WF-500 appliance software CLI operates in two modes:

- **Operational mode**—View the state of the system, navigate the WF-500 appliance software CLI, and enter configuration mode.
- **Configuration mode**—View and modify the configuration hierarchy.

WF-500 Appliance Software CLI Command Conventions

The basic command prompt incorporates the user name and hostname of the appliance:

```
username@hostname>
```

Example:

```
admin@WF-500>
```

When entering Configuration mode, the prompt changes from > to #:

```
username@hostname>(Operational mode)
username@hostname> configure
Entering configuration mode
[edit]
username@hostname# (Configuration mode)
```

In Configuration mode, the current hierarchy context is shown by the `[edit...]` banner presented in square brackets when a command is issued.

WF-500 Appliance CLI Command Messages

Messages may be displayed when issuing a command. The messages provide context information and can help in correcting invalid commands. In the following examples, the message is shown in bold.

Example: Unknown command

```
username@hostname# application-group
Unknown command: application-group
[edit network]
username@hostname#
```

Example: Changing modes

```
username@hostname# exit
Exiting configuration mode
username@hostname>
```

Example: Invalid syntax

```
username@hostname> debug 17
Unrecognized command
Invalid syntax.
username@hostname>
```

The CLI checks the syntax of each command. If the syntax is correct, it executes the command and the candidate hierarchy changes are recorded. If the syntax is incorrect, an invalid syntax message is presented, as in the following example:

```
username@hostname# set deviceconfig setting wildfire cloud-intelligence
submit-sample yes
Unrecognized command
Invalid syntax.
[edit]
username@hostname#
```

Command Option Symbols

The symbol preceding an option can provide additional information about command syntax.

Symbol	Description
*	This option is required.
>	There are additional nested options for this command.
+	There are additional command options for this command at this level.
|	There is an option to specify an “except value” or a “match value” to restrict the command.

Symbol	Description
“ “	Although the double quote is not a command option symbol, it must be used when entering multi-word phrases in CLI commands. For example, to create an address group named Test Group and to add the user named user1 to this group, you must surround the group name with double quotes as follows: set address-group “Test Group” user1. If you do not put a double quote surrounding the group name, the CLI would interpret the word Test as the group name and Group as the username and the following error wold be displayed: “test is not a valid name”.  A single quote would also be invalid in this example.

The following examples show how these symbols are used.

Example: In the following command, the keyword `from` is required:

```
  username@hostname> scp import configuration ?
  + remote-port   SSH port number on remote host
  * from          Source (username@host:path)
  username@hostname> scp import configuration
  Example: This command output shows options designated with + and >.
  username@hostname# set rulebase security rules rule1 ?
  + action               action
  + application          application
  + destination          destination
  + disabled             disabled
  + from                 from
  + log-end              log-end
  + log-setting          log-setting
  + log-start            log-start
  + negate-destination   negate-destination
  + negate-source        negate-source
  + schedule             schedule
  + service              service
  + source               source
  + to                   to
  > profiles             profiles
    <Enter>              Finish input
  [edit]
username@hostname# set rulebase security rules rule1
```

Each option listed with + can be added to the command.

The profiles keyword (with >) has additional options:

```
  username@hostname# set rulebase security rules rule1 profiles ?
```

```
  + virus           Help string for virus
  + spyware         Help string for spyware
  + vulnerability   Help string for vulnerability
  + group           Help string for group
    <Enter>         Finish input
[edit]
username@hostname# set rulebase security rules rule1 profiles
```

Privilege Levels

Privilege levels determine which commands the user is permitted to execute and the information the user is permitted to view.

Level	Description
superreader	Has complete read-only access to the appliance.
superuser	Has complete read-write access to the appliance.

WildFire CLI Command Modes

This section describes the modes used to interact with the WF-500 appliance software CLI:

- Configuration Mode
- Operational Mode

Configuration Mode

Entering commands in configuration mode modifies the candidate configuration. The modified candidate configuration is stored in the appliance memory and maintained while the appliance is running.

Each configuration command involves an action, and may also include keywords, options, and values.

This section describes Configuration mode and the configuration hierarchy:

- Configuration Mode Command Usage
- Configuration Hierarchy
- Navigate the Hierarchy

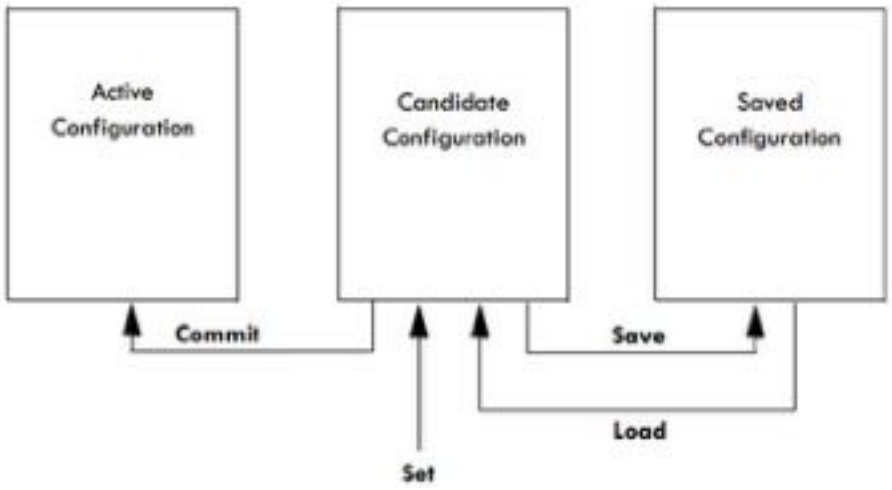
Configuration Mode Command Usage

Use the following commands to store and apply configuration changes:

- `save`—Saves the candidate configuration in the non-volatile storage on the appliance. The saved configuration is retained until overwritten by subsequent **save** commands. Note that this command does not make the configuration active.
- `commit`—Applies the candidate configuration to the appliance. A committed configuration becomes the active configuration for the device.
- `set`—Changes a value in the candidate configuration.
- `load`—Assigns the last saved configuration or a specified configuration to be the candidate configuration.



When exiting configuration mode without issuing the `save` or `commit` command, the configuration changes could be lost if the appliance loses power.



Maintaining a candidate configuration and separating the save and commit steps confers important advantages when compared with traditional CLI architectures:

- Distinguishing between the **save** and **commit** concepts allows multiple changes to be made at the same time and reduces system vulnerability.
- Commands can easily be adapted for similar functions. For example, when configuring two Ethernet interfaces, each with a different IP address, you can edit the configuration for the first interface, copy the command, modify only the interface and IP address, and then apply the change to the second interface.
- The command structure is always consistent.

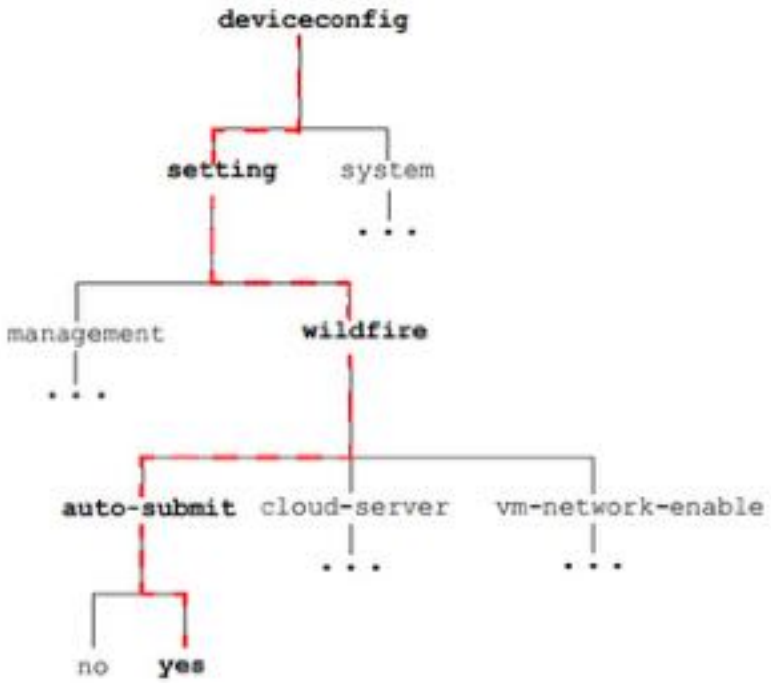
Because the candidate configuration is always unique, all authorized changes to the candidate configuration are consistent with each other.

Configuration Hierarchy

The configuration for the appliance is organized in a hierarchical structure. To display a segment of the current hierarchy level, use the `show` command. Entering show displays the complete hierarchy, while entering `show` with keywords displays a segment of the hierarchy. For example, when running the command `show` from the top level of configuration mode, the entire configuration is displayed. When running the command `edit mgt-config` and you enter `show`, or by running **show mgt-config**, only the mgt-config part of the hierarchy displays.

Hierarchy Paths

When entering commands, the path is traced through the hierarchy as follows:

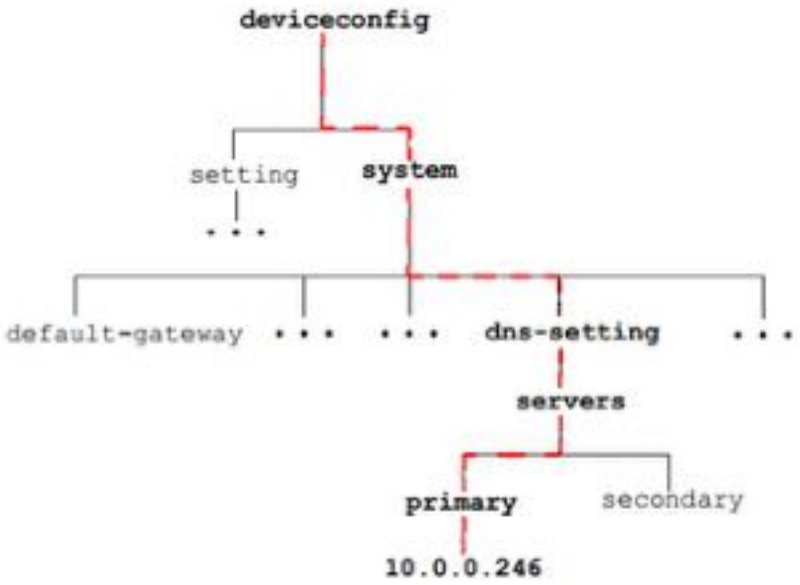


For example, the following command assigns the primary DNS server 10.0.0.246 for the appliance:

```
     [edit]
username@hostname# set deviceconfig system dns-setting servers primary
10.0.0.246
```

This command generates a new element in the hierarchy and in the output of the following `show` command:

```
[edit]
username@hostname# show deviceconfig system dns-settings
dns-setting {
    servers {
      primary 10.0.0.246
      }
  }
[edit]
username@hostname#
```



Navigate the Hierarchy

The [edit...] banner presented below the Configure mode command prompt line shows the current hierarchy context.

`[edit]`

indicates that the relative context is the top level of the hierarchy, whereas

`[edit deviceconfig]`

indicates that the relative context is at the deviceconfig level.

Use the commands listed in to navigate through the configuration hierarchy.

Level	Description
edit	Sets the context for configuration within the command hierarchy.
up	Changes the context to the next higher level in the hierarchy.
top	Changes the context to the highest level in the hierarchy.

The `set` command issued after using the `up` and `top` commands starts from the new context.

Operational Mode

At the initial login to the device, the WF-500 appliance software CLI opens in Operational mode. Operational mode commands involve actions that are executed immediately. They do not involve changes to the configuration, and do not need to be saved or committed.

Operational mode commands are of several types:

- **Network access**—Open a window to another host. SSH is supported.
- **Monitoring and troubleshooting**—Perform diagnosis and analysis. Includes **debug** and **ping** commands.
- **Display commands**—Display or clear current information. Includes **clear** and **show** commands.
- **WF-500 appliance software CLI navigation commands**—Enter Configure mode or exit the WF-500 appliance software CLI. Includes **configure**, **exit**, and **quit** commands.
- **System commands**—Make system-level requests or restart. Includes **set** and **request** commands.

Access the CLI

This section describes how to access and begin using the WF-500 appliance software CLI:

- Establish a Direct Console Connection
- Establish an SSH Connection

Establish a Direct Console Connection

Use the following settings for direct console connection:

- Data rate: 9600
- Data bits: 8
- Parity: none
- Stop bits: 1
- Flow control: None

Establish an SSH Connection

To access the WF-500 appliance software CLI:

Launch the WildFire CLI	
1.	Use terminal emulation software to establish an SSH console connection with the WF-500 appliance.
2.	Enter the administrative user name. The default is admin.
3.	Enter the administrative password. The default is admin. The WF-500 appliance software CLI opens in Operational mode, and the CLI prompt is displayed: `username@hostname>`

Use the CLI

- Access Operational and Configuration Modes
- Display WF-500 Appliance Software CLI Command Options
- Restrict Command Output
- Set the Output Format for Configuration Commands

Access Operational and Configuration Modes

When logging in, the WF-500 appliance software CLI opens in Operational mode. You can navigate between Operational and Configuration modes at any time.

- 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, to show system resources from configure mode, use `run show system resources`.

Display WF-500 Appliance Software CLI Command Options

Use ? (or **Meta-H**) to display a list of command options, based on context:

- To display a list of operational commands, enter **?** at the command prompt.

```
username@hostname> ?
  clear       Clear runtime parameters
  configure   Manipulate software configuration information
  debug       Debug and diagnose
  exit        Exit this session
  grep        Searches file for lines containing a pattern match
  less        Examine debug file content
  ping        Ping hosts and networks
  quit        Exit this session
  request     Make system-level requests
  scp         Use ssh to copy file to another host
  set         Set operational parameters
  show        Show operational parameters
  ssh         Start a secure shell to another host
  tail        Print the last 10 lines of debug file content
```

```
username@hostname>
```

- To display the available options for a specified command, enter the command followed by **?**.

Example:

```
username@hostname> ping ?
+ bypass-routing    Bypass routing table, use specified interface
+ count             Number of requests to send (1..2000000000 packets)
+ do-not-fragment   Don't fragment echo request packets (IPv4)
+ inet              Force to IPv4 destination
+ interface         Source interface (multicast, all-ones, unrouted packets)
+ interval          Delay between requests (seconds)
+ no-resolve        Don't attempt to print addresses symbolically
+ pattern           Hexadecimal fill pattern
+ record-route      Record and report packet's path (IPv4)
+ size              Size of request packets (0..65468 bytes)
+ source            Source address of echo request
+ tos               IP type-of-service value (0..255)
+ ttl               IP time-to-live value (IPv6 hop-limit value) (0..255 hops)
+ verbose           Display detailed output
+ wait              Delay after sending last packet (seconds)
  <host>            Hostname or IP address of remote host
```

Restrict Command Output

Some operational commands include an option to restrict the displayed output. To restrict the output, enter a pipe symbol followed by `except` or `match` and the value that is to be excluded or included:

Example:

The following sample output is for the `show system info` command:

```
username@hostname> show system info
hostname: WF-500
ip-address: 192.168.2.20
netmask: 255.255.255.0
default-gateway: 192.168.2.1
mac-address: 00:25:90:95:84:76
vm-interface-ip-address: 10.16.0.20
vm-interface-netmask: 255.255.252.0
vm-interface-default-gateway: 10.16.0.1
vm-interface-dns-server: 10.0.0.247
time: Mon Apr 15 13:31:39 2013
uptime: 0 days, 0:02:35
family: m
model: WF-500
serial: 009707000118
sw-version: 5.1.0
logdb-version: 5.0.2
platform-family: m

username@hostname>
```

```
  The following sample displays only the system model information:

  username@hostname> show system info | match model
  model: WF-500

username@hostname>
```

Set the Output Format for Configuration Commands

Change the output format for the configuration commands by using the `set cli config-output-format` command in Operational mode. Options include the default format, JSON (JavaScript Object Notation), set format, and XML format. The default format is a hierarchal format where configuration sections are indented and enclosed in curly brackets.

Configuration Mode Command Reference

This section contains command reference information for the following Configuration mode commands that are specific to the WF-500 appliance software. All other commands that are part of the WF-500 appliance software are identical to PAN-OS as described in the PAN-OS 7.0 CLI Quick Start.

- set deviceconfig setting wildfire
- set deviceconfig system update-schedule
- set deviceconfig system vm-interface

set deviceconfig setting wildfire

Description

Configure Wildfire settings on the WF-500 appliance. You can configure forwarding of malicious files, define the cloud server that receives malware infected files, and enable or disable the vm-interface.

Hierarchy Location

```
set deviceconfig settings
```

Syntax

```
wildfire {
  active-vm;
  cloud-server <value>;
  vm-network-enable {no | yes};
  vm-network-use-tor {enable | disable};
  cloud-intelligence {
     submit-report {no | yes};
     submit-sample {no | yes};
  signature-generation {
     av {no | yes};
     dns {no | yes};
     url {no | yes};
     {
  {
{
```

Options

+ active-vm — Select the virtual machine environment that WildFire will use for sample analysis. Each vm has a different configuration, such as Windows XP, a specific versions of Flash, Adobe reader, etc. To view which VM is selected, run the following command: **admin@WF-500> show wildfire status** and view the Selected VM field. To view the VM environment information, run the following command: **admin@WF-500> show wildfire vm-images**.
+ cloud-server — Hostname for the cloud server that the appliance will forward malicious samples/reports to for a re-analysis. The default cloud server is wildfire-public-cloud. To configure forwarding, use the following command: **set deviceconfig setting wildfire cloud-intelligence**.
+ vm-network-enable — Enable or disable the vm-network. When enabled, sample files running in the virtual machine sandbox can access the Internet. This helps WildFire better analyze the behavior of the malware to look for things like phone home activity.
+ vm-network-use-tor — Enable or disable the Tor network for the vm-interface. When this option is enabled, any malicious traffic coming from the sandbox systems on the WF-500 appliance during sample analysis is sent through the Tor network. The Tor network will mask your public facing IP address, so the owners of the malicious site cannot determine the source of the traffic.
+ cloud-intelligence — Configure the appliance to submit WildFire reports or samples to the Palo Alto Networks WildFire cloud. The submit report option will send reports for malicious samples to the cloud for statistical gathering. The submit sample option will send malicious samples to the cloud. If submit-sample enabled, there is no need to enable submit-report because the sample is re-analyzed in the cloud and a new report and signature is generated if the sample is malicious.
+ signature-generation — Enable the appliance to generate signatures locally, eliminating the need to send any data to the public cloud in order to block malicious content. The WF-500 appliance will analyze files forwarded to it from Palo Alto Networks firewalls or from the WildFire API and generate antivirus and DNS signatures that block both the malicious files as well as associated command and control traffic. When the appliance detects a malicious URL, it sends the URL to PAN-DB and PAN-DB assigns it the malware category.

Sample Output

The following shows an example output of the WildFire settings.

```
admin@WF-500# show deviceconfig setting wildfire
wildfire {
  active-vm vm-5;
  cloud-intelligence {
    submit-sample yes;
    submit-report no;
  }
  cloud-server wildfire-public-cloud;
  signature-generation {
    av yes;
    dns yes;
    url yes;
  }
}
```

Required Privilege Level

▲ superuser, deviceadmin

set deviceconfig system update-schedule

Description

Schedule content updates on a WF-500 appliance. These content updates equip the appliance with the most up-to-date threat information for accurate malware detection and improve the appliance's ability to differentiate the malicious from the benign.

Hierarchy Location

```
set deviceconfig system update-schedule
```

Syntax

```
wf-content recurring {
  daily at <value> action {download-and-install | download-only};
  weekly {
     action {download-and-install | download-only};
     at <value>;
     day-of-week {friday | monday | saturday | sunday | thursday | tuesday | wednesday};
  }
}
```

Options

```
> wf-content — WF-500 content updates
> daily — Schedule update every day
+ action — Specify the action to take. You can schedule the appliance to download and
install the update or download only and then you install manually
+ at — Time specification hh:mm (e.g. 20:10)
> hourly — Schedule update every hour
+ action — Specify the action to take. You can schedule the appliance to download and
install the update or download only and then you install manually
+ at — Minutes past the hour
> weekly — Schedule update once a week
+ action — Specify the action to take. You can schedule the appliance to download and
install the update or download only and then you install manually
+ at — Time specification hh:mm (e.g. 20:10)
+ day-of-week — Day of the week (Friday, Monday, Saturday, Sunday, Thursday, Tuesday,
Wednesday)
```

Sample Output

```
admin@WF-500# show
update-schedule {
  wf-content {
    recurring {
      weekly {
        at 19:00;
        action download-and-install;
        day-of-week friday;
      }
    }
  }
}
```

Required Privilege Level

superuser, deviceadmin

set deviceconfig system vm-interface

Description

The vm-interface is used by malware running on the WF-500 appliance virtual machine sandbox to access the Internet. Activating this port is recommended and will help WildFire better identify malicious activity if the malware accesses the Internet for phone-home or other activity. It is important that this interface has an isolated connection to the Internet. For more information, see Set Up the VM Interface.

After configuring the vm-interface, enable it by running the following command:

set deviceconfig setting wildfire vm-network-enable yes

Hierarchy Location

set deviceconfig system

Syntax

```
set vm-interface {
   default-gateway <ip_address>;
   dns-server <ip_address>;
   ip-address <ip_address>;
   link-state;
   mtu;
   netmask <ip_address>;
```

```
    speed-duplex;
{
```

Options

```
admin@WF-500# set vm-interface
+ default-gateway — Default gateway for the VM interface
+ dns-server — dns server for the VM interface
+ ip-address — IP address for VM interface
+ link-state — Set the link state to up or down
+ mtu — Maximum Transmission Unit for the VM interface
+ netmask — IP netmask for the VM interface
+ speed-duplex — Speed and duplex for the VM interface
```

Sample Output

The following shows a configured vm-interface.

```
vm-interface {
  ip-address 10.16.0.20;
  netmask 255.255.252.0;
  default-gateway 10.16.0.1;
  dns-server 10.0.0.246;
}
```

Required Privilege Level

```
superuser, deviceadmin
```

Operational Mode Command Reference

This section contains command reference information for the following Operational mode commands that are specific to the WF-500 appliance software. All other commands that are part of the WF-500 appliance software are identical to PAN-OS; refer to the PAN-OS 7.0 CLI Quick Start for information on those commands.

- create wildfire api-key
- delete wildfire api-key
- delete wildfire-metadata
- edit wildfire api-key
- load wildfire api-key
- request system raid
- request system wildfire-vm-image
- request wf-content
- save wildfire api-key
- set wildfire portal-admin
- show system raid
- show wildfire
- test wildfire registration

create wildfire api-key

Description

Generate API keys on a WF-500 appliance that you will use on an external system to submit samples to the appliance, query reports, or retrieve samples and Packet Captures (PCAPS) from the appliance.

Syntax

```
create {
  wildfire {
     api-key {
        key <value>;
        name <value>;
     {
  {
{
```

Options

+ key — Create an API key by manually entering a key value. The value must be 64 alpha characters (a-z) or numbers (0-9). If you do not specify the key option, the appliance generates a key automatically.
+ name — Optionally enter a name for the API key. An API key name is simply used to label the keys to make it easier to identify keys assigned for specific uses and has no impact on the functionality of the key.

Sample Output

The following output shows that the appliance has three API keys and one key is named my-api-key.

admin@WF-500> **show wildfire api-keys all**

```
+------------------------------------------------------------------+------------
----+---------+--------------------+--------------------+
| Apikey                                                           | Name
    | Status  | Create Time        | Last Used Time     |
+------------------------------------------------------------------+------------
----+---------+--------------------+--------------------+
| C625DE87CBFB6EF0B1A8183A74AB5B61287F7F63B6E14E2FFC704AABF5640D62 | my-api-key
| Enabled | 2014-06-24 16:38:50 |                    |
| D414CC910E93E9E05942A5E6F94DA36777B444543E71761CF5E9ACFA547F7D6F |
    | Enabled | 2014-06-25 09:05:30 | 2014-06-26 14:49:35 |
| 73585ACAFEC0109CB65EB944B8DFC0B341B9B73A6FA7F43AA9862CAD47D0884C |
    | Enabled | 2014-08-04 17:00:42 |                    |
+------------------------------------------------------------------+------------
----+---------+--------------------+--------------------+
```

Required Privilege Level

superuser, deviceadmin

delete wildfire api-key

Description

Delete an API key from the WF-500 appliance. Systems configured to use the API to perform API functions on the appliance will no longer be able to access the appliance after you delete the key.

Syntax

```
delete {
  wildfire {
     api-key {
        key <value>;
     {
  {
{
```

Options

+ key <value> — The key value for the key that you want to delete. To view a list of API keys, run the following command: **admin@WF-500> show wildfire api-keys all**

Sample Output

```
admin@WF-500> delete wildfire api-key key
A0418F8EADABA4C78CD3106D71147321462C5AA085B2979136447B1EC334655A

APIKey A0418F8EADABA4C78CD3106D71147321462C5AA085B2979136447B1EC334655A
deleted
```

Required Privilege Level

superuser, deviceadmin

delete wildfire-metadata

Description

Delete content updates on the WF-500 appliance. For more information on content updates and how to install them, see request wf-content.

Syntax

```
delete {
  wildfire-metadata update <value>;
     {
```

Options

+ update <value> — Define the content update that you want to delete.

Sample Output

The output that follows shows the deletion of an update named panup-all-wfmeta-2-181.candidate.tgz.

```
admin@WF-500> delete wildfire-metadata update panup-all-wfmeta-2-181.candidate.tgz

successfully removed panup-all-wfmeta-2-181.candidate.tgz
```

Required Privilege Level

```
superuser, deviceadmin
```

edit wildfire api-key

Description

Modify an API key name or the key status (enabled/disabled) on a WF-500 appliance.

Syntax

```
edit {
  wildfire {
     api-key [name | status] key <value>;
     {
  {
```

Options

```
+ name — Change the name of an API key
+ status — Enable or disable an API key
* key — Specify the key to modify
```

Sample Output

The key value in this command is required. For example, to change the name of a key named `stu` to `stu-key1`, enter the following command:



In the following command, you do not need to enter the old key name; only enter the new key name.

```
admin@WF-500> edit wildfire api-key name stu-key1 key
B870210A6BDF2615D5A40B2DE515A6F5E66186BE28E4FFAC4405F22E83329288
```

To change the status of stu-key1 to disabled, enter the following command:

```
admin@WF-500> edit wildfire api-key status disable key
B870210A6BDF2615D5A40B2DE515A6F5E66186BE28E4FFAC4405F22E83329288
```

Example output that shows that stu-key1 is disabled:

```
admin@WF-500> show wildfire api-keys all
+-------------------------------------------------------------------+----------+----------+--------------------
+---------------------+
```

```
| Apikey                                                                  | Name     | Status   | Create Time         |
Last Used Time      |
+-------------------------------------------------------------------------+----------+----------+---------------------
+---------------------+
|
|B870210A6BDF2615D5A40B2DE515A6F5E66186BE28E4FFAC4405F22E83329288|stu-key1|Disabled|2014-08-21 07:23:34|
|
+-------------------------------------------------------------------------+----------+----------+---------------------
+---------------------+
```

Required Privilege Level

superuser, deviceadmin

load wildfire api-key

Description

After importing API keys to the WF-500 appliance, you must use the load command to make the keys available for use. Use this command to replace all existing API keys, or you can merge the keys in the import file with the existing key database.

Syntax

```
load {
  wildfire {
     from <value> mode [merge | replace];
     {
  {
```

Options

```
* from — Specify the API key filename that you want to import. The key files use the
.keys file extension. For example, my-api-keys.keys. To view a list of keys that are
available for import, enter the following command:
admin@WF-500> load wildfire api-key from ?
+ mode — Optionally enter the mode for the import (merge/replace). For example, to
replace the key database on the appliance with the contents of the contents of the new
key file, enter the following command:
admin@WF-500> load wildfire api-key mode replace from my-api-keys.keys
```

If you do not specify the mode option, the default action will merge the keys.

Required Privilege Level

```
superuser, deviceadmin
```

request system raid

Description

Use this option to manage the RAID pairs installed in the WF-500 appliance. The WF-500 appliance ships with four drives in the first four drive bays (A1, A2, B1, B2). Drives A1 and A2 are a RAID 1 pair and drives B1 and B2 are a second RAID 1 pair.

Hierarchy Location

request system

Syntax

```
raid {
      remove <value>;
      OR...
      copy {
        from <value>;
        to <value>;
      }
      OR...
        add {
```

Options

```
> add      Add a drive into the corresponding RAID Disk Pair
> copy     Copy and migrate from one drive to other drive in the bay
> remove   drive to remove from RAID Disk Pair
```

Sample Output

The following output shows a WF-500 appliance with a correctly configured RAID.

```
admin@WF-500> show system raid

Disk Pair A                          Available
   Disk id A1                          Present
   Disk id A2                          Present
Disk Pair B                          Available
   Disk id B1                          Present
   Disk id B2                          Present
```

Required Privilege Level

superuser, deviceadmin

request system wildfire-vm-image

Perform upgrades on the WF-500 appliance virtual machine (VM) sandbox images used to analyze files. To retrieve new VM images from the Palo Alto Networks Update Server, you must first download the image manually, host it on an SCP enabled server, and then retrieve the image from the appliance using the SCP client. After downloading the image to the appliance, you can then install it using this command.

Hierarchy Location

request system

Syntax

```
request {
  system {
     wildfire-vm-image {
        upgrade install file <value>;
        }
     }
  }
```

Options

```
> wildfire-vm-image —— Install Virtual Machine (VM) images.
+ upgrade install file —— Perform an upgrade to the VM image. After the file option,
type ? to view a list of available VM images. For example, run the following command to
list available images: admin@WF-500> request system wildfire-vm-image
upgrade install file ?
```

Sample Output

To list available VM images, run the following command:

```
admin@WF-500> request system wildfire-vm-image upgrade install file ?
```

To install a VM image (Windows 7 64-bit in this example), run the following command:

```
admin@WF-500> request system wildfire-vm-image upgrade install file
WFWin7_64Base_m-1.0.0_64base
```

Required Privilege Level

```
superuser, deviceadmin
```

request wf-content

Perform content updates on a WF-500 appliance. These content updates equip the appliance with the most up-to-date threat information for accurate malware detection and improve the appliance's ability to differentiate the malicious from the benign. To schedule content updates to install automatically, see set deviceconfig system update-schedule and to delete content updates on the WF-500 appliance, see delete wildfire-metadata.

Hierarchy Location

```
request
```

Syntax

```
request wf-content
  {
  downgrade install {previous | <value>};
  upgrade
     {
     check
     download latest
     info
     install {
        file <filename>
        version latest;
        }
     }
  }
```

Options

```
> downgrade — Installs a previous content version. Use the previous option to install
the previously installed content package or enter a value to downgrade to a specific
content package number.
> upgrade — Performs content upgrade functions
> check — Obtain information on available content packages from the Palo Alto Networks
Update Server
> download — Download a content package
> info — Show information about available content packages
> install — Install a content package
> file — Specify the name of the file containing the content package
> version — Download or upgrade based on the version number of the content package
```

Sample Output

To list available content updates, run the following command:

```
admin@WF-500> request wf-content upgrade check

Version               Size              Released on Downloaded  Installed
-------------------------------------------------------------------------
2-217                 58MB  2014/07/29 13:04:55 PDT        yes    current
2-188                 58MB  2014/07/01 13:04:48 PDT        yes   previous
2-221                 59MB  2014/08/02 13:04:55 PDT         no         no
```

Required Privilege Level

superuser, deviceadmin

save wildfire api-key

Description

Use the save command to save all API keys on the WF-500 appliance to a file. You can then export the key file for backup purposes or to modify the keys in bulk. For details on using the WildFire API on a WF-500 appliance, see About WildFire Subscriptions and API Keys.

Hierarchy Location

save

Syntax

```
save {
  wildfire {
     api-key to <value>;
     {
  {
```

Options

* to — Enter the filename for key export. For example, to export all of the API keys on the WF-500 to a file named my-wf-keys, enter the following command:

```
admin@WF-500> save wildfire api-key to my-wf-keys
```

Required Privilege Level

```
superuser, deviceadmin
```

set wildfire portal-admin

Description

Sets the portal admin account password that an administrator will use to view WildFire analysis reports generated by a WF-500 appliance. The account name (admin) and password is required when viewing the report on the firewall or from Panorama in **Monitor > WildFire Submissions > View WildFire Report**. The default username and password is admin/admin.



The portal admin account is the only account that you configure on the appliance to view reports from the firewall or Panorama. You cannot create new accounts or change the account name. This is not the same admin account used to manage the appliance.

Hierarchy Location

```
set wildfire
```

Syntax

```
set {
  wildfire {
    portal-admin {
      password <value>;
    }
  }
```

Sample Output

The following shows the output of this command.

```
admin@WF-500> set wildfire portal-admin password
Enter password:
Confirm password:
```

Required Privilege Level

```
superuser, deviceadmin
```

show system raid

Description

Show the RAID configuration of the appliance. The WF-500 appliance ships with four drives in the first four drive bays (A1, A2, B1, B2). Drives A1 and A2 are a RAID 1 pair and drives B1 and B2 are a second RAID 1 pair.

Hierarchy Location

```
show system
```

Syntax

```
raid {
        detail;
  {
```

Options

```
No additional options.
```

Sample Output

The following shows the RAID configuration on a functioning WF-500 appliance.

```
admin@WF-500> show system raid detail

Disk Pair A                             Available
   Status                                   clean
   Disk id A1                             Present
       model        : ST91000640NS
       size         : 953869 MB
       partition_1  : active sync
       partition_2  : active sync
   Disk id A2                             Present
       model        : ST91000640NS
       size         : 953869 MB
       partition_1  : active sync
       partition_2  : active sync
Disk Pair B                             Available
   Status                                   clean
   Disk id B1                             Present
       model        : ST91000640NS
       size         : 953869 MB
       partition_1  : active sync
```

```
      partition_2  : active sync
   Disk id B2                          Present
      model        : ST91000640NS
      size         : 953869 MB
      partition_1  : active sync
      partition_2  : active sync
```

Required Privilege Level

superuser, superreader

show wildfire

Description

Shows various information about the WF-500 appliance, such as available API keys, registration information, activity, recent samples that the appliance analyzed, and the virtual machine that is selected to perform analysis.

Hierarchy Location

show wildfire

Syntax

```
api-keys
  all {
     details;
     }
  key <value>;
  }
last-device-registration all |
latest {
      analysis {
        filter malicious|benign;
        sort-by SHA256|Submit Time|Start Time|Finish Time|Malicious|Status;
        sort-direction asc|desc;
        limit 1-20000;
        days 1-7;
      }
      OR...
       samples {
        filter malicious|benign;
        sort-by SHA256|Create Time|File Name|File Type|File Size|Malicious|Status;
        sort-direction asc|desc;
```

```
        limit 1-20000;
        days 1-7;
      }
      OR...
       sessions {
        filter malicious|benign;
        sort-by SHA256|Create Time|Src IP|Src Port|Dst Ip|Dst Port|File|Device
        ID|App|Malicious|Status;
        sort-direction asc|desc;
        limit 1-20000;
        days 1-7;
      }
     OR...
      uploads {
        sort-by SHA256|Create Time|Finish Time|Status;
        sort-direction asc|desc;
        limit 1-20000;
        days 1-7;
      }
sample-status {
      sha256 {
        equal <value>;
       }
     }
statistics days <1-31>;
status |
vm-images |
}
```

Options

```
admin@WF-500> show wildfire
> api-keys — Show details about the API keys generated on the WF-500 appliance. You can
view the last time the key was used, the key name, status (Enabled or Disabled), and the
date/time the key was generated.
> last-device-registration — Show list of latest registration activities.
> latest — Show latest 30 activities, which include the last 30 analysis activities, the
last 30 files that were analyzed, network session information on files that were
analyzed and files that were uploaded to the public cloud server.
> sample-status — Show wildfire sample status. Enter the SHA or MD5 value of the file to
view the current analysis status.
> statistics — Display basic wildfire statistics.
> status — Display the status of the appliance as well as configuration information such
as the Virtual Machine (VM) used for sample analysis, whether or not samples/reports are
sent to the cloud, vm network, and registration information.
> vm-images — Display the attributes of the available virtual machine images used for
sample analysis. To view the current active image, run the following command:
admin@WF-500> show wildfire status and view the Select VM field.
```

Sample Output

The following shows the output for this command.

```
admin@WF-500> show wildfire api-keys all
+------------------------------------------------------------------+---------------+-
--------+--------------------+--------------------+
| Apikey                                                           | Name          |
Status  | Create Time        | Last Used Time     |
+------------------------------------------------------------------+---------------+-
--------+--------------------+--------------------+
| C625DE87CBFB6EF0B1A8183A74AB5B61287F7F63B6E14E2FFC704AABF5640D62 | my-api-key-stu |
Enabled | 2014-06-24 16:38:50 |                    |
| D414CC910E93E9E05942A5E6F94DA36777B444543E71761CF5E9ACFA547F7D6F |               |
Enabled | 2014-06-25 09:05:30 | 2014-06-26 14:49:35 |
+------------------------------------------------------------------+---------------+-
--------+--------------------+--------------------+

admin@WF-500> show wildfire last-device-registration all
+--------------+--------------------+-------------+------------+----------+--------+
| Device ID    | Last Registered    | Device IP   | SW Version | HW Model | Status |
+--------------+--------------------+-------------+------------+----------+--------+
| 001606000114 | 2014-07-31 12:35:53 | 10.43.14.24 | 6.1.0-b14  | PA-200   | OK     |
+--------------+--------------------+-------------+------------+----------+--------+

admin@WF-500> show wildfire latest
> analysis   Show latest 30 analysis
> samples    Show latest 30 samples
> sessions   Show latest 30 sessions
> uploads    Show latest 30 uploads

admin@WF-500> show wildfire sample-status sha256 equal
809bad2d3fbdf1c18ef47ba9c5a0feca691103f094bc8d7e0cbed480870fd78c

Sample information:
+--------------------+-----------------------------------------------------------------
+------------------+----------+----------+-------------------+
| Create Time        | File Name                                                       |
File Type        | File Size | Malicious | Status            |
+--------------------+-----------------------------------------------------------------
+------------------+----------+----------+-------------------+
| 2014-08-04 11:49:41 | 25047801_20130919175646000_970x66_Adobe_Marketing_RM_AUTO.swf |
Adobe Flash File | 64502     | No        | analysis complete |
+--------------------+-----------------------------------------------------------------
+------------------+----------+----------+-------------------+

Session information:
+--------------------+--------------+---------+-------------+----------+----------
----------------------------------------------------+--------------+-------+
-----------+-----------+
| Create Time     | Src IP    | Src Port | Dst IP    | Dst Port | File
| Device ID     | App    |
 Malicious | Status    |
```

```
+--------------------+--------------+----------+--------------+----------+----------
-------------------------------------------------+--------------+--------+
-----------+-----------+
| 2014-08-04 11:49:41 | 10.10.10.50 | 80       | 192.168.2.10 | 64108    |
25047801_20130919175646000_970x66_Adobe_Marketing_RM_AUTO.swf | 001606000114 | flash |
 No        | completed |
+--------------------+--------------+----------+--------------+----------+----------
-------------------------------------------------+--------------+--------+
-----------+-----------+

Analysis information:
+--------------------+--------------------+--------------------+-----------+-------
-----------------------------------------------------+-----------+
| Submit Time    | Start Time     | Finish Time    | Malicious | VM Image
| Status    |
+--------------------+--------------------+--------------------+-----------+-------
-----------------------------------------------------+-----------+
| 2014-08-04 11:49:41 | 2014-08-04 11:49:41 | 2014-08-04 11:56:52 | No        | Windows
7 x64 SP1, Adobe Reader 11, Flash 11, Office 2010 | completed |
+--------------------+--------------------+--------------------+-----------+-------
-----------------------------------------------------+-----------+

admin@WF-500> show wildfire statistics

Last one hour statistics :
Total sessions submitted :            0
Samples submitted        :            0
  analyzed               :            0
  pending                :            0
  malicious              :            0
  benign                 :            0
  error                  :            0
  uploaded               :            0

Last 24 hours statistics :
Total sessions submitted :           13
Samples submitted        :           13
  analyzed               :           13
  pending                :            0
  malicious              :            0
  benign                 :           13
  error                  :            0
  uploaded               :            0

admin@WF-500> show wildfire status

Connection info:
  Wildfire cloud:               s1.wildfire.paloaltonetworks.com
  Status:                       Idle
  Submit sample:                disabled
  Submit report:                disabled
  Selected VM:                  vm-5
  VM internet connection:       disabled
```

```
  VM network using Tor:          disabled
  Best server:                   s1.wildfire.paloaltonetworks.com
  Device registered:             yes
  Service route IP address:      10.3.4.99
  Signature verification:        enable
  Server selection:              enable
  Through a proxy:               no
```

Required Privilege Level

`superuser, superreader`

test wildfire registration

Description

Performs a test to check the registration status of a WF-500 appliance or Palo Alto Networks firewall to a WildFire server. If the test is successful, the IP address or server name of the WildFire server is displayed. A successful registration is required before a WF-500 appliance or firewall can forward files to the WildFire server.

Syntax

```
test {
  wildfire {
     registration;
    }
  }
```

Options

No additional options.

Sample Output

The following shows a successful output on a firewall that can communicate with a WF-500 appliance. If this is a WF-500 appliance pointing to the Palo Alto Networks WildFire cloud, the server name of one of the cloud servers is displayed in the `select the best server:` field.

```
Test wildfire
        wildfire registration:         successful
        download server list:          successful
        select the best server:        ca-s1.wildfire.paloaltonetworks.com
```

Required Privilege Level

superuser, superreader

WildFire Example

The following example scenario summarizes the full WildFire™ lifecycle. In this example, a sales representative from Palo Alto Networks downloads a new software sales tool that a sales partner uploaded to Dropbox. The sales partner unknowingly uploaded an infected version of the sales tool install file and the sales rep then downloads the infected file.

This example will demonstrate how a Palo Alto Networks firewall in conjunction with WildFire can discover zero-day malware downloaded by an end user, even if the traffic is SSL encrypted. After WildFire identifies the malware a log is sent to the firewall and the firewall alerts the administrator who then contacts the user to eradicate the malware. WildFire then generates a new signature for the malware and firewalls with a Threat Prevention or WildFire subscription automatically downloads the signature to protect against future exposure. Although some file sharing web sites have an antivirus feature that checks files as they are uploaded, they can only protect against known malware.

This example uses a web site that uses SSL encryption. In this case, the firewall has decryption enabled, including the option to forward decrypted content for analysis. To enable decrypted content to be forwarded to the WildFire cloud, see Forward Files for WildFire Analysis.

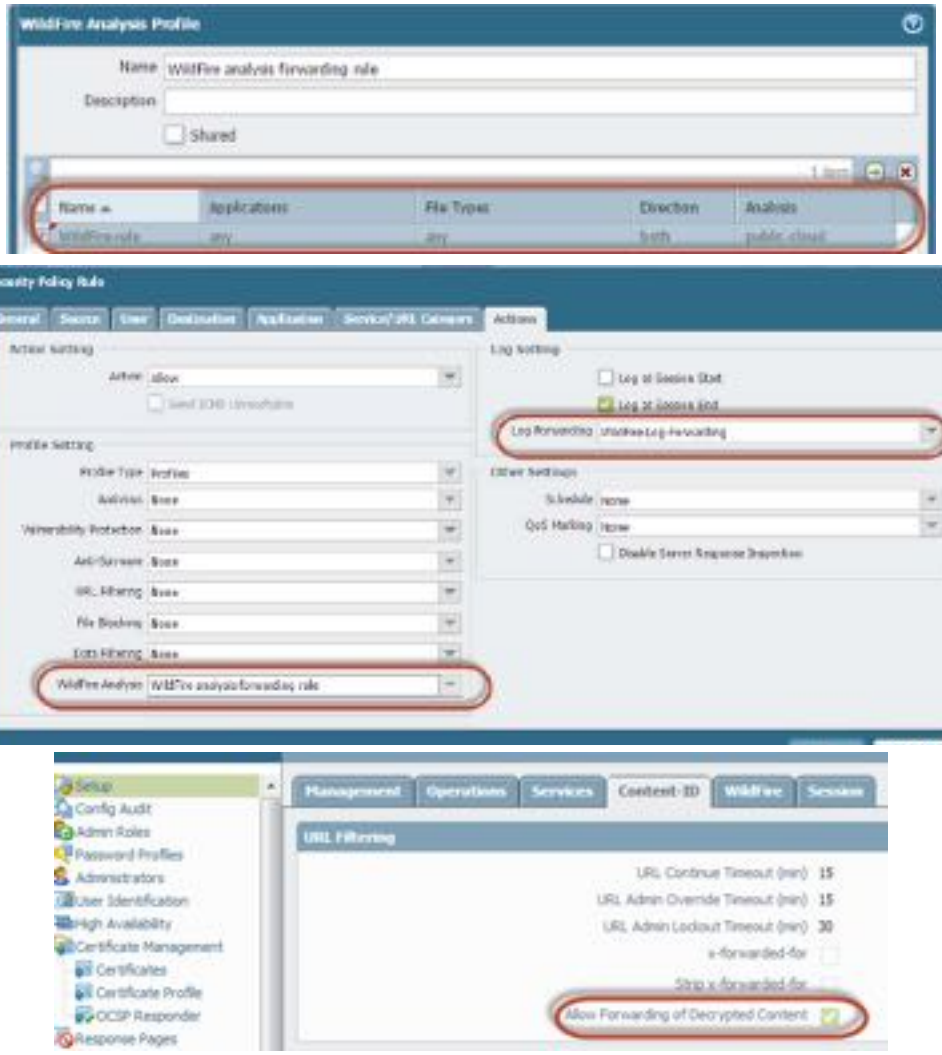
WildFire Workflow	
Step 1	The sales person from the partner company uploads a sales tool file named sales-tool.exe to his Dropbox account and then sends an email to the Palo Alto Networks sales person with a link to the file.
Step 2	The Palo Alto sales person receives the email from the sales partner and clicks the download link, which takes her to the Dropbox site. She then clicks **Download** to save the file to her desktop.





WildFire Workflow (Continued)

Step 3 The firewall that is protecting the Palo Alto sales rep has a WildFire Analysis profile rule attached to a security policy rule that will look for files in any application that is used to download or upload any of the supported file types. The firewall can also be configured to forward the email-link file type, which enables the firewall to extract HTTP/HTTPS links contained in SMTP and POP3 email messages. As soon as the sales rep clicks download, the firewall forwards the sales-toole.exe file to WildFire, where the file is analyzed for zero-day malware. Even though the sales rep is using Dropbox, which is SSL encrypted, the firewall is configured to decrypt traffic, so all traffic can be inspected. The following screen shots show the WildFire Analysis profile rule, the security policy rule configured with the WildFire analysis profile rule attached, and the option to allow forwarding of decrypted content enabled.



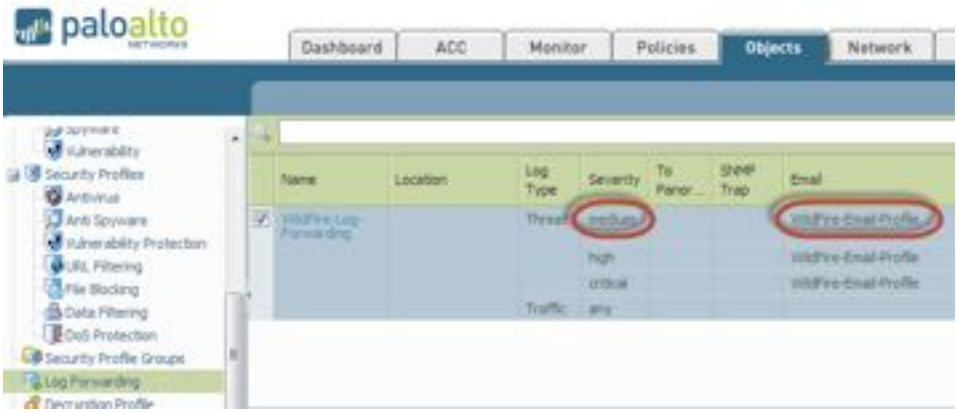
Step 4 At this point, WildFire has received the file and is analyzing it for more than 200 different malicious behaviors. Verify File Forwarding to check that the firewall has correctly forwarded a file or email links for WildFire analysis.

WildFire Workflow (Continued)

Step 5 Within approximately five minutes, WildFire has completed the file analysis and then sends a WildFire log back to the firewall with the analysis results. In this example, the WildFire log shows that the file is malicious.



Step 6 The firewall is configured with a log forwarding profile that will send WildFire alerts to the security administrator when malware is discovered.



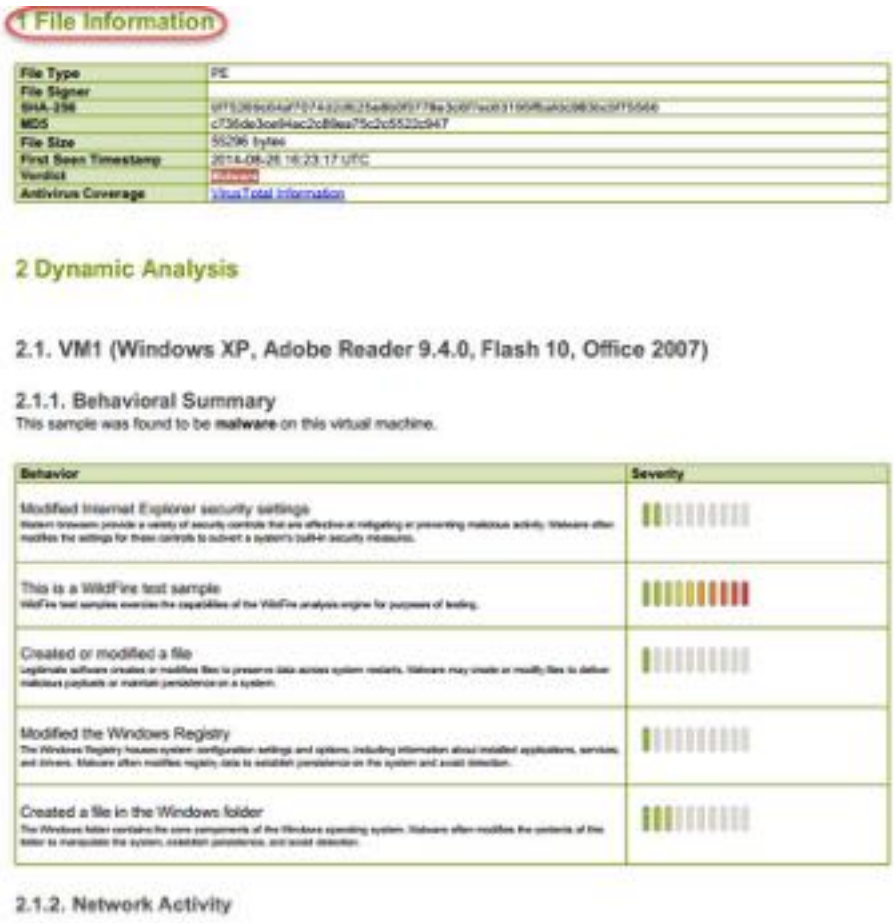
WildFire Workflow (Continued)

Step 7 The security administrator identifies the user by name (if User-ID is configured), or by IP address if User-ID is not enabled. At this point, the administrator can shut down the network or VPN connection that the sales representative is using and will then contact the desktop support group to work with the user to check and clean the system.

By using the WildFire detailed analysis report, the desktop support person can determine if the user system is infected with malware by looking at the files, processes, and registry information detailed in the WildFire analysis report. If the user runs the malware, the support person can attempt to clean the system manually or re-image it.

For details on the WildFire report fields, see WildFire Analysis Reports—Close Up.

Figure: Partial View of the WildFire Analysis Report in PDF



1 File Information

File Type	PE
File Signer	
SHA-256	[redacted]
MD5	c736de3ce9fac2c89ea75c2c5522c947
File Size	55296 bytes
First Seen Timestamp	2014-08-26 16:23:17 UTC
Verdict	[redacted]
Antivirus Coverage	VirusTotal Information

2 Dynamic Analysis

2.1. VM1 (Windows XP, Adobe Reader 9.4.0, Flash 10, Office 2007)

2.1.1. Behavioral Summary

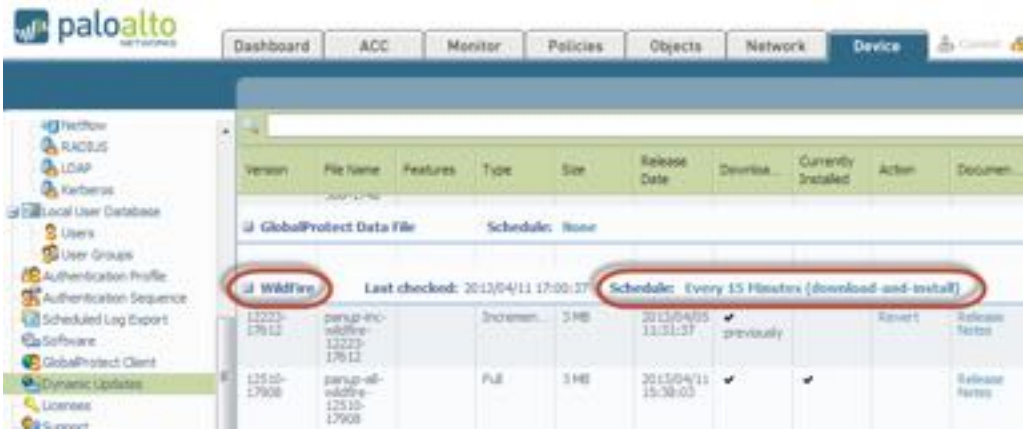
This sample was found to be **malware** on this virtual machine.

Behavior	Severity
Modified Internet Explorer security settings	
This is a WildFire test sample	
Created or modified a file	
Modified the Windows Registry	
Created a file in the Windows folder	

2.1.2. Network Activity

Step 8 Now that the administrator has identified the malware and the user system is being checked, how do you protect from future exposure? Answer: In this example, the administrator set a schedule on the firewall to download and install WildFire signatures every 15 minutes and to download and install Antivirus updates once per day. In less than an hour and a half after the sales rep downloaded the infected file, WildFire identified the zero-day malware, generated a signature, added it to the WildFire update signature database provided by Palo Alto Networks, and the firewall downloaded and installed the new signature. This firewall and any other Palo Alto Networks firewall configured to download WildFire and antivirus signatures is now protected against this newly discovered malware. The following screenshot shows the WildFire update schedule:

WildFire Workflow (Continued)



All of this occurs well before most antivirus vendors are even aware of the zero-day malware. In this example, within a very short period of time, the malware is no longer considered zero-day because Palo Alto Networks has already discovered it and has provided protection to customers to prevent future exposure.

WILDFIRE

Protection from Targeted and Unknown Attacks

WildFire™ cloud-based malware-analysis environment is an advanced threat intelligence service that identifies unknown malware, zero-day exploits, and advanced persistent threats (APTs) through static and dynamic analysis in a scalable, virtual environment. Once deployed, WildFire automatically disseminates updated protections in near-real time to immediately halt threats from spreading – without manual intervention. This closed-loop, automated process gives organizations the assurance that their networks, endpoints and cloud are armed with the absolute latest threat intelligence at all times.

Benefits

- Our Next-Generation Security Platform provides up-to-date protection throughout your organization to reduce the attack surface across multiple attack vectors.
- WildFire identifies unknown malware and zero-day exploits using advanced static and dynamic analysis techniques against multiple OS and application versions.
- WildFire automatically generates malware, URL and DNS signatures and distributes them within minutes to all global, WildFire-subscribed Palo Alto Networks platforms.
- Information about indicators of compromise (IOCs) from WildFire analysis reports is used by our Next- Generation Security Platform and our technology partners to identify infected hosts and prevent secondary downloads.

Advanced cyberattacks are employing stealthy and persistent methods to evade traditional security measures. Skilled adversaries demand that modern security teams re-evaluate their prevention tactics to better address the volume and sophistication of today's attacks. Purpose-built for high fidelity hardware emulation, WildFire analyzes suspicious samples as they execute. When new threats emerge, Palo Alto Networks® Next-Generation Security Platform automatically routes suspicious files and URLs to WildFire for deep analysis.

WildFire inspects millions of samples per week from its global network of customers and threat intelligence partners looking for new forms of previously unknown malware, exploits, malicious domains and outbound command-and-control activity. The cloud-based service creates new protections that are capable of blocking targeted and unknown malware, exploits, and outbound C2 activity by observing their actual "behavior," rather than relying on pre-existing signatures. The protections are shared globally in minutes.

Next-Generation Security Platform

WildFire is built on our industry-leading Next-Generation Security Platform, benefiting from full visibility into all network traffic, including stealthy attempts to evade detection, such as the use of non-standard ports or SSL encryption. Known threats are proactively blocked with our Next-Generation Firewall, Threat Prevention, URL Filtering, Traps and Aperture, providing baseline defenses against known exploits, malware, malicious URLs and command-and-control (C2) activity. Unknown files, and email links are forwarded and analyzed by WildFire in a scalable sandbox environment, where new threats are identified and protections are automatically developed and delivered to the security elements in your organization in the form of signatures and verdict updates. The result is a unique, closed-loop approach to preventing cyberthreats that includes: positive security controls to reduce the attack surface; inspection of all traffic, ports and protocols

to block all known threats; rapid detection of unknown threats by observing the actions of malware in a cloud-based execution environment; and automatic deployment of new protections back to the frontline to ensure threats are known to all and blocked across the attack lifecycle.

Behavior-Based Cyberthreat Discovery

To find unknown malware and exploits, WildFire executes suspicious content in the Windows® XP, Windows 7, Android® and Mac® OS X® operating systems, with full visibility into common file types, including: EXE, DLL, ZIP, PDF, as well as Microsoft® Office documents, Java® files, Android APKs, Adobe® Flash® applets, and webpages, including high-risk, embedded content, such as Java and Adobe Flash files and images.

WildFire identifies hundreds of potentially malicious behaviors to uncover the true nature of malicious files based on their actions, including:

- **Changes made to host:** WildFire observes all processes for modifications to the host, including file and registry activity, code injection, memory heap spray (exploit) detection, addition of auto-run programs, mutexes, Windows services, and other suspicious activities.
- **Suspicious network traffic:** WildFire performs analysis of all network activity produced by the suspicious file, including backdoor creation, downloading of next-stage malware, visiting low-reputation domains, network reconnaissance, and much more.
- **Anti-analysis detection:** WildFire monitors techniques used by advanced malware that is designed to avoid VM-based analysis, such as debugger detection, hypervisor detection, code injection into trusted processes, disabling of host-based security features, and more.

Natively integrated with our Next-Generation Security Platform, which is capable of classifying all traffic across hundreds of applications, WildFire uniquely applies this behavioral analysis to web traffic, email protocols (SMTP, IMAP, POP) and FTP, regardless of ports or encryption.

WildFire Architecture

To support static and dynamic malware analysis at scale, across the organization, WildFire is built on a cloud-based architecture that can be leveraged by your existing Palo Alto Networks Next-Generation Security Platform with no additional hardware. Where regulatory or privacy requirements prevent the use of a cloud infrastructure, a local hardware solution can be deployed on premises using the WF-500 appliance. You can leverage both the cloud and local versions of WildFire within the same environment. This hybrid solution enables you to define which files types are sent to the cloud and which are sent to the local appliance, based on data sensitivity. Analysis results in the form of new protections that are shared globally.

Threat Prevention with Global Intelligence Sharing

When an unknown threat is discovered, WildFire automatically generates protections to block it across the cyberattack lifecycle, sharing these updates with all global subscribers in as little as five minutes. These quick updates are able to stop rapidly spreading malware; since these updates are payload based, they can block proliferation of future variants without any additional action or analysis.

In addition to protecting organizations from malicious and exploitive files and links, WildFire looks deeply into malicious outbound communication, disrupting command-and-control activity with anti-C2 signatures and DNS-based callback signatures. The information is also fed into URL Filtering with PAN-DB, where newly discovered malicious URLs are automatically blocked. This correlation of threat data and automated protections are key to identifying and blocking ongoing intrusion attempts and future attacks on your organization.

Integrated Logging, Reporting and Forensics

WildFire users receive integrated logs, analysis and visibility into WildFire events through the management interface, Panorama™, AutoFocus™ or the WildFire portal, enabling teams to quickly investigate and correlate events observed in their networks. This allows security staff to rapidly locate the data needed for timely investigations and incident response. Host-based and network-based indicators of compromise become actionable through log analysis and custom signatures.

To aid security and IR staff in discovering infected hosts, WildFire also provides:

- Detailed analysis of every malicious file sent to WildFire across multiple operating system environments, including both host-based and network-based activity.
- Session data associated with the delivery of the malicious file, including source, destination, application, User-ID™, URL, etc.
- Access to the original malware sample for reverse-engineering and full PCAPs of dynamic analysis sessions.
- An open API for integration with best-in-class SIEM tools, such as the Palo Alto Networks application for Splunk, and leading endpoint agents. This analysis provides a wealth of indicators of compromise (IOCs) that can be applied across the attack lifecycle.
- Native integration with Traps™ advanced endpoint protection and Aperture; advanced SaaS protection.
- Access to the actionable intelligence and global context provided by Palo Alto Networks AutoFocus threat intelligence service.
- Natively integrated with the correlation engine in our next- generation firewalls.

Maintaining the Privacy of Your Files

WildFire leverages a public cloud environment managed directly by Palo Alto Networks. All suspicious files are securely transferred between the Next-Generation Security Platform and the WildFire data center over encrypted connections, signed on both sides by Palo Alto Networks. Any files that are found to be benign are destroyed, while malware files are archived for further analysis.

WildFire Requirements:

- PAN-OS® 4.1+
- DF, Java, Office, and APK analysis require PAN-OS 6.0+
- Adobe Flash and webpage analysis require PAN-OS 6.1+

Licensing Information:

Basic WildFire functionality is available as a standard feature on all platforms running PAN-OS 4.1 or greater:

- Windows XP and Windows 7 image analysis
- EXE and DLL file types, including compressed (zip) and encrypted (SSL) content
- Automatic submission of suspicious files
- Automatic protections are delivered with regular threat prevention content updates (Threat Prevention license is required) every 24-48 hours

The WildFire subscription adds near-real time protection from advanced threats, including these additional features:

- Automatic WildFire signature updates every 15 minutes for all new malware detected anywhere in the world
- Enhanced file-type support, including: PE files (EXE, DLL, and others), all Microsoft Office file types, PDF files, and Java applets (JAR and CLASS)

WF-500

The WF-500 is an optional hardware appliance to support customers who choose to deploy WildFire as a private cloud for additional data privacy. The WF-500 is sized to accommodate most mid-range to large-scale networks, with the option of deploying additional appliances as traffic volumes increase or for networks that require geographic distribution.

The WF-500 can be deployed in a hybrid mode with the global WildFire services.

WF-500 Specifications

Processor	Memory	System Disk
Dual 6-Core Intel® Processor with Hyper -Threading Technology	128 GB RA 1	20GB SSD

Hardware Specifications

I/O Storage	Capacity	Power Supply
4x10/100/1,000 DB9 Console serial port, USB HDD for 2 TB of RAID storage	2TB RAID1: 4 x 1TB RAID Certified	Dual 920W power supplies in hot swap, redundant configuration
Max Power Consumption	**Rack Mountable (Dimensions)**	
510 Watts	2U, 19" standard rack (3.5"H x 21"D x 17.5"W)	
Max Btu/Hr	**Input Voltage (Input Frequency)**	
1,740 BTU/hr	100-240VAC (50-60Hz)	
Max Current Consumption	**Safety**	
11A @ 100VAC	UL, CSA, CB	
EMI	**Environment**	
FCC Class A, CE Class A, VCCI Class A	Operating Temperature: 32 to 95 F, 5 to 35 C Non-operating Temperature: -4 to 158 F, -40 to 65 C	

To view additional information about the WF-500 security features and associated capacities, please visit www.paloaltonetworks.com/products.

4401 Great America Parkway
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

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. **pan-ds-wildfire-040416**