

PAN-OS[®] Administrator's Guide

Version 9.0

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

© 2019-2020 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

August 3, 2020

Table of Contents

Getting Started.....	17
Integrate the Firewall into Your Management Network.....	19
Determine Your Management Strategy.....	19
Perform Initial Configuration.....	20
Set Up Network Access for External Services.....	23
Register the Firewall.....	28
Create a New Support Account and Register a Firewall.....	28
Register a Firewall.....	30
(Optional) Perform Day 1 Configuration.....	32
Segment Your Network Using Interfaces and Zones.....	35
Network Segmentation for a Reduced Attack Surface.....	35
Configure Interfaces and Zones.....	35
Set Up a Basic Security Policy.....	38
Assess Network Traffic.....	42
Enable Free WildFire Forwarding.....	44
Best Practices for Completing the Firewall Deployment.....	46
Best Practices for Securing Administrative Access.....	47
Isolate the Management Network.....	47
Use Service Routes to Access External Services.....	48
Restrict Access to the Management Interface.....	48
Manage Administrator Access.....	50
Create Strong Administrator Passwords.....	51
Scan All Traffic Destined for the Management Interface.....	51
Replace the Certificate for Inbound Management Traffic.....	52
Keep Content and Software Updates Current.....	53
 Subscriptions.....	 55
Subscriptions You Can Use With the Firewall.....	57
Activate Subscription Licenses.....	59
What Happens When Licenses Expire?.....	60
Enhanced Application Logs for Palo Alto Networks Cloud Services.....	62
 Software and Content Updates.....	 65
PAN-OS Software Updates.....	67
Dynamic Content Updates.....	68
Install Content Updates.....	70
Applications and Threats Content Updates.....	73
Deploy Applications and Threats Content Updates.....	73
Tips for Content Updates.....	74
Best Practices for Applications and Threats Content Updates.....	76
Best Practices for Content Updates—Mission-Critical.....	76
Best Practices for Content Updates—Security-First.....	79
Content Delivery Network Infrastructure.....	82
 Firewall Administration.....	 83
Management Interfaces.....	85
Use the Web Interface.....	86

Launch the Web Interface.....	86
Configure Banners, Message of the Day, and Logos.....	86
Use the Administrator Login Activity Indicators to Detect Account Misuse.....	88
Manage and Monitor Administrative Tasks.....	90
Commit, Validate, and Preview Firewall Configuration Changes.....	90
Export Configuration Table Data.....	92
Use Global Find to Search the Firewall or Panorama Management Server.....	93
Manage Locks for Restricting Configuration Changes.....	94
Manage Configuration Backups.....	96
Save and Export Firewall Configurations.....	96
Revert Firewall Configuration Changes.....	97
Manage Firewall Administrators.....	100
Administrative Role Types.....	100
Configure an Admin Role Profile.....	101
Administrative Authentication.....	101
Configure Administrative Accounts and Authentication.....	102
Reference: Web Interface Administrator Access.....	109
Web Interface Access Privileges.....	109
Panorama Web Interface Access Privileges.....	165
Reference: Port Number Usage.....	170
Ports Used for Management Functions.....	170
Ports Used for HA.....	171
Ports Used for Panorama.....	172
Ports Used for GlobalProtect.....	173
Ports Used for User-ID.....	173
Reset the Firewall to Factory Default Settings.....	176
Bootstrap the Firewall.....	177
USB Flash Drive Support.....	177
Sample init-cfg.txt Files.....	178
Prepare a USB Flash Drive for Bootstrapping a Firewall.....	179
Bootstrap a Firewall Using a USB Flash Drive.....	182

Authentication..... 183

Authentication Types.....	185
External Authentication Services.....	185
Multi-Factor Authentication.....	185
SAML.....	186
Kerberos.....	187
TACACS+.....	188
RADIUS.....	188
LDAP.....	190
Local Authentication.....	190
Plan Your Authentication Deployment.....	191
Configure Multi-Factor Authentication.....	192
Configure MFA Between RSA SecurID and the Firewall.....	195
Configure MFA Between Okta and the Firewall.....	201
Configure MFA Between Duo and the Firewall.....	210
Configure SAML Authentication.....	218
Configure Kerberos Single Sign-On.....	222
Configure Kerberos Server Authentication.....	224
Configure TACACS+ Authentication.....	225
Configure RADIUS Authentication.....	227
Configure LDAP Authentication.....	230
Connection Timeouts for Authentication Servers.....	232

Guidelines for Setting Authentication Server Timeouts.....	232
Modify the PAN-OS Web Server Timeout.....	233
Modify the Captive Portal Session Timeout.....	233
Configure Local Database Authentication.....	234
Configure an Authentication Profile and Sequence.....	235
Test Authentication Server Connectivity.....	238
Authentication Policy.....	240
Authentication Timestamps.....	240
Configure Authentication Policy.....	241
Troubleshoot Authentication Issues.....	244

Certificate Management..... 247

Keys and Certificates.....	249
Default Trusted Certificate Authorities (CAs).....	252
Certificate Revocation.....	253
Certificate Revocation List (CRL).....	253
Online Certificate Status Protocol (OCSP).....	254
Certificate Deployment.....	255
Set Up Verification for Certificate Revocation Status.....	256
Configure an OCSP Responder.....	256
Configure Revocation Status Verification of Certificates.....	257
Configure Revocation Status Verification of Certificates Used for SSL/TLS	
Decryption.....	257
Configure the Master Key.....	259
Obtain Certificates.....	261
Create a Self-Signed Root CA Certificate.....	261
Generate a Certificate.....	261
Import a Certificate and Private Key.....	263
Obtain a Certificate from an External CA.....	264
Deploy Certificates Using SCEP.....	265
Export a Certificate and Private Key.....	268
Configure a Certificate Profile.....	269
Configure an SSL/TLS Service Profile.....	271
Replace the Certificate for Inbound Management Traffic.....	272
Configure the Key Size for SSL Forward Proxy Server Certificates.....	273
Revoke and Renew Certificates.....	274
Revoke a Certificate.....	274
Renew a Certificate.....	274
Secure Keys with a Hardware Security Module.....	275
Set Up Connectivity with an HSM.....	275
Encrypt a Master Key Using an HSM.....	280
Store Private Keys on an HSM.....	281
Manage the HSM Deployment.....	282

High Availability..... 283

HA Overview.....	285
HA Concepts.....	286
HA Modes.....	286
HA Links and Backup Links.....	287
Device Priority and Preemption.....	291
Failover.....	291
LACP and LLDP Pre-Negotiation for Active/Passive HA.....	292
Floating IP Address and Virtual MAC Address.....	292

ARP Load-Sharing.....	293
Route-Based Redundancy.....	295
HA Timers.....	296
Session Owner.....	298
Session Setup.....	299
NAT in Active/Active HA Mode.....	301
ECMP in Active/Active HA Mode.....	301
Set Up Active/Passive HA.....	302
Prerequisites for Active/Passive HA.....	302
Configuration Guidelines for Active/Passive HA.....	302
Configure Active/Passive HA.....	305
Define HA Failover Conditions.....	309
Verify Failover.....	310
Set Up Active/Active HA.....	311
Prerequisites for Active/Active HA.....	311
Configure Active/Active HA.....	312
Determine Your Active/Active Use Case.....	316
Refresh HA1 SSH Keys and Configure Key Options.....	330
HA Firewall States.....	335
Reference: HA Synchronization.....	337
What Settings Don't Sync in Active/Passive HA?.....	337
What Settings Don't Sync in Active/Active HA?.....	339
Synchronization of System Runtime Information.....	342

Monitoring.....345

Use the Dashboard.....	347
Use the Application Command Center.....	349
ACC—First Look.....	349
ACC Tabs.....	351
ACC Widgets.....	352
Widget Descriptions.....	354
ACC Filters.....	358
Interact with the ACC.....	359
Use Case: ACC—Path of Information Discovery.....	362
Use the App Scope Reports.....	368
Summary Report.....	368
Change Monitor Report.....	369
Threat Monitor Report.....	370
Threat Map Report.....	371
Network Monitor Report.....	371
Traffic Map Report.....	372
Use the Automated Correlation Engine.....	374
Automated Correlation Engine Concepts.....	374
View the Correlated Objects.....	375
Interpret Correlated Events.....	375
Use the Compromised Hosts Widget in the ACC.....	377
Take Packet Captures.....	378
Types of Packet Captures.....	378
Disable Hardware Offload.....	378
Take a Custom Packet Capture.....	379
Take a Threat Packet Capture.....	383
Take an Application Packet Capture.....	384
Take a Packet Capture on the Management Interface.....	387
Monitor Applications and Threats.....	389

View and Manage Logs.....	390
Log Types and Severity Levels.....	390
View Logs.....	395
Filter Logs.....	396
Export Logs.....	397
Configure Log Storage Quotas and Expiration Periods.....	397
Schedule Log Exports to an SCP or FTP Server.....	398
Monitor Block List.....	399
View and Manage Reports.....	400
Report Types.....	400
View Reports.....	400
Configure the Expiration Period and Run Time for Reports.....	401
Disable Predefined Reports.....	401
Custom Reports.....	402
Generate Custom Reports.....	404
Generate Botnet Reports.....	406
Generate the SaaS Application Usage Report.....	407
Manage PDF Summary Reports.....	410
Generate User/Group Activity Reports.....	412
Manage Report Groups.....	413
Schedule Reports for Email Delivery.....	414
Manage Report Storage Capacity.....	415
View Policy Rule Usage.....	416
Use External Services for Monitoring.....	420
Configure Log Forwarding.....	421
Configure Email Alerts.....	424
Use Syslog for Monitoring.....	425
Configure Syslog Monitoring.....	425
Syslog Field Descriptions.....	427
SNMP Monitoring and Traps.....	465
SNMP Support.....	465
Use an SNMP Manager to Explore MIBs and Objects.....	466
Enable SNMP Services for Firewall-Secured Network Elements.....	468
Monitor Statistics Using SNMP.....	469
Forward Traps to an SNMP Manager.....	470
Supported MIBs.....	472
Forward Logs to an HTTP/S Destination.....	480
NetFlow Monitoring.....	484
Configure NetFlow Exports.....	484
NetFlow Templates.....	485
Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors.....	491

User-ID.....493

User-ID Overview.....	495
User-ID Concepts.....	497
Group Mapping.....	497
User Mapping.....	497
Enable User-ID.....	501
Map Users to Groups.....	504
Map IP Addresses to Users.....	509
Create a Dedicated Service Account for the User-ID Agent.....	509
Configure User Mapping Using the Windows User-ID Agent.....	525
Configure User Mapping Using the PAN-OS Integrated User-ID Agent.....	535
Configure Server Monitoring Using WinRM.....	538

Configure User-ID to Monitor Syslog Senders for User Mapping.....	544
Map IP Addresses to Usernames Using Captive Portal.....	552
Configure User Mapping for Terminal Server Users.....	557
Send User Mappings to User-ID Using the XML API.....	565
Enable User- and Group-Based Policy.....	566
Enable Policy for Users with Multiple Accounts.....	567
Verify the User-ID Configuration.....	569
Deploy User-ID in a Large-Scale Network.....	571
Deploy User-ID for Numerous Mapping Information Sources.....	571
Redistribute User Mappings and Authentication Timestamps.....	575
Share User-ID Mappings Across Virtual Systems.....	578

App-ID..... 581

App-ID Overview.....	583
App-ID and HTTP/2 Inspection.....	584
Manage Custom or Unknown Applications.....	586
Manage New and Modified App-IDs.....	587
Workflow to Best Incorporate New and Modified App-IDs.....	587
See the New and Modified App-IDs in a Content Release.....	588
See How New and Modified App-IDs Impact Your Security Policy.....	589
Ensure Critical New App-IDs are Allowed.....	590
Monitor New App-IDs.....	591
Disable and Enable App-IDs.....	592
Use Application Objects in Policy.....	594
Create an Application Group.....	594
Create an Application Filter.....	594
Create a Custom Application.....	595
Safely Enable Applications on Default Ports.....	599
Applications with Implicit Support.....	601
Security Policy Rule Optimization.....	604
Policy Optimizer Concepts.....	605
Migrate Port-Based to App-ID Based Security Policy Rules.....	609
Rule Cloning Migration Use Case: Web Browsing and SSL Traffic.....	615
Add Applications to an Existing Rule.....	619
Identify Security Policy Rules with Unused Applications.....	622
High Availability for Application Usage Statistics.....	626
How to Disable Policy Optimizer.....	627
Application Level Gateways.....	628
Disable the SIP Application-level Gateway (ALG).....	630
Use HTTP Headers to Manage SaaS Application Access.....	632
Understand SaaS Custom Headers.....	632
Domains used by the Predefined SaaS Application Types.....	634
Create HTTP Header Insertion Entries using Predefined Types.....	634
Create Custom HTTP Header Insertion Entries.....	635
Maintain Custom Timeouts for Data Center Applications.....	637

Threat Prevention..... 639

Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions.....	641
Set Up Antivirus, Anti-Spyware, and Vulnerability Protection.....	647
DNS Security.....	650
About DNS Security.....	650
Domain Generation Algorithm (DGA) Detection.....	651
DNS Tunneling Detection.....	651

Cloud-Delivered DNS Signatures and Protections.....	652
Enable DNS Security.....	652
Use DNS Queries to Identify Infected Hosts on the Network.....	656
How DNS Sinkholing Works.....	656
Configure DNS Sinkholing.....	657
Configure DNS Sinkholing for a List of Custom Domains.....	658
Configure the Sinkhole IP Address to a Local Server on Your Network.....	660
See Infected Hosts that Attempted to Connect to a Malicious Domain.....	662
Set Up Data Filtering.....	665
Create a Data Filtering Profile.....	665
Predefined Data Filtering Patterns.....	667
Set Up File Blocking.....	670
Prevent Brute Force Attacks.....	673
Customize the Action and Trigger Conditions for a Brute Force Signature.....	674
Enable Evasion Signatures.....	677
Prevent Credential Phishing.....	678
Methods to Check for Corporate Credential Submissions.....	678
Configure Credential Detection with the Windows-based User-ID Agent.....	680
Set Up Credential Phishing Prevention.....	682
Monitor Blocked IP Addresses.....	685
Threat Signature Categories.....	687
Create Threat Exceptions.....	694
Custom Signatures.....	696
Monitor and Get Threat Reports.....	697
Monitor Activity and Create Custom Reports Based on Threat Categories.....	697
Learn More About Threat Signatures.....	698
AutoFocus Threat Intelligence for Network Traffic.....	701
Share Threat Intelligence with Palo Alto Networks.....	706
What Telemetry Data Does the Firewall Collect?.....	706
Passive DNS Monitoring.....	707
Enable Telemetry.....	708
Threat Prevention Resources.....	710

Decryption.....711

Decryption Overview.....	713
Decryption Concepts.....	714
Keys and Certificates for Decryption Policies.....	714
SSL Forward Proxy.....	716
SSL Forward Proxy Decryption Profile.....	717
SSL Inbound Inspection.....	719
SSL Inbound Inspection Decryption Profile.....	720
SSL Protocol Settings Decryption Profile.....	722
SSH Proxy.....	723
SSH Proxy Decryption Profile.....	724
Decryption Profile for No Decryption.....	725
SSL Decryption for Elliptical Curve Cryptography (ECC) Certificates.....	726
Perfect Forward Secrecy (PFS) Support for SSL Decryption.....	726
SSL Decryption and Subject Alternative Names (SANs).....	727
High Availability Support for Decrypted Sessions.....	728
Decryption Mirroring.....	729
Prepare to Deploy Decryption.....	730
Work with Stakeholders to Develop a Decryption Deployment Strategy.....	730
Develop a PKI Rollout Plan.....	732
Size the Decryption Firewall Deployment.....	733

Plan a Staged, Prioritized Deployment.....	734
Define Traffic to Decrypt.....	736
Create a Decryption Profile.....	736
Create a Decryption Policy Rule.....	738
Configure SSL Forward Proxy.....	741
Configure SSL Inbound Inspection.....	745
Configure SSH Proxy.....	747
Configure Server Certificate Verification for Undecrypted Traffic.....	748
Decryption Exclusions.....	749
Palo Alto Networks Predefined Decryption Exclusions.....	749
Exclude a Server from Decryption for Technical Reasons.....	750
Create a Policy-Based Decryption Exclusion.....	752
Enable Users to Opt Out of SSL Decryption.....	754
Temporarily Disable SSL Decryption.....	756
Configure Decryption Port Mirroring.....	757
Verify Decryption.....	760
Decryption Broker.....	763
How Decryption Broker Works.....	763
Decryption Broker Concepts.....	764
Layer 3 Security Chain Guidelines.....	770
Configure Decryption Broker with One or More Layer 3 Security Chain.....	771
Transparent Bridge Security Chain Guidelines.....	773
Configure Decryption Broker with a Single Transparent Bridge Security Chain.....	773
Configure Decryption Broker with Multiple Transparent Bridge Security Chains.....	775
Activate Free Licenses for Decryption Features.....	777

URL Filtering..... 779

About URL Filtering.....	781
How URL Filtering Works.....	782
URL Filtering Vendors.....	784
Enable BrightCloud URL Filtering.....	784
Enable PAN-DB URL Filtering.....	786
URL Filtering Use Cases.....	788
URL Categories.....	791
Security-Focused URL Categories.....	791
Malicious URL Categories.....	792
Policy Actions You Can Take Based on URL Categories.....	793
Plan Your URL Filtering Deployment.....	796
URL Filtering Best Practices.....	798
Configure URL Filtering.....	801
Monitor Web Activity.....	804
Monitor Web Activity of Network Users.....	804
View the User Activity Report.....	806
Configure Custom URL Filtering Reports.....	807
Log Only the Page a User Visits.....	810
Create a Custom URL Category.....	811
URL Category Exceptions.....	813
Basic Guidelines For URL Category Exception Lists.....	813
Wildcard Guidelines for URL Category Exception Lists.....	813
URL Category Exception List—Wildcard Examples.....	814
Use an External Dynamic List in a URL Filtering Profile.....	816
Allow Password Access to Certain Sites.....	818
Safe Search Enforcement.....	820
Safe Search Settings for Search Providers.....	820

Block Search Results when Strict Safe Search is not Enabled.....	822
Transparently Enable Safe Search for Users.....	824
URL Filtering Response Pages.....	828
Customize the URL Filtering Response Pages.....	832
HTTP Header Logging.....	833
Request to Change the Category for a URL.....	834
Make a Change Request Online.....	834
Make a Bulk Change Request.....	835
Make a Change Request from the Firewall.....	836
Troubleshoot URL Filtering.....	837
Problems Activating PAN-DB.....	837
PAN-DB Cloud Connectivity Issues.....	837
URLs Classified as Not-Resolved.....	838
Incorrect Categorization.....	839
PAN-DB Private Cloud.....	841
M-600 Appliance for PAN-DB Private Cloud.....	841
Set Up the PAN-DB Private Cloud.....	842

Quality of Service.....851

QoS Overview.....	853
QoS Concepts.....	855
QoS for Applications and Users.....	855
QoS Policy.....	855
QoS Profile.....	855
QoS Classes.....	856
QoS Priority Queuing.....	856
QoS Bandwidth Management.....	857
QoS Egress Interface.....	857
QoS for Clear Text and Tunneled Traffic.....	858
Configure QoS.....	859
Configure QoS for a Virtual System.....	864
Enforce QoS Based on DSCP Classification.....	869
QoS Use Cases.....	871
Use Case: QoS for a Single User.....	871
Use Case: QoS for Voice and Video Applications.....	873

VPNs.....877

VPN Deployments.....	879
Site-to-Site VPN Overview.....	880
Site-to-Site VPN Concepts.....	881
IKE Gateway.....	881
Tunnel Interface.....	881
Tunnel Monitoring.....	882
Internet Key Exchange (IKE) for VPN.....	882
IKEv2.....	884
Set Up Site-to-Site VPN.....	888
Set Up an IKE Gateway.....	888
Define Cryptographic Profiles.....	893
Set Up an IPSec Tunnel.....	895
Set Up Tunnel Monitoring.....	898
Enable/Disable, Refresh or Restart an IKE Gateway or IPSec Tunnel.....	899
Test VPN Connectivity.....	900
Interpret VPN Error Messages.....	901

Site-to-Site VPN Quick Configs.....	903
Site-to-Site VPN with Static Routing.....	903
Site-to-Site VPN with OSPF.....	906
Site-to-Site VPN with Static and Dynamic Routing.....	909
Large Scale VPN (LSVPN).....	915
LSVPN Overview.....	917
Create Interfaces and Zones for the LSVPN.....	918
Enable SSL Between GlobalProtect LSVPN Components.....	920
About Certificate Deployment.....	920
Deploy Server Certificates to the GlobalProtect LSVPN Components.....	920
Deploy Client Certificates to the GlobalProtect Satellites Using SCEP.....	922
Configure the Portal to Authenticate Satellites.....	925
Configure GlobalProtect Gateways for LSVPN.....	927
Configure the GlobalProtect Portal for LSVPN.....	930
GlobalProtect Portal for LSVPN Prerequisite Tasks.....	930
Configure the Portal.....	930
Define the Satellite Configurations.....	931
Prepare the Satellite to Join the LSVPN.....	934
Verify the LSVPN Configuration.....	936
LSVPN Quick Configs.....	937
Basic LSVPN Configuration with Static Routing.....	937
Advanced LSVPN Configuration with Dynamic Routing.....	939
Advanced LSVPN Configuration with iBGP.....	941
Networking.....	947
Configure Interfaces.....	949
Tap Interfaces.....	949
Virtual Wire Interfaces.....	950
Layer 2 Interfaces.....	957
Layer 3 Interfaces.....	962
Configure an Aggregate Interface Group.....	971
Use Interface Management Profiles to Restrict Access.....	973
Virtual Routers.....	975
Service Routes.....	977
Static Routes.....	979
Static Route Overview.....	979
Static Route Removal Based on Path Monitoring.....	979
Configure a Static Route.....	981
Configure Path Monitoring for a Static Route.....	982
RIP.....	985
OSPF.....	987
OSPF Concepts.....	987
Configure OSPF.....	989
Configure OSPFv3.....	991
Configure OSPF Graceful Restart.....	993
Confirm OSPF Operation.....	994
BGP.....	996
BGP Overview.....	996
MP-BGP.....	996
Configure BGP.....	998
Configure a BGP Peer with MP-BGP for IPv4 or IPv6 Unicast.....	1002
Configure a BGP Peer with MP-BGP for IPv4 Multicast.....	1004

BGP Confederations.....	1005
IP Multicast.....	1011
IGMP.....	1011
PIM.....	1012
Configure IP Multicast.....	1017
View IP Multicast Information.....	1022
Route Redistribution.....	1024
GRE Tunnels.....	1027
GRE Tunnel Overview.....	1027
Create a GRE Tunnel.....	1028
DHCP.....	1031
DHCP Overview.....	1031
Firewall as a DHCP Server and Client.....	1032
DHCP Messages.....	1032
DHCP Addressing.....	1033
DHCP Options.....	1034
Configure an Interface as a DHCP Server.....	1036
Configure an Interface as a DHCP Client.....	1039
Configure the Management Interface as a DHCP Client.....	1041
Configure an Interface as a DHCP Relay Agent.....	1043
Monitor and Troubleshoot DHCP.....	1043
DNS.....	1046
DNS Overview.....	1046
DNS Proxy Object.....	1047
DNS Server Profile.....	1048
Multi-Tenant DNS Deployments.....	1048
Configure a DNS Proxy Object.....	1049
Configure a DNS Server Profile.....	1051
Use Case 1: Firewall Requires DNS Resolution.....	1052
Use Case 2: ISP Tenant Uses DNS Proxy to Handle DNS Resolution for Security Policies, Reporting, and Services within its Virtual System.....	1053
Use Case 3: Firewall Acts as DNS Proxy Between Client and Server.....	1055
DNS Proxy Rule and FQDN Matching.....	1056
Dynamic DNS Overview.....	1060
Configure Dynamic DNS for Firewall Interfaces.....	1062
NAT.....	1064
NAT Policy Rules.....	1064
Source NAT and Destination NAT.....	1066
NAT Rule Capacities.....	1073
Dynamic IP and Port NAT Oversubscription.....	1073
Dataplane NAT Memory Statistics.....	1074
Configure NAT.....	1075
NAT Configuration Examples.....	1083
NPTv6.....	1089
NPTv6 Overview.....	1089
How NPTv6 Works.....	1090
NDP Proxy.....	1092
NPTv6 and NDP Proxy Example.....	1093
Create an NPTv6 Policy.....	1094
NAT64.....	1096
NAT64 Overview.....	1096
IPv4-Embedded IPv6 Address.....	1097
DNS64 Server.....	1097
Path MTU Discovery.....	1097
IPv6-Initiated Communication.....	1098

Configure NAT64 for IPv6-Initiated Communication.....	1099
Configure NAT64 for IPv4-Initiated Communication.....	1101
Configure NAT64 for IPv4-Initiated Communication with Port Translation.....	1103
ECMP.....	1106
ECMP Load-Balancing Algorithms.....	1106
ECMP Model, Interface, and IP Routing Support.....	1107
Configure ECMP on a Virtual Router.....	1107
Enable ECMP for Multiple BGP Autonomous Systems.....	1109
Verify ECMP.....	1110
LLDP.....	1111
LLDP Overview.....	1111
Supported TLVs in LLDP.....	1112
LLDP Syslog Messages and SNMP Traps.....	1113
Configure LLDP.....	1113
View LLDP Settings and Status.....	1115
Clear LLDP Statistics.....	1116
BFD.....	1117
BFD Overview.....	1117
Configure BFD.....	1119
Reference: BFD Details.....	1124
Session Settings and Timeouts.....	1128
Transport Layer Sessions.....	1128
TCP.....	1128
UDP.....	1132
ICMP.....	1133
Control Specific ICMP or ICMPv6 Types and Codes.....	1134
Configure Session Timeouts.....	1135
Configure Session Settings.....	1137
Session Distribution Policies.....	1139
Prevent TCP Split Handshake Session Establishment.....	1143
Tunnel Content Inspection.....	1144
Tunnel Content Inspection Overview.....	1144
Configure Tunnel Content Inspection.....	1147
View Inspected Tunnel Activity.....	1152
View Tunnel Information in Logs.....	1153
Create a Custom Report Based on Tagged Tunnel Traffic.....	1154

Policy..... 1155

Policy Types.....	1157
Security Policy.....	1158
Components of a Security Policy Rule.....	1158
Security Policy Actions.....	1161
Create a Security Policy Rule.....	1161
Policy Objects.....	1164
Security Profiles.....	1166
Create a Security Profile Group.....	1172
Set Up or Override a Default Security Profile Group.....	1173
Track Rules Within a Rulebase.....	1175
Rule Numbers.....	1175
Rule UUIDs.....	1176
Enforce Policy Rule Description, Tag, and Audit Comment.....	1180
Move or Clone a Policy Rule or Object to a Different Virtual System.....	1182
Use an Address Object to Represent IP Addresses.....	1183
Address Objects.....	1183

Create an Address Object.....	1184
Use Tags to Group and Visually Distinguish Objects.....	1186
Create and Apply Tags.....	1186
Modify Tags.....	1187
View Rules by Tag Group.....	1187
Use an External Dynamic List in Policy.....	1190
External Dynamic List.....	1190
Formatting Guidelines for an External Dynamic List.....	1192
Built-in External Dynamic Lists.....	1194
Configure the Firewall to Access an External Dynamic List.....	1195
Retrieve an External Dynamic List from the Web Server.....	1197
View External Dynamic List Entries.....	1198
Exclude Entries from an External Dynamic List.....	1199
Enforce Policy on an External Dynamic List.....	1199
Find External Dynamic Lists That Failed Authentication.....	1201
Disable Authentication for an External Dynamic List.....	1202
Register IP Addresses and Tags Dynamically.....	1204
Monitor Changes in the Virtual Environment.....	1206
Enable VM Monitoring to Track Changes on the Virtual Network.....	1206
Attributes Monitored on Virtual Machines in Cloud Platforms.....	1208
Use Dynamic Address Groups in Policy.....	1212
CLI Commands for Dynamic IP Addresses and Tags.....	1216
Identify Users Connected through a Proxy Server.....	1218
Use XFF Values for Policies and Logging Source Users.....	1218
Use the IP Address in the XFF Header to Troubleshoot Events.....	1219
Policy-Based Forwarding.....	1221
PBF.....	1221
Create a Policy-Based Forwarding Rule.....	1223
Use Case: PBF for Outbound Access with Dual ISPs.....	1225
Test Policy Rules.....	1231

Virtual Systems..... 1233

Virtual Systems Overview.....	1235
Virtual System Components and Segmentation.....	1235
Benefits of Virtual Systems.....	1236
Use Cases for Virtual Systems.....	1236
Platform Support and Licensing for Virtual Systems.....	1236
Administrative Roles for Virtual Systems.....	1237
Shared Objects for Virtual Systems.....	1237
Communication Between Virtual Systems.....	1238
Inter-VSYS Traffic That Must Leave the Firewall.....	1238
Inter-VSYS Traffic That Remains Within the Firewall.....	1238
Inter-VSYS Communication Uses Two Sessions.....	1240
Shared Gateway.....	1241
External Zones and Shared Gateway.....	1241
Networking Considerations for a Shared Gateway.....	1242
Configure Virtual Systems.....	1243
Configure Inter-Virtual System Communication within the Firewall.....	1248
Configure a Shared Gateway.....	1249
Customize Service Routes for a Virtual System.....	1250
Customize Service Routes to Services for Virtual Systems.....	1250
Configure a PA-7000 Series Firewall for Logging Per Virtual System.....	1251
Configure Administrative Access Per Virtual System or Firewall.....	1253
Virtual System Functionality with Other Features.....	1255

Zone Protection and DoS Protection.....	1257
Network Segmentation Using Zones.....	1259
How Do Zones Protect the Network?.....	1260
Zone Defense.....	1261
Zone Defense Tools.....	1261
How Do the Zone Defense Tools Work?.....	1263
Firewall Placement for DoS Protection.....	1263
Baseline CPS Measurements for Setting Flood Thresholds.....	1264
Zone Protection Profiles.....	1265
Packet Buffer Protection.....	1268
DoS Protection Profiles and Policy Rules.....	1270
Configure Zone Protection to Increase Network Security.....	1275
Configure Reconnaissance Protection.....	1275
Configure Packet Based Attack Protection.....	1275
Configure Protocol Protection.....	1277
Configure Packet Buffer Protection.....	1280
DoS Protection Against Flooding of New Sessions.....	1282
Multiple-Session DoS Attack.....	1282
Single-Session DoS Attack.....	1285
Configure DoS Protection Against Flooding of New Sessions.....	1286
End a Single Session DoS Attack.....	1288
Identify Sessions That Use an Excessive Percentage of the Packet Buffer.....	1288
Discard a Session Without a Commit.....	1290
Certifications.....	1293
Enable FIPS and Common Criteria Support.....	1295
Access the Maintenance Recovery Tool (MRT).....	1295
Change the Operational Mode to FIPS-CC Mode.....	1296
FIPS-CC Security Functions.....	1298
Scrub the Swap Memory on Firewalls or Appliances Running in FIPS-CC Mode.....	1299

Getting Started

The following topics provide detailed steps to help you deploy a new Palo Alto Networks next-generation firewall. They provide details for integrating a new firewall into your network and how to set up a basic security policy. For guidance on continuing to deploy the security platform features to address your network security needs, review the Best Practices for Completing the Firewall Deployment.

- > [Integrate the Firewall into Your Management Network](#)
- > [Register the Firewall](#)
- > [Segment Your Network Using Interfaces and Zones](#)
- > [Set Up a Basic Security Policy](#)
- > [Assess Network Traffic](#)
- > [Enable Free WildFire Forwarding](#)
- > [Best Practices for Completing the Firewall Deployment](#)
- > [Best Practices for Securing Administrative Access](#)

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.



Do not enable access to your management interface from the internet or from other untrusted zones inside your enterprise security boundary. This applies whether you use the dedicated management port (MGT) or you configured a data port as your management interface. When integrating your firewall into your management network, follow the [Best Practices for Securing Administrative Access](#) to ensure that you are securing administrative access to your firewalls and other security devices in a way that prevents successful attacks.

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.

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.

STEP 1 | Install your firewall and connect power to it.



If your firewall model has dual power supplies, connect the second power supply for redundancy. Refer to the [hardware reference guide](#) for your model for details.

STEP 2 | Gather the required information from your network administrator.

- IP address for MGT port
- Netmask
- Default gateway
- DNS server address

STEP 3 | 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-220 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**.



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, to access this URL.

STEP 4 | 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 5 | Set a secure password for the admin account.



Starting with PAN-OS 9.0.4, the predefined, default administrator password (admin/admin) must be changed on the first login on a device. The new password must be a minimum of eight characters and include a minimum of one lowercase and one uppercase character, as well as one number or special character.

Be sure to use the [best practices for password strength](#) to ensure a strict password and review the [password complexity settings](#).

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 6 | Configure the MGT interface.

1. Select **Device > Setup > Interfaces** and edit the **Management** interface.
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 Client**. 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**.

STEP 7 | 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 8 | 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**.

STEP 9 | (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 10 | 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 11 | 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 12 | Open an SSH management session to the firewall.

Using a terminal emulation software, such as PuTTY, launch an SSH session to the firewall using the new IP address you assigned to it.

STEP 13 | 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 Subscription Licenses](#).
1. Use update server connectivity test to verify network connectivity to the Palo Alto Networks Update server as shown in the following example:
 1. Select **Device > Troubleshooting**, and select **Update Server Connectivity** from the Select Test drop-down.
 2. **Execute** the update server connectivity test.



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. If your firewall is not yet registered, the update server returns 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

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.



Do not enable management access from the internet or from other untrusted zones inside your enterprise security boundary. Follow the [Best Practices for Securing Administrative Access](#) to ensure that you are properly securing your firewall.



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](#).

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.

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, don't enable HTTP or Telnet because those protocols transmit in plaintext and therefore aren't secure. Or if you plan to use the MGT interface for firewall configuration tasks through the web interface or CLI, you don't enable HTTP, HTTPS, SSH, or Telnet so that you prevent unauthorized access through the interface (if you must allow HTTPS or SSH in this scenario, 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**.

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 [Customize Service Routes to Services for Virtual Systems](#).

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 Networks Services, URL Updates, 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. To limit the drop-down list for Source Address, select **Source Interface** and select the interface you just configured. Then select a Source Address (from that interface) as the service route.

If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you to select an IP address.
 - To create a service route for a custom destination, select **Destination**, and click **Add**. Enter a **Destination** IP address. An incoming packet with a destination address that matches this address will use as its source the Source Address you specify for this service route. To limit the drop-down for Source Address, select a **Source Interface**. If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you to select an IP address.
3. Click **OK** to save the settings.
4. Repeat Steps 5.2 - 5.3 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).

Name	Type	Source		Destination		Application	Service	Action
		Zone	Address	Zone	Address			
Palo Alto Networks Services	universal	Management	any	Internet	any	paloalto-updates paloalto-willflow-check pan-rtb-check	application-default	Allow

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.

Name	Original Packet		Translated Packet	
	Source Zone	Destination Zone	Source Translation	Destination Translation
Source NAT	Management	Internet	dynamic-ip-and-port	none

STEP 8 | Select **Device > Troubleshooting** and verify that you have connectivity from the data port to the external services, including the default gateway, using the **Ping** connectivity test, and the Palo Alto Networks Update Server using the **Update Server Connectivity** test. In this example, the firewall connectivity to the Palo Alto Networks Update Server is tested.

After you verify you have the required network connectivity, continue to [Register the Firewall](#) and [Activate Subscription Licenses](#).

1. Select **Update Server** from the Select Test drop-down.
2. **Execute** the Palo Alto Networks Update Server connectivity test.



3. Access the firewall CLI, and use the following 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. Before you can register a firewall, though, you must first have an active support account. Perform one of the following tasks depending on whether you have an active support account:

- If you don't have an active support account, then [Create a New Support Account and Register a Firewall](#).
- If you already have an active support account, then you are ready to [Register a Firewall](#).
- (Optional) [Perform Day 1 Configuration](#) on a registered firewall.




If you are [registering a VM-Series firewall](#), refer to the [VM-Series Deployment Guide](#) for instructions.

Create a New Support Account and Register a Firewall

If you do not already have an active Palo Alto Networks support account, then you need to register your firewall when you create your new support account.

STEP 1 | Go to the [Palo Alto Networks Customer Support Portal](#).

STEP 2 | Click **Create my account**.



STEP 3 | Enter **Your Email Address**, check **I'm not a robot**, and click **Submit**.

The screenshot shows the 'Create a New Support Account' page on the Palo Alto Networks Customer Support portal. The page has a dark blue header with the Palo Alto Networks logo and 'Customer Support' text. A left sidebar contains navigation links: 'Support Home', 'Resources', 'Contact Us', 'Feedback', 'Help', 'Privacy Policy', and 'Terms of Service'. The main content area is titled 'Create a New Support Account' and contains a form with the following fields: 'Account Email' (with a dropdown arrow), 'First Name', 'Last Name', and a checkbox labeled 'I'm not a robot'. A red asterisk indicates a required field. A 'Submit' button is located at the bottom right of the form.

STEP 4 | Select **Register device using Serial Number or Authorization Code** and click **Next**.

The screenshot shows the 'DEVICE REGISTRATION' page on the Palo Alto Networks Customer Support portal. The page has a dark blue header with the Palo Alto Networks logo and 'Customer Support' text. A left sidebar contains navigation links: 'Support Home', 'Support Cases', 'Contact Us', 'Feedback', 'Help', 'Privacy Policy', and 'Terms of Service'. The main content area is titled 'DEVICE REGISTRATION' and contains a 'Select Device Type' section with a dropdown menu. Below the dropdown, there is a link: 'Register device using Serial Number or Authorization Code'. A 'Next' button is located at the bottom right of the page.

STEP 5 | Complete the registration form.

1. Enter the contact details for the person in your organization who will own this account. Required fields are indicated by red asterisks.
2. Create a UserID and Password for the account. Required fields are indicated by red asterisks.
3. Enter the **Device Serial Number** or **Auth Code**.
4. Enter your **Sales Order Number** or **Customer Id**.
5. To ensure that you are always alerted to the latest updates and security advisories, **Subscribe to Content Update Emails**, **Subscribe to Security Advisories**, and **Subscribe to Software Update Emails**.
6. Select the check box to agree to the End User Agreement and **Submit**.

Register a Firewall

If you already have an active Palo Alto Networks Customer Support account, perform the following task to register your firewall.


STEP 1 | Log in to the firewall 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, if not already logged in, **Sign In** now.



STEP 4 | Register the firewall.

1. On the Support Home page, click **Register a Device**.



2. Select **Register device using Serial Number or Authorization Code**, and then click **Next**.



3. Enter the firewall **Serial Number** (you can copy and paste it from the firewall Dashboard).
4. (Optional) Enter the **Device Name** and **Device Tag**.
5. (Optional) If the device will not have a connection to the internet, select the **Device will be used offline** check box and then, from the drop-down, select the **OS Release** you plan to use.
6. Provide information about where you plan to deploy the firewall including the **Address**, **City**, **Postal Code**, and **Country**.
7. Read the End User License Agreement (EULA) and the Support Agreement, then **Agree and Submit**.

You can view the entry for the firewall you just registered under **Devices**.

(Optional) Perform Day 1 Configuration

After you register your firewall, you have the option of running Day 1 Configuration. The Day 1 Configuration tool provides configuration templates informed by Palo Alto Networks best practices, which you can use as a starting point to build the rest of your configuration.

The benefits of Day 1 Configuration templates include:

- Faster implementation time
- Reduced configuration errors
- Improved security posture

Perform Day 1 Configuration by following these steps:

STEP 1 | From the page that displays after you have registered your firewall, select **Run Day 1 Configuration**.




If you've already registered your firewall but haven't run Day 1 Configuration, you can also run it from the Customer Support Portal home page by navigating to Tools > Run Day 1 Configuration.

STEP 2 | Enter the **Hostname** and **Pan OS Version** for your new device, and optionally, the **Serial Number** and **Device Type**.



STEP 3 | Under **Management**, select either **Static** or **DHCP Client** for your **Management Type**.

Selecting **Static** will require you fill out the **IPV4**, **Subnet Mask**, and **Default Gateway** fields.



Selecting **DHCP Client** only requires that you enter the **Primary DNS** and **Secondary DNS**. A device configured in DHCP client mode will ensure the management interface receives an IP address from the local DHCP server, or it will fill out all the parameters if they are known.

Management

Management Type* OnCP Cluster

Primary DNS* 1.1.1.1

Secondary DNS* 1.0.0.1

STEP 4 | Fill out all fields under **Logging**.

STEP 5 | Click **Generate Config File**.

Logging

SMTP Server IP* 10.0.0.25

From* firewall@mycompany.com

To* admin@mycompany.com

Logging Server IP* 10.0.0.100

Generate Config File

STEP 6 | To import and load the Day 1 Configuration file you just downloaded to your firewall:

1. Log into your firewall web interface.
2. Navigate to **Device > Setup > Operations**.
3. Click **Import named configuration snapshot**.
4. Select the file.

paloalto

Dashboard ACC Monitor Policies Objects Network Device

Setup

High Availability

Config Audit

Password Profiles

Administrators

Admin Roles

Authentication Profile

Authentication Sequence

User Identification

URL Information Sources

Troubleshooting

Certificate Management

Certificates

Certificate Profile

OCSP Responder

SSL/TLS Service Profile

SCP

SSL Certificate Features

Management Operations Service Interface Scheduling Content-ID WAF IPS Session

Configuration Management

Revert Revert to last saved configuration

Revert to running configuration

Save Save named configuration snapshot

Save certificate configuration

Load Load named configuration snapshot

Load configuration version

Export Export named configuration snapshot

Export configuration version

Export device state

Import Import named configuration snapshot

Import device state


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 [Network Segmentation Using Zones](#). 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. [Configure Interfaces](#) on the firewall to support 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 as [Virtual Wire Interfaces](#) or as [Layer 2 Interfaces](#)), 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 Network Access for External Services](#).

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**.

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 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, you should 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**.

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.



Set Up a Basic Security Policy

Now that you 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 that allows 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.

Even though a Security policy rule allows a packet, this does not mean that the traffic is free of threats. To enable the firewall to scan the traffic that it allows based on a Security policy rule, you must also attach [Security Profiles](#)—including URL Filtering, Antivirus, Anti-Spyware, File Blocking, and WildFire Analysis—to each rule (the profiles you can use depend on which [Subscriptions](#) you purchased). When creating your basic Security policy, use the predefined security profiles to ensure that the traffic you allow into your network is being scanned for threats. You can customize these profiles later as needed for your environment.

Use the following workflow set up a very basic Security policy that enables access to the network infrastructure, to data center applications, and to the internet. This enables you to get the firewall up and running so that you can verify that you have successfully configured the firewall. However, this initial policy is not comprehensive enough to protect your network. After you verify that you successfully configured the firewall and integrated it into your network, proceed with creating a [Best Practice Internet Gateway Security Policy](#) that safely enables application access while protecting your network from attack.

STEP 1 | (Optional) Delete the default Security policy rule.

By default, the firewall includes a Security policy 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 | Allow access to your network infrastructure resources.

1. Select **Policies > Security** and click **Add**.
2. In the **General** tab, enter a descriptive **Name** for the rule.
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, use 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**, and **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. Set **Profile Type** to **Profiles** and select the following security profiles to attach to the policy rule:

- For **Antivirus**, select **default**
 - For **Vulnerability Protection**, select **strict**
 - For **Anti-Spyware**, select **strict**
 - For **URL Filtering**, select **default**
 - For **File Blocking**, select **basic file blocking**
 - For **WildFire Analysis**, select **default**
9. Verify that **Log at Session End** is enabled. Only traffic that matches a Security policy rule will be logged.
 10. Click **OK**.

		Source		Destination		Application	Service	Action	Profile	Options
Name	Type	Zone	Address	Zone	Address					
Network Infrastructure	universal	Users	any	Internet	any	DC-Infrastructure	any	application-default	Allow	

STEP 3 | 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 which applications your users need to access, you can make informed decisions about which applications to allow and create more granular application-based rules for each user group.

1. Select **Policies > Security** and **Add** a rule.
2. In the **General** tab, enter a descriptive **Name** for the rule.
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. Set **Profile Type** to **Profiles** and select the following security profiles to attach to the policy rule:
 - For **Antivirus**, select **default**
 - For **Vulnerability Protection**, select **strict**
 - For **Anti-Spyware**, select **strict**
 - For **URL Filtering**, select **default**
 - For **File Blocking**, select **strict file blocking**
 - For **WildFire Analysis**, select **default**
9. Verify that **Log at Session End** is enabled. Only traffic that matches a security rule will be logged.
10. Click **OK**.

		Source		Destination		Application	Service	Action	Profile	Options
Name	Type	Zone	Address	Zone	Address					
Internet Access	universal	Users	any	Internet	any	Internet	application-default	Allow		

STEP 4 | Enable access to data center applications.

1. Select **Policies > Security** and **Add** a rule.
2. In the **General** tab, Enter a descriptive **Name** for the rule.
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. Set **Profile Type** to **Profiles** and select the following security profiles to attach to the policy rule:
 - For **Antivirus**, select **default**
 - For **Vulnerability Protection** select **strict**
 - For **Anti-Spyware** select **strict**
 - For **URL Filtering** select **default**
 - For **File Blocking** select **basic file blocking**
 - For **WildFire Analysis** select **default**
9. Verify that **Log at Session End** is enabled. Only traffic that matches a security rule will be logged.
10. Click **OK**.

Name	Type	Zone	Address	Port	Address	Port	Applications	Service	Action	Profile	Options
Data Center Applications	universal	any	any	any	any	any	activesync, imap, kerberos, ldap, ms-exchange, ms-lync	application-default	allow	Antivirus, Vulnerability Protection, Anti-Spyware, URL Filtering, File Blocking, WildFire Analysis	

STEP 5 | Save your policy rules to the running configuration on the firewall.

Click **Commit**.

STEP 6 | 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.

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:

1. Select **Device > Troubleshooting** and select **Security Policy Match (Select Test)**.
2. Enter the **Source** and **Destination** IP addresses.
3. Enter the **Protocol**.
4. Select **dns (Application)**
5. **Execute** the Security policy match test.

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.

- [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 how to [Create a Security Policy Rule](#) and [Security Profiles](#).

- [View 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.

- [Configure Log Storage Quotas and Expiration Periods](#).

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.

Enable Free 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, phishing, 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 subscription within minutes. This enables all Palo Alto next-generation firewalls worldwide to detect and prevent malware found by a single firewall. Malware signatures often match multiple variants of the same malware family, and as such, block new malware variants that the firewall has never seen before. The Palo Alto Networks threat research team uses the threat intelligence gathered from malware variants to block malicious IP addresses, domains, and URLs.

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 you 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 within a minute of availability—new signatures are released every five minutes.
- Forward advanced file types and email links for analysis.
- Use the WildFire API.
- Use a WildFire 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:

STEP 1 | Confirm that your firewall is registered and that you have a valid support account as well as any subscriptions you require.

1. Log in to the [Palo Alto Networks Customer Support Portal\(CSP\)](#) and on the left-hand side navigation pane, select **Assets > Devices**.
2. Verify that the firewall is listed. If it is not listed, select **Register New Device** and continue to [Register the Firewall](#).
3. (Optional) If you have a Threat Prevention subscription, be sure to [Activate Subscription Licenses](#).

STEP 2 | Log in to the firewall and configure WildFire forwarding settings.

1. Select **Device > Setup > WildFire** and edit the General Settings.
2. Set the **WildFire Public Cloud** field to forward files to the WildFire global cloud at: **wildfire.paloaltonetworks.com**.



You can also forward files to a [regional cloud](#) or a [private cloud](#) based on your location and your organizational requirements.

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. **Add** a forwarding rule and enter a **Name** for it.
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 rule or create a new policy rule 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 rule 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 | **Commit** 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. (WildFire only assigns the phishing verdict to email links). The Action column indicates whether the firewall allowed or blocked the sample. The [Severity](#) column indicates how much of a threat a sample poses to an organization using the following values: critical, high, medium, low, information.

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 scheduled to download, and install Antivirus updates.

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:

- ❑ Follow the [Best Practices for Securing Administrative Access](#) to make sure you are properly securing the management interfaces.
- ❑ Configure a best-practice security policy rulebase to safely enable applications and protect your network from attack. Go to the [Best Practices](#) page and select security policy best practice for your firewall deployment.
- ❑ 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.
- ❑ 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.
- ❑ Follow the [Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions](#).
- ❑ [Share Threat Intelligence with Palo Alto Networks](#)—Permit the firewall to periodically collect and send information about applications, threats, and device health to Palo Alto Networks. Telemetry includes options to enable passive DNS monitoring and to allow experimental test signatures to run in the background with no impact to your security policy rules, firewall logs, or firewall performance. All Palo Alto Networks customers benefit from the intelligence gathered from telemetry, which Palo Alto Networks uses to improve the threat prevention capabilities of the firewall.

Best Practices for Securing Administrative Access

Protecting your network from cyberattacks begins with a secure firewall deployment. If the network you use to manage your sensitive IT devices—including your Palo Alto Networks next-generation firewalls and Panorama—is not secured properly, you can't detect and defend against vulnerability exploits that could lead to infiltration or the loss of sensitive data. The ultimate goal when securing firewall access is to ensure that even if an attacker gains access to privileged credentials, you can still thwart their ability to get in and do damage. Follow these best practice guidelines to ensure that you secure administrative access to your firewalls and other security devices in a way that prevents successful attacks.

- [Isolate the Management Network](#)
- [Use Service Routes to Access External Services](#)
- [Restrict Access to the Management Interface](#)
- [Manage Administrator Access](#)
- [Create Strong Administrator Passwords](#)
- [Scan All Traffic Destined for the Management Interface](#)
- [Replace the Certificate for Inbound Management Traffic](#)
- [Keep Content and Software Updates Current](#)

Isolate the Management Network

All Palo Alto Networks firewalls provide an out-of-band management (MGT) port that you can use to perform firewall administration functions. Alternatively, you can use the MGT port for initial configuration and then configure a data port for management access to the firewall. Either way, because the management interface provides access to your Security configuration, you must take the following precautions to safeguard access to this interface:




Do not enable access to your management interface from the internet or from other untrusted zones inside your enterprise security boundary. This applies whether you use the dedicated management (MGT) port or a data port as your management interface.

- ❑ Isolate the management interface on a dedicated management VLAN.
- ❑ Limit the source IP addresses allowed in to the management network to IP addresses of your dedicated management devices, such as a jump server or a bastion host.
- ❑ Use a jump server or bastion host (with screen recording) to provide secure access from your corporate network in to your dedicated management network and require that users authenticate and are authorized to access your management network.
- ❑ If you don't have a bastion host, use [Authentication Policy](#) with multi-factor authentication (MFA) to require administrators to successfully authenticate before you allow them to continue to the firewall web interface login page or CLI login prompt. This prevents access to the management interface using stolen credentials or through vulnerability exploits.
- ❑ Limit access to users in your security administrator, network administrator, or IT user groups as appropriate for your organization.
- ❑ If you must enable remote access to the management network, require access through a VPN tunnel using Prisma™ Access. After administrators successfully establish a VPN tunnel into your VPN zone, they must still authenticate in to the management network through your bastion host.

- ❑ Do not use an interface management profile that allows HTTP, HTTPS, Telnet, or SSH on the interface where you configured a GlobalProtect portal or gateway because this configuration exposes the management interface to internet access. Also do not use HTTP or Telnet internally because these protocols transmit in cleartext.
- ❑ If you are using a template to deploy a VM-Series firewall that includes a field for restricting management access to a specific IP address, supply a CIDR block that corresponds to your dedicated management IP addresses or network. If necessary, modify the corresponding security group to add additional hosts or networks after the template launch. Do not make the allowed source network range larger than necessary and never configure the allowed source using 0.0.0.0/0.

Use Service Routes to Access External Services

By default, the firewall uses the management (MGT) port to access services that are outside of the management network on potentially untrusted networks, such as DNS servers, NTP servers, and authentication servers. This includes services that require internet access, such as Palo Alto Networks services and AutoFocus™. Because your management interface—whether on the MGT port or a data port—must be isolated on the management network, you must use service routes (**Device > Setup > Services > Service Route Configuration**) to enable access to these services. When you configure a service route, the firewall uses the specified source interface and address to access the services you need. Specify the source IP address or interface for your service route on an interface that doesn't allow management access (an interface that is not using HTTPS or SSH).




Restrict Access to the Management Interface

- ❑ Restrict the IP addresses that you permit to access the management interface.

Even if your firewall is on a dedicated management network that is only accessible by a device on the same VLAN or through a bastion host or VPN tunnel, you can further secure the firewall by restricting the source IP addresses to administrators who are allowed to access the management interface. This prevents access from unexpected IP addresses or subnets and prevents access through stolen credentials.

- ❑ Restrict the services that are available on the management interface.


- ❑ Do not allow access over Telnet and HTTP because these services use plaintext, are not as secure as the other services, and could compromise administrator credentials. Instead, require administrators to access firewall interfaces over SSH or HTTPS.
- ❑ Enable ping if you want to test connectivity to the interface but do not enable any other services on the management interface.
- ❑ The way you configure these settings depends on whether you are using the MGT port or a data port for access to the firewall management interfaces:
 - If you are using the MGT port as your management interface, select **Device > Setup > Interfaces** and select the **Management** interface to configure the settings to restrict who can access the management interface and what services the interface allows.



- If you are using a data port as your management interface, after you [configure the interface](#), select **Network > Network Profiles > Interface Mgmt** and **Add** an interface management [profile](#) to restrict who can access the management interface and what services the interface allows.



Do not attach an interface management profile that allows Telnet, SSH, HTTP, or HTTPS to an interface where you have configured a GlobalProtect portal or gateway because this will expose the management interface to the internet. Also, do not use HTTP or Telnet for any management interface profile because those protocols transmit in cleartext.



Manage Administrator Access

- ❑  **Change the default admin account password** (*Device > Administrators > admin*) immediately after initial configuration. The firewall comes preconfigured with a default administrative account (*admin*) that provides full read-write access (also known as superuser access) to the firewall but you should not leave it with the default account password.

If compliance, audit, or security requirements stipulate that you must remove the default administrative account from your devices, you can remove it only after you create at least one other superuser administrative account on the device.

- ❑ **Configure a Firewall Administrator Account** for each individual 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.
- ❑ Assign each administrator account to an admin **role** profile that limits management privileges to only those functions the individual administrator needs.
- ❑ For administrators with change privileges, require MFA using external authentication and authorization using RADIUS or SAML ([Configure Local or External Authentication for Firewall Administrators](#) with MFA).



If you have a strong authentication infrastructure using smart cards, [Configure Certificate-Based Administrator Authentication to the Web Interface](#) and [Configure SSH Key-Based Administrator Authentication to the CLI](#).

If available, use privileged account management (PAM) or privileged identity management (PIM) solutions (or both) to secure administrator credentials externally.

- ❑ Monitor the System logs to identify abnormal account activity on any of your administrator accounts. For example, if the logs show excessive login attempts or repeated logins at certain times of day, this may indicate that an administrative account is compromised. Also, educate all administrators about how to [Use the Administrator Login Activity Indicators to Detect Account Misuse](#).

Create Strong Administrator Passwords

Configure a strict password policy that requires frequent password changes (**Device > Setup > Management > Minimum Password Complexity**).

You are responsible for assessing the appropriate password requirements for your organization; however, the following characteristics are best practices for creating strong passwords. Passwords should:

- Be a minimum of eight characters
- Not be based on a single dictionary word
- Not include context specific words (for example, the name of a website)
- Not include a username or derivatives of a username (for example, @dmin, Johnny)
- Not have repetitive or sequential characters (for example, aaaaaa, 1234abcd)
- Include uppercase and lowercase characters, numbers, and special characters (including spaces)

One way to create a strong password is to create a long passphrase rather than a complex password. Industry standards recommend creating long, unique passphrases that you will remember easily (using whatever characters you want, including dictionary words) instead of creating convoluted and complex passwords that are easy to forget. Longer passwords with a minimum of 15 characters are believed to compensate for the use of dictionary words. Try to create a passphrase based on long, familiar phrases that only you know or string together at least four words.

For more information on how to determine the appropriate password requirements for your organization, we recommend the following resources:

- [NIST SP 800-63B, Digital Identity Guidelines](#)
- [NIST, Easy Ways to Build a Better P@\\$5w0rd](#)

Scan All Traffic Destined for the Management Interface



Because security policy and decryption policy do not evaluate management plane traffic, you cannot directly scan the MGT port for threats. If you are using the MGT port as your management interface, consider routing traffic destined for the MGT port through a data port or through another firewall so that you can apply these important security checks to your management traffic.

- ❑ Create Security policy rules to allow access to the management interfaces of the firewall and Panorama (web interface or CLI). How you define the policy depends on whether or not you are using a bastion host to enable access to the management network.
 - If you are not using a bastion host to isolate your management network, create a Security policy rule to allow access from the users zone to the IT Infrastructure zone. This Security policy rule must be very granular and must specify the source zone, source IP address (if available), and source user group of the user attempting to access the management interface. This rule must also specify the destination zone, IP address of the appliance (firewall or Panorama), and the App-ID to identify the specific management application (web interface or CLI) running on the application default port. For example, use the panos-web-interface App-ID to allow access to the web interface and the ssh App-ID to allow access to the CLI. You must also attach a Vulnerability Protection profile to the rule as described later in this section.

The following example rule allows access from the users zone directly to the IT Infrastructure zone, and restricts access to users in the IT-admins group who are attempting to access the management

interface IP address to access the panos-web-interface application on the application-default port only.

Name	Source			Destination		Application	Service	Action	Profile
	Zone	Address	User	Zone	Address				
1	any	any	any	any	any	any	any	any	any
2	FW-mgt	any	IT-admins	IT-infrastructure	192.168.1.1	panos-web-interface	application-default	Allow	

- If you are using a bastion host to enable access to your management network, you need two Security policy rules: one rule to allow access from the user zone to the bastion host zone and a second rule to allow access from the bastion host zone to the IT infrastructure zone. Again, both of these Security policy rules should be as granular as possible and include source zone, address (if available), user and destination zone and address, and the App-ID. If you are using a bastion host, the user IP address is usually the IP address of the bastion host so you cannot identify the User-ID for the administrator unless you are using the Terminal Server [agent](#) on the bastion host to identify individual users. In this case, you must also attach a Vulnerability Protection profile to both rules as described later in this section.

In the example rules that follow, the first rule allows access from the user zone to the Bastion-host zone for users in the IT-admins group who are attempting to access the specified bastion server IP address over SSH or RDP. The second rule allows access to users from the Bastion-host zone to the IT-infrastructure zone attempting to access the panos-web-interface application on the default port on the firewall with the specified destination address.

Name	Source			Destination		Application	Service	Action	Profile
	Zone	Address	User	Zone	Address				
1	any	any	any	any	any	any	any	any	any
2	Bastion-host-access	any	IT-admins	Bastion-host	192.168.2.5	any-rdp	application-default	Allow	
3	FW-mgt	Bastion-host	IT-admins	IT-infrastructure	192.168.1.1	panos-web-interface	application-default	Allow	

- Attach a [best practice Vulnerability Protection profile](#) to the Security policy rules that allow access in to your management network to protect against buffer overflows, illegal code execution, and other attempts to exploit client- and server-side vulnerabilities. To create a profile for the purpose of protecting your management interface, clone the strict profile and then enable extended packet capture to help you track down the source of any potential attacks.

Name	Location	Count	Rule Name	Threat Name	Host Type	Severity	Action	Packet Capture
Best Practice High (Strict) Profile	mgt-mgmt-1	Rules: 10	single-client-strict	any	client	critical	reset both	extended-capture
			single-client-high	any	client	high	reset both	extended-capture
			single-client-medium	any	client	medium	reset both	extended-capture
			single-client-informational	any	client	informational	default	disable
			single-client-low	any	client	low	default	single-packet
			single-server-critical	any	server	critical	reset both	extended-capture
			single-server-high	any	server	high	reset both	extended-capture
			single-server-medium	any	server	medium	reset both	extended-capture
			single-server-informational	any	server	informational	default	disable
			single-server-low	any	server	low	default	single-packet

- [Configure SSL Inbound Inspection](#) or [Configure SSL Forward Proxy](#) for traffic to or from the management interface to ensure that you can decrypt and scan the traffic for threats. Attach a [best practice decryption profile](#) to the decryption policy rule to ensure that you block vulnerable SSL/TLS versions such as TLS 1.0 and SSLv3 and that you reject sessions using weak encryption algorithms such as RC4 and 3DES and weak authentication algorithms such as MD5 and SHA1.

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. Use certificates signed by your enterprise CA so that users won't learn to ignore certificate warnings. In addition, in the SSL/TLS profile, set the **Min version** to **TLSv1.2** (the strongest protocol) and set the **Max version** to **Max** so that you continue to use the strongest protocol as stronger versions become available.

Keep Content and Software Updates Current

Current content and software updates ensure that you are always protected by the latest security patches and threat updates.

- ❑ To ensure that you are always alerted to the latest updates and security advisories, go to the [Palo Alto Networks Support Portal](#), select **Edit Profile**, and **Subscribe to Content Update Emails**, **Subscribe to Security Advisories**, and **Subscribe to Software Update Emails**. Make sure you **Save Edits**.



- ❑ Follow the [Best Practices for Applications and Threats Content Updates](#) when updating to the latest content release version.
- ❑ Before you [upgrade PAN-OS](#), read the latest [release notes](#).

Subscriptions

Learn about all the subscriptions and services that work with the firewall, and get started by activating subscription licenses:

- > Subscriptions You Can Use With the Firewall
- > Activate Subscription Licenses
- > What Happens When Licenses Expire?
- > Enhanced Application Logs for Palo Alto Networks Cloud Services



Certain cloud services, like Cortex XDR™, do not integrate with the firewall directly, but rely on data stored in Cortex Data Lake for visibility into network activity. Enhanced application logging is a feature that comes with a Cortex Data Lake subscription—it allows the firewall to collect data specifically for Cortex XDR to use to detect anomalous network activity. Turning on enhanced application logging is a Cortex XDR best practice.

Subscriptions You Can Use With the Firewall

The following Palo Alto Networks subscriptions unlock certain firewall features or enable the firewall to leverage a Palo Alto Networks cloud-delivered service (or both). Here you can read more about each service or feature that requires a subscription to work with the firewall. To enable a subscription, you must first [Activate Subscription Licenses](#); once active, most subscription services can use [Dynamic Content Updates](#) to provide new and updated functionality to the firewall.

Subscriptions You Can Use With the Firewall

Threat Prevention	Threat Prevention provides: • Antivirus, anti-spyware (command-and-control), and vulnerability protection. • Built-in external dynamic lists that you can use to secure your network against malicious hosts. • Ability to identify infected hosts that try to connect to malicious domains. • Get Started with Threat Prevention
DNS Security	Provides enhanced DNS sinkholing capabilities by querying DNS Security, an extensible cloud-based service capable of generating DNS signatures using advanced predictive analytics and machine learning. This service provides full access to the continuously expanding DNS-based threat intelligence produced by Palo Alto Networks. To set up DNS Security, you must first purchase and install a Threat Prevention license. • Get Started with DNS Security
URL Filtering	Provides the ability to not only control web-access, but how users interact with online content based on dynamic URL categories. You can also prevent credential theft by controlling the sites to which users can submit their corporate credentials. To set up URL Filtering, 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. • Get Started with URL Filtering
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 an on-premise WF-500 appliance. • Get Started with WildFire

Subscriptions You Can Use With the Firewall

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. • Get Started with AutoFocus
 Cortex Data Lake <i>Cortex Data Lake was previously called the Logging Service. The Customer Support Portal and firewall web interface both still reference the Logging Service in some places, including the device license name that's displayed in the firewall web interface (Device > Licenses).</i>	Provides cloud-based, centralized log storage and aggregation. The Logging Service is required or highly-recommended to support several other cloud-delivered services, including Magnifier, GlobalProtect cloud service, and Traps management service. • Get Started with Cortex Data Lake
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 advanced GlobalProtect features (HIP checks and related content updates, the GlobalProtect Mobile App, IPv6 connections, or a GlobalProtect Clientless VPN) you will need a GlobalProtect license (subscription) for each gateway. • Get Started with GlobalProtect
Virtual Systems	This license is required to enable support for multiple virtual systems on PA-3200 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-5200 Series, and PA-7000 Series firewalls (the base number varies by platform). The PA-800 Series, PA-220, and VM-Series firewalls do not support virtual systems. • Get Started with Virtual Systems

Activate Subscription Licenses

Follow these steps to activate a new license on the firewall.

Certain decryption features like [Decryption Mirroring](#) and [Decryption Broker](#) require you to activate a free license to unlock feature functionality. For those features, you should instead follow the steps to [Activate Free Licenses for Decryption Features](#).

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**.

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 Portal](#). 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 is successfully activated

On the **Device > Licenses** page, verify that the license is successfully activated. For example, after activating the WildFire license, you should see that the license is valid:



Threat Prevention	
Date Issued:	December 20, 2016
Date Expires:	April 19, 2020
Description:	Threat prevention subscription

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.

What Happens When Licenses Expire?

Palo Alto Networks [subscriptions](#) provide the firewall with added functionality and/or access to a Palo Alto Networks cloud-delivered service. Upon license expiration, some subscriptions continue to function in a limited capacity, and others stop operating completely. Here you can find out what happens when each subscription expires.

Subscription	Expiry Behavior
Threat Prevention	Alerts appear in the System Log indicating that the license has expired. You can still: • Use signatures that were installed at the time the license expired. • Use and modify Custom App-ID™ and threat signatures. You can no longer: • Install new signatures. • Roll signatures back to previous versions.
DNS Security	You can still: • Use local DNS signatures if you have an active Threat Prevention license. You can no longer: • Get new DNS signatures.
URL Filtering	PAN-DB You can still: • Enforce policy using custom URL categories. You can no longer: • Enforce policy using PAN-DB categories. BrightCloud You can still: • Filter URLs, but you must choose to either allow all categories or block all categories. Do this by following Step 1 in Enable BrightCloud Filtering. You can no longer: • Get BrightCloud updates.
WildFire	You can still: • Forward PEs for analysis. • Get signature updates every 24-48 hours if you have an active Threat Prevention subscription. You can no longer:

Subscription	Expiry Behavior
	• Get five-minute updates through the WildFire public and private clouds. • Forward advanced file types such as APKs, Flash files, PDFs, Microsoft Office files, Java Applets, Java files (.jar and .class), and HTTP/HTTPS email links contained in SMTP and POP3 email messages. • Use the WildFire API. • Use the WildFire appliance to host a WildFire private cloud or a WildFire hybrid cloud.
AutoFocus	You can still: • Use an external dynamic list with AutoFocus data for a grace period of three months. You can no longer: • Access the AutoFocus portal. • View the AutoFocus Intelligence Summary for Monitor log or ACC artifacts.
Cortex Data Lake	You can still: • Store log data for a 30-day grace period, after which it is deleted. • Forward logs to Cortex Data Lake until the end of the 30-day grace period.
GlobalProtect	You can still: • Use the app for endpoints running Windows and macOS. • Configure single or multiple internal/external gateways. You can no longer: • Access the Linux OS app and mobile app for iOS, Android, Chrome OS, and Windows 10 UWP. • Use IPv6 for external gateways. • Run HIP checks. • Use Clientless VPN. • Enable split tunneling.
VM-Series	You can still: • Configure and use the firewall(s) you had deployed when the license expired.
Support	You can no longer: • Receive software updates. • Download VM images. • Benefit from technical support.

Enhanced Application Logs for Palo Alto Networks Cloud Services

The firewall can collect data that increases visibility into network activity for Palo Alto Networks apps and services, like Cortex XDR. These enhanced application logs are designed strictly for Palo Alto Networks apps and services to consume and process; you cannot view enhanced application logs on the firewall or Panorama. Only firewalls sending logs to [Cortex Data Lake](#) can generate enhanced application logs.

Examples of the types of data that enhanced application logs gather includes records of DNS queries, the HTTP header User Agent field that specifies the web browser or tool used to access a URL, and information about DHCP automatic IP address assignment. With DHCP information, for example, [Cortex XDR™](#) can alert on unusual activity based on hostname instead of IP address. This allows the security analyst using Cortex XDR to meaningfully assess whether the user's activity is within the scope of his or her role, and if not, to more quickly take action to stop the activity.

To benefit from the most comprehensive set of enhanced application logs, you should enable [User-ID](#); deployments for the Windows-based User-ID agent and the PAN-OS integrated User-ID agent both collect some data that is not reflected in the firewall User-ID logs but that is useful towards associating network activity with specific users.

To start forwarding enhanced application logs to Cortex Data Lake, turn on enhanced application logging globally, and then enable it on a per-security rule basis (using a Log Forwarding profile). The global setting is required and captures data for traffic that is not session-based (ARP requests, for example). The per-security policy rule setting is strongly recommended; the majority of enhanced application logs are gathered from the session-based traffic that your security policy rules enforce.

STEP 1 | Enhanced application logging requires a Cortex Data Lake subscription and User-ID is also recommended. Here are steps to [get started with Cortex Data Lake](#) and [enable User-ID](#).

STEP 2 | To **Enable Enhanced Application Logging** on the firewall, select **Device > Setup > Management > Logging Service** and edit Cortex Data Lake Settings.




Cortex Data Lake was previously called the Logging Service; you might continue to see references to the Logging Service in the firewall web interface.



STEP 3 | Continue to enable enhanced application logging for the security policy rules that control the traffic into which you want extended visibility.

1. Select **Objects > Log Forwarding** and **Add** or modify a log forwarding profile.
2. Update the profile to **Enable Enhanced Application Logging to the Logging Service**.



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.

3. Click **OK** to save the profile and continue to update as many profiles as needed.
4. Ensure that the Log Forwarding profile that you've updated is attached to a security policy rule, to trigger log generation and forwarding for the traffic matched to the rule.
 1. Select **Policies > Security** to view the profiles attached to each security policy rule.
 2. To update the log forwarding profile attached to a rule, **Add** or edit a rule and select **Policies > Security > Actions > Log Forwarding** and select the Log Forwarding profile enabled with enhanced application logging.

Software and Content Updates

PAN-OS is the software that runs all Palo Alto Networks next-generation firewalls. Palo Alto Networks also frequently publishes updates to equip the firewall with the latest security features. The firewall can enforce policy based on the applications and threat signatures (and more) that content updates provide, without requiring you to update the firewall configuration.

- > [PAN-OS Software Updates](#)
- > [Dynamic Content Updates](#)
- > [Install Content Updates](#)
- > [Applications and Threats Content Updates](#)
- > [Best Practices for Applications and Threats Content Updates](#)
- > [Content Delivery Network Infrastructure](#)

PAN-OS Software Updates

PAN-OS is the software that runs all Palo Alto Networks next-generation firewalls. The PAN-OS software version that a firewall is running is displayed on the firewall **Dashboard**.

You can check for new PAN-OS releases directly in the firewall, or on the [Palo Alto Networks support portal](#). To upgrade the firewall to the latest version of PAN-OS:

STEP 1 | Review the latest [PAN-OS Release Notes](#) to see what's new. Also take a look at [PAN-OS Upgrade/Downgrade Considerations](#) to make sure you understand all potential changes the PAN-OS release might introduce.

STEP 2 | Check for new PAN-OS releases:

- **On the firewall**—Select **Device** > **Software** and **Check Now** for the firewall to check with the Palo Alto Networks Update Server for new PAN-OS release versions.
- **On the support portal**—Go to support.paloaltonetworks.com and, on the left menu bar, select **Updates** > **Software Updates**. Download and save the release you want to use to upgrade the firewall.

STEP 3 | Once you've decided the release version you want, follow the complete workflow to [upgrade the firewall to a new PAN-OS version](#). The steps you'll take might depend on the release version you're currently running, if you're using HA, and whether or not you're using Panorama to manage firewalls.

Dynamic Content Updates

Palo Alto Networks frequently publishes updates that the firewall can use to enforce security policy, without requiring you to upgrade PAN-OS software or change the firewall configuration. These updates equip the firewall with the very latest security features and threat intelligence.

Except for application updates and some antivirus updates—which any firewall can receive—dynamic content updates available to you might depend on your [Subscriptions](#). You can set a schedule for each dynamic content update to define the frequency at which the firewall checks for and downloads or installs new updates (**Device > Dynamic Updates**).

Dynamic Content Update	What's in this package?
Antivirus	Antivirus updates are released every 24 hours and include: • WildFire signatures for newly-discovered malware. To get these updates every five minutes instead of once daily, you'll need a WildFire subscription. • (Requires Threat Prevention) Automatically-generated command-and-control (C2) signatures that detect certain patterns in C2 traffic. These signatures enable the firewall to detect C2 activity even when the C2 host is unknown or changes rapidly. • (Requires Threat Prevention) New and updated list entries for built-in external dynamic lists. These lists include malicious, high-risk, and bulletproof host-provided IP addresses, and can help to protect you against malicious hosts. • (Requires Threat Prevention) Updates to the local set of DNS signatures that the firewall uses to identify known malicious domains. If you've set up DNS sinkholing, the firewall can identify hosts on your network that try to connect to these domains. To allow the firewall to check domains against the complete database of DNS signatures, set up DNS Security.
Applications	Application updates provide new and modified application signatures, or App-IDs. This update does not require any additional subscriptions, but it does require a valid maintenance/support contract. New application updates are published only on the third Tuesday of every month, to give you time to prepare any necessary policy updates in advance; modifications to App-ID are released more frequently. While new and modified App-IDs enable the firewall to enforce your security policy with ever-increasing precision, resulting changes in security policy enforcement that can impact application availability. To get the most out of application updates, follow our tips to Manage New and Modified App-IDs.
Applications and Threats	Includes new and updated application and threat signatures. This update is available if you have a Threat Prevention subscription (in this case, you will get this update instead of the Applications update). New threat updates are published frequently, sometimes several times a week, along with updated App-IDs. New App-IDs are published only on the third Tuesday of every month. The firewall can retrieve the latest threat and application updates within as little as 30 minutes of availability. For guidance on how to best enable application and threat updates to ensure both application availability and protection against the latest threats, review the Best Practices for Applications and Threats Content Updates.

Dynamic Content Update	What's in this package?
GlobalProtect Data File	Contains the vendor-specific information for defining and evaluating host information profile (HIP) data returned by GlobalProtect apps. You must have a GlobalProtect gateway subscription in order to receive these updates. In addition, you must create a schedule for these updates before GlobalProtect will function.
GlobalProtect Clientless VPN	Contains new and updated application signatures to enable Clientless VPN access to common web applications from the GlobalProtect portal. You must have a GlobalProtect subscription to receive these updates. In addition, you must create a schedule for these updates before GlobalProtect Clientless VPN will function. As a best practice, it is recommended to always install the latest content updates for GlobalProtect Clientless VPN.
WildFire	Provides near real-time malware and antivirus signatures created as a result of the analysis done by the WildFire public cloud. WildFire signature updates are made available every five minutes. You can set the firewall to check for new updates as frequently as every minute to ensure that the firewall retrieves the latest WildFire signatures within a minute of availability. Without the WildFire subscription, you must wait at least 24 hours for the signatures to be provided in the Antivirus update.
WF-Private	Provides near real-time malware and antivirus signatures created as a result of the analysis done by a WildFire appliance. To receive content updates from a WildFire appliance, the firewall and appliance must both be running PAN-OS 6.1 or a later release and the firewall must be configured to forward files and email links to the WildFire Private Cloud.
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.

Install Content Updates

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 [Dynamic Content Updates](#) available to you depend on which [Subscriptions](#) you have.

Follow these steps to install content updates. You can also set a schedule for content updates, to define the frequency at which the firewall retrieves and installs updates.

Applications and Threats content updates work a little differently than other update types—to get the most out of the latest application knowledge and threat prevention, follow the guidelines to [Deploy Applications and Threats Content Updates](#) instead of the steps here.

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 [Content Delivery Network Infrastructure for Dynamic Updates](#). If the firewall has restricted access to the Internet, set the update server address to use the hostname **staticupdates.paloaltonetworks.com** 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. This is enabled by default.
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**.

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**.

WAFFile	Last checked: 2018/01/12 14:00:08 PST	Schedule: None
PAN-OS 7.1.1	PAN-OS 7.1.1 - Full and later	4 MB 2018/01/12 14:00:08 PST



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.



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.

URL Filtering		Schedule: Every day at 01:02 (download-and-install)					
2040							Upgrade
2032							



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 10 minutes on a PA-220 firewall and up to two minutes on a PA-5200 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.

2040-20307	panos2-all-vuln-2040-20307-complete	PAN-OS 7.1 and later	Full	1 MB	2040/01/12 14:13:04 PST	✓	Install

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.

WildFire	Last checked: 2020/01/12 14:13:04 PST	Schedule: None
2040-20307	panos2-all-vuln-2040-20307-complete	PAN-OS 7.1 and later

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 **Weekly**, **Daily**, **Hourly**, or **Every 30 Minutes** 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.



If you have mission critical applications that must be 100% available, set the threshold for Applications or Applications and Threats updates to a minimum of 24 hours or more and follow the [Best Practices for Applications and Threats Content Updates](#). Additionally, While scheduling content updates is a one-time or infrequent task, after

you've set the schedule, you'll need to continue to [Manage New and Modified App-IDs](#) that are included in content releases, as these App-IDs can change how security policy is enforced.



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](#).

Applications and Threats Content Updates

Applications and Threats content updates deliver the very latest application and threat signatures to the firewall. The applications portion of the package includes new and modified App-IDs and does not require a license. The full Applications and Threats content package, which also includes new and modified threat signatures, requires a Threat Prevention license. As the firewall automatically retrieves and installs the latest application and threat signatures (based on your custom settings), it starts enforcing security policy based on the latest App-IDs and threat protection without any additional configuration.

New and modified threat signatures and modified App-IDs are released at least weekly and, often, more frequently. New App-IDs are released on the third Tuesday of every month. Because new App-IDs can change how the security policy enforces traffic, this more limited release of new App-IDs is intended to provide you with a predictable window in which you can prepare and update your security policy. Additionally, content updates are cumulative; this means that the latest content update always includes the application and threat signatures released in previous versions.

Because application and threat signatures are delivered in a single package—the same decoders that enable application signatures to identify applications also enable threat signatures to inspect traffic—you need to consider whether you want to deploy the signatures together or separately. How you choose to deploy content updates depends on your organization's network security and application availability requirements. As a starting point, identify your organization as having one of the following postures (or perhaps both, depending on firewall location):

- An organization with a *security-first* posture prioritizes protection using the latest threat signatures over application availability. You're primarily using the firewall for its threat prevention capabilities. Any changes to App-ID that impact how security policy enforces application traffic is secondary.
- A *mission-critical* network prioritizes application availability over protection using the latest threat signatures. Your network has zero tolerance for downtime. The firewall is deployed inline to enforce security policy and if you're using App-ID in security policy, any change a content releases introduces that affects App-ID could cause downtime.

You can take a mission-critical or security-first approach to deploying content updates, or you can apply a mix of both approaches to meet the needs of the business. Review and consider [Best Practices for Applications and Threats Content Updates](#) to decide how you want to implement application and threat updates. Then:

- ❑ [Deploy Applications and Threats Content Updates.](#)
- ❑ Follow our [Tips for Content Updates.](#)



While scheduling content updates is a one-time or infrequent task, after you've set the schedule, you'll need to continue to [Manage New and Modified App-IDs](#) that are included in content releases, as these App-IDs can change how security policy is enforced.

Deploy Applications and Threats Content Updates

Before you take the steps to configure application and threat content updates, learn about how [Applications and Threats Content Updates](#) work and decide how you want to implement [Best Practices for Applications and Threats Content Updates](#).

Additionally, Panorama enables you to deploy content updates to firewalls easily and rapidly. If you're using Panorama to manage firewalls, follow [these steps to deploy content updates](#) instead of the ones below.

STEP 1 | To unlock the full Applications and Threats content package, get a Threat Prevention license and [activate the license](#) on the firewall.

1. Select **Device > Licenses**.

-
2. Manually upload the license key or retrieve it from the Palo Alto Networks license server.
 3. Verify that the Threat Prevention license is active.

STEP 2 | Set the schedule for the firewall to retrieve and install content updates.

As you complete the following steps, it's particularly important that you consider whether your organization is [mission-critical or security-first](#) (or a mix of both), and that you have reviewed the [Best Practices for Applications and Threats Content Updates](#).

1. Select **Device > Dynamic Updates**.
2. Select the **Schedule** for Applications and Threat content updates.
3. Set how frequently (the **Recurrence**) the firewall checks with the Palo Alto Networks update server for new Applications and Threat content releases, and on what **Day** and **Time**.
4. Set the **Action** for the firewall to take when it finds and retrieves a new content release.
5. Set an installation **Threshold** for content releases. Content releases must be available on the Palo Alto Networks update server at least this amount of time before the firewall can retrieve the release and perform the Action you configured in the last step.
6. If yours is a mission-critical network, where you have zero tolerance for application downtime (application availability is tantamount even to the latest threat prevention), you can set a **New App-ID Threshold**. The firewall only retrieves content updates that contain new App-IDs after they have been available for this amount of time.
7. Click **OK** to save the Applications and Threats content update schedule, and **Commit**.

STEP 3 | Set up log forwarding to send Palo Alto Networks critical content alerts to external services that you use for monitoring network and firewall activity. This allows you to ensure that the appropriate personnel is notified about critical content issues, so that they can take action as needed. Critical content alerts are logged as system log entries with the following Type and Event: (subtype eq content) and (eventid eq palo-alto-networks-message).

STEP 4 | While scheduling content updates is a one-time or infrequent task, after you've set the schedule, you'll need to continue to [Manage New and Modified App-IDs](#) that are included in content releases, as these App-IDs can change how security policy is enforced.

Tips for Content Updates

Palo Alto Networks application and threat content releases 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 an unexpected way. Follow these tips to mitigate or troubleshoot an issue with a content release, so that there is as little impact to your network as possible.

❑ Follow the best practices for Application and Threat Content Updates.

Review and implement the [Best Practices for Applications and Threats Content Updates](#). How you choose to deploy content updates might depend on your network security and application availability requirements.

❑ Ensure that you're running the latest content.

Get the latest content update, if you haven't configured the firewall to download and install it automatically.

The firewall validates that downloaded content updates are 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 removed, even if the firewall has already downloaded it. If you see an error message that the content update you're attempting to install is no longer valid, **Check Now** to get the most recent content update and install that version instead (**Device > Dynamic Updates**).

❑ **Turn on threat intelligence telemetry.**

Turn on the [threat intelligence telemetry](#) that the firewall sends to Palo Alto Networks. We use telemetry data to identify and troubleshoot issues with content updates.

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. The more quickly we can identify an issue, the more quickly we can help you to avoid the issue altogether or mitigate impact to your network.

To enable 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.

❑ **Forward Palo Alto Networks content update alerts to the right people.**

Enable log forwarding for Palo Alto Networks critical content alerts, so that important messages about content release issues go directly to the appropriate personnel.

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. Critical content alerts describe the issue so that you can understand how it affects you, and include steps to take action if needed.

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 **dynamic-updates** and the Event **palo-alto-networks-message**. Use the following filter to view these log entries: (subtype eq dynamic-updates) and (eventid eq palo-alto-networks-message).

❑ **If needed, use Panorama to rollback to an earlier content release.**

After being notified about an issue with a content update, you can use Panorama to quickly revert managed firewalls to the last content update version, instead of manually reverting the content version for individual firewalls: [Revert Content Updates on Managed Firewalls](#).

Best Practices for Applications and Threats Content Updates

The best practices to deploy content updates helps to ensure seamless policy enforcement as the firewall is continually equipped with new and modified application and threat signatures. Even though application and threat signatures are delivered together in a single content update package (read more about [Applications and Threats Content Updates](#)), you have the flexibility to deploy them differently based on your network security and availability requirements:

- An organization with a *security-first* posture prioritizes protection using the latest threat signatures over application availability. You're primarily using the firewall for its threat prevention capabilities.
- A *mission-critical* network prioritizes application availability over protection using the latest threat signatures. Your network has zero tolerance for downtime. The firewall is deployed inline to enforce security policy and if you're using App-ID in security policy, any change to content that affects App-ID could cause downtime.

You can take a mission-critical or security-first approach to deploying content updates, or you can apply a mix of both approaches to meet the needs of the business. Consider your approach as you apply the following best practices to most effectively leverage new and modified threat and application signatures:


- [Best Practices for Content Updates—Mission-Critical](#)
- [Best Practices for Content Updates—Security-First](#)

Best Practices for Content Updates—Mission-Critical


The [Best Practices for Applications and Threats Content Updates](#) help to ensure seamless policy enforcement as new application and threat signatures are released. Follow these best practices to deploy content updates in a *mission-critical network*, where you have zero tolerance for application downtime.

- ❑ Always review Content Release Notes for the list of newly-identified and modified application and threat signatures that the content release introduces. Content Release Notes also describe how the update might impact existing security policy enforcement and provides recommendations on how you can modify your security policy to best leverage what's new.

To subscribe to get notifications for new content updates, visit the [Customer Support Portal](#), edit your **Preferences**, and select **Subscribe to Content Update Emails**.



You can also review [Content Release Notes for apps and threats](#) on the Palo Alto Networks Support Portal or directly in the firewall web interface: select **Device** > **Dynamic Updates** and open the **Release Note** for a specific content release version.




The Notes section of Content Release Notes highlights future updates that Palo Alto Networks has identified as possibly significantly impacting coverage: for example, new App-IDs or decoders. Check for these future updates, so that you can account for any policy impact in advance of the release.

- ❑ Create a security policy rule to always [allow certain categories of new App-IDs](#), like authentication or software development applications on which critical business functions rely. This means that when a content release introduces or changes coverage for an important business application, the firewall continues to seamlessly allow the application without requiring you to update your security policy. This eliminates any potential availability impact for App-IDs in critical categories, and gives you thirty days (new App-IDs are released on a monthly basis) to adjust your security policy to allow the mission-critical App-ID(s).

To do this, create an [application filter for new App-IDs in critical categories](#)(**Objects > Application Filters**), and add the application filter to a security policy rule.



- ❑ To mitigate any impact to security policy enforcement that is associated with enabling new application and threat signatures, stagger the roll-out of new content. Provide new content to locations with less business risk (fewer users in satellite offices) before deploying them to locations with more business risk (such as locations with critical applications). Confining the latest content updates to certain firewalls before deploying them across your network also makes it easier to troubleshoot any issues that arise. You can use Panorama to push staggered schedules and installation thresholds to firewalls and device groups based on organization or location ([Use Panorama to Deploy Updates to Firewalls](#)).
- ❑ Schedule content updates so that they **download-and-install** automatically. Then, set a **Threshold** that determines the amount of time the firewall waits before installing the latest content. In a mission-critical network, schedule up to a 48 hour threshold.




The installation delay ensures that the firewall only installs content that has been available and functioning in customer environments for the specified amount of time. To [schedule content updates](#), select **Device > Dynamic Updates > Schedule**.

- ❑ Give yourself additional time to adjust your security policy based on new App-IDs before you install them. To do this, set an installation threshold that applies only to content updates that contain new App-IDs. Content updates with new App-IDs are released only once a month, and installation threshold triggers only at that time. [Schedule content updates](#) to configure a **New App-ID Threshold (Device > Dynamic Updates > Schedule)**.




- ❑ Always review the new and modified App-IDs that a content release introduces, in order to assess how the changes might impact your security policy. The following topic describes the options you can use to update your security policy both before and after installing new App-IDs: [Manage New and Modified App-IDs](#).



- ❑ [Set up log forwarding](#) to send Palo Alto Networks critical content alerts to external services that you use for monitoring network and firewall activity. This allows you to ensure that the appropriate personnel is

notified about critical content issues, so that they can take action as needed. Critical content alerts are logged as system log entries with the following Type and Event: **(subtype eq dynamic-updates)** and **(eventid eq palo-alto-networks-message)**.



PAN-OS 8.1.2 changed the log type for critical content alerts from **general** to **dynamic-updates**. If you're using PAN-OS 8.1.0 or PAN-OS 8.1.1, critical content are logged as system log entries with the following Type and Event, and you should set up forwarding for these alerts using the following filter: **(subtype eq general)** and **(eventid eq palo-alto-networks-message)**.


- ❑ Test new Applications and Threats content updates in a dedicated staging environment before enabling them in your production environment. The easiest way to test new applications and threats is to use a test firewall to tap into production traffic. Install the latest content on the test firewall and monitor the firewall as it processes the traffic copied from your production environment. You can also use test clients and a test firewall or packet captures (PCAPs) to simulate production traffic. Using PCAPs works well to simulate traffic for diverse deployments where firewall security policy varies depending on location.

Best Practices for Content Updates—Security-First

The [Best Practices for Applications and Threats Content Updates](#) help to ensure seamless policy enforcement as new application and threat signatures are released. Follow these best practices to deploy content updates in a *security-first network*, where you're primarily using the firewall for its threat prevention capabilities and your first priority is attack defense.

- ❑ Always review Content Release Notes for the list of newly-identified and modified application and threat signatures that the content release introduces. Content Release Notes also describe how the update might impact existing security policy enforcement and provides recommendations on how you can modify your security policy to best leverage what's new.

To subscribe to get notifications for new content updates, visit the [Customer Support Portal](#), edit your **Preferences**, and select **Subscribe to Content Update Emails**.




You can also review [Content Release Notes for apps and threats](#) on the Palo Alto Networks Support Portal or directly in the firewall web interface: select **Device > Dynamic Updates** and open the **Release Note** for a specific content release version.

Version	Release Date	Downloaded	Currently Installed	Action	Documentation
Applications and Threats					
8524-4686	2018/05/04 14:00:07 PDT	✓	✓	Download	Release Notes
788-4228	2018/03/02 14:00:01 PST	✓ previously	✓	Download	Release Notes
8004-4694	2018/04/12 17:47:08 PDT			Download	Release Notes
8005-4696	2018/04/12 23:00:11 PDT			Download	Release Notes
8006-4698	2018/04/12 23:00:52 PDT			Download	Release Notes
8007-4699	2018/04/12 23:00:52 PDT			Download	Release Notes
8008-4656	2018/04/12 23:00:52 PDT			Download	Release Notes
8009-4658	2018/04/12 23:00:52 PDT			Download	Release Notes
8010-4660	2018/04/12 23:00:52 PDT			Download	Release Notes



The Notes section of Content Release Notes highlights future updates that Palo Alto Networks has identified as possibly significantly impacting coverage: for example, new App-IDs or decoders. Check for these future updates, so that you can account for any policy impact in advance of the release.

- ❑ To mitigate any impact to security policy enforcement that is associated with enabling new application and threat signatures, stagger the roll-out of new content. Provide new content to locations with less business risk (fewer users in satellite offices) before deploying them to locations with more business risk (such as locations with critical applications). Confining the latest content updates to certain firewalls before deploying them across your network also makes it easier to troubleshoot any issues that arise. You can use Panorama to push staggered schedules and installation thresholds to firewalls and device groups based on organization or location ([Use Panorama to Deploy Updates to Firewalls](#)).
- ❑ Schedule content updates so that they **download-and-install** automatically. Then, set a **Threshold** that determines the amount of time the firewall waits before installing the latest content. In a security-first network, schedule a six to twelve hour threshold.




The installation delay ensures that the firewall only installs content that has been available and functioning in customer environments for the specified amount of time. To [schedule content updates](#), select **Device > Dynamic Updates > Schedule**.



Do not schedule a New App-ID Threshold. This threshold allows mission-critical organizations extra time to adjust security policy enforcement based on new App-IDs. However, because this threshold also delays delivery of the latest threat prevention updates, it is not recommended for organizations with a security-first posture.

- ❑ Review the new and modified App-IDs that a content release introduces, in order to assess how the changes might impact your security policy. The following topic describes the options you can use to

update your security policy both before and after installing new App-IDs: [Manage New and Modified App-IDs](#).



- **Set up log forwarding** to send Palo Alto Networks critical content alerts to external services that you use for monitoring network and firewall activity. This allows you to ensure that the appropriate personnel is notified about critical content issues, so that they can take action as needed. Critical content alerts are logged as system log entries with the following Type and Event: (subtype eq dynamic-updates) and (eventid eq palo-alto-networks-message).



PAN-OS 8.1.2 changed the log type for critical content alerts from **general** to **dynamic-updates**. If you're using PAN-OS 8.1.0 or PAN-OS 8.1.1, critical content are logged as system log entries with the following Type and Event, and you should set up forwarding for these alerts using the following filter: (**subtype eq general**) and (**eventid eq palo-alto-networks-message**).

Content Delivery Network Infrastructure

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 content and application identification functions.

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 - proditpdownloads.paloaltonetworks.com:443	staticupdates.paloaltonetworks.com
Threat/ Antivirus Database	- updates.paloaltonetworks.com:443 - downloads.paloaltonetworks.com:443 - proditpdownloads.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
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.

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:



Do not enable management access from the internet or from other untrusted zones inside your enterprise security boundary. Follow the [Best Practices for Securing Administrative Access](#) to ensure that you are properly securing your firewall.

- [Use the Web Interface](#) to perform configuration and monitoring tasks with relative ease. This graphical interface allows you to access the firewall using HTTPS (recommended) or HTTP and it is the best way to perform administrative tasks.
- [Use the Command Line Interface \(CLI\)](#) to perform a series of tasks by entering commands in rapid succession over SSH (recommended), Telnet, or the console port. The CLI is a no-frills interface that supports two command modes, operational and configure, each with a distinct 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.
- [Use Panorama](#) to perform web-based management, reporting, and log collection for multiple firewalls. The Panorama web interface resembles the firewall web interface but with additional functions for centralized management.

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](#)
- [Export Configuration Table Data](#)
- [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 11+
- Firefox 3.6+
- Safari 5+
- Chrome 11+

Perform the following tasks to 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 > Interfaces and edit the Management interface.

STEP 2 | Log in to the firewall according to the type of authentication used for your account. If logging in to the firewall for the first time, use the default value **admin** for your username and password.

- **SAML**—Click **Use Single Sign-On (SSO)**. If the firewall performs authorization (role assignment) for administrators, enter your **Username** and **Continue**. If the SAML identity provider (IdP) performs authorization, **Continue** without entering a **Username**. In both cases, the firewall redirects you to the IdP, which prompts you to enter a username and password. After you authenticate to the IdP, the firewall web interface displays.
- **Any other type of authentication**—Enter your user **Name** and **Password**. Read the login banner and select **I Accept and Acknowledge the Statement Below** if the login page has the banner and check box. Then click **Login**.

STEP 3 | Read and **Close** the messages of the day.

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.

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).

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. The supported file types are png, gif, and jpg. The firewall does not support image files that are interlaced or that contain alpha channels.

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:

1. Click upload .
2. Select a logo image and click **Open**.



You can preview the image to see how PAN-OS will crop it to fit by clicking the magnifying glass icon.

3. 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.

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.

1. If you see the caution symbol, hover over it to display the number of failed login attempts.




2. 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.

STEP 2 | Locate hosts that are continually attempting to log in to your firewall or Panorama management server.

1. Click the failed login caution symbol to view the failed login attempts summary.
2. 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.



3. 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 3 | 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.

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 pending changes to the firewall configuration. You can filter pending changes by administrator or *location* and then preview, validate, or commit only those changes. The locations can be specific virtual systems, shared policies and objects, or shared device and network settings.

The firewall queues commit requests so that you can initiate a new commit while a previous commit is in progress. The firewall performs the commits in the order they are initiated but prioritizes auto-commits that are initiated by the firewall (such as FQDN refreshes). However, if the queue already has the maximum number of administrator-initiated commits, you must wait for the firewall to finish processing a pending commit before initiating a new one. To cancel pending commits or view details about commits of any status, see [Manage and Monitor Administrative Tasks](#).

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 (warnings). For example, validation could indicate an invalid route destination that you need to fix for the commit to succeed. The validation process enables you to find and fix errors before you commit (it makes no changes to the running configuration). This is useful if you have a fixed commit window and want to be sure the commit will succeed without errors.



The commit, validate, preview, save, and revert operations apply only to changes made after the last commit. To restore configurations to the state they were in before the last commit, you must [load a previously backed up configuration](#).

To prevent multiple administrators from making configuration changes during concurrent sessions, see [Manage Locks for Restricting Configuration Changes](#).

STEP 1 | Configure the scope of configuration changes that you will commit, validate, or preview.

1. Click **Commit** at the top of the web interface.
2. Select one of the following options:
 - **Commit All Changes** (default)—Applies the commit to all changes for which you have administrative privileges. You cannot manually filter the commit scope when you select this option. Instead, the administrator role assigned to the account you used to log in determines the commit scope.
 - **Commit Changes Made By**—Enables you to filter the commit scope by administrator or location. The administrative role assigned to the account you used to log in determines which changes you can filter.



To commit the changes of other administrators, the account you used to log in must be assigned the Superuser role or an [Admin Role profile](#) with the Commit For Other Admins privilege enabled.

3. (Optional) To filter the commit scope by administrator, select **Commit Changes Made By**, click the adjacent link, select the administrators, and click **OK**.
4. (Optional) To filter by location, select **Commit Changes Made By** and clear any changes that you want to exclude from the Commit Scope.



If dependencies between the configuration changes you included and excluded cause a validation error, perform the commit with all the changes included. For example, when you commit changes to a virtual system, you must include the changes of all administrators who added, deleted, or repositioned rules for the same rulebase in that virtual system.

STEP 2 | 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 compares the configurations you selected in the Commit Scope to the running configuration. The preview window displays the configurations side-by-side and uses color coding to indicate which changes are additions (green), modifications (yellow), or deletions (red).

Preview Changes and 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. Close the preview window when you finish reviewing the changes.



Because the preview results display in a new browser window, your browser must allow pop-ups. If the preview window does not open, refer to your browser documentation for the steps to allow pop-ups.

STEP 3 | Preview the individual settings for which you are committing changes.

This can be useful if you want to know details about the changes, such as the types of settings and who changed them.

1. Click **Change Summary**.
2. (Optional) **Group By** a column name (such as the **Type** of setting).
3. **Close** the Change Summary dialog when you finish reviewing the changes.

STEP 4 | Validate the changes before you commit to ensure the commit will succeed.

1. **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.

STEP 5 | Commit your configuration changes.

Commit your changes to validate and activate them.



To view details about commits that are pending (which you can still cancel), in progress, completed, or failed, see [Manage and Monitor Administrative Tasks](#).

Export Configuration Table Data

Export policy rules, configuration 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 your firewalls and appliances, to take screen shots or to access 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 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 into Panorama or the firewall. When you export data as a PDF file and the table data exceeds 50,000 rows, the data is split in to multiple PDF files (for example, <report-name>_part1.pdf and <report-name>_part2.pdf). When you export data as a CSV file, the data is exported as a single file. These export formats allow you 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.

STEP 1 | Launch the [Web Interface](#) and identify the configuration data you need to export.

STEP 2 | Apply filters as needed to produce the configuration data you need to export and click **PDF/CSV**.



STEP 3 | Configure the Configuration Table Export report:

1. Enter a **File Name**.
2. Select the **File Type**.
3. (Optional) Enter a report Description.
4. Confirm the configuration table data matches the filters you applied.



Select **Show All Columns** to show all filters applied.

STEP 4 | Export the configuration table data.

Configuration table export works like a print function—you cannot import generated files back in to Panorama or the firewall.



STEP 5 | Select a location to save the exported file.

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, UUID, or application name. In addition to searching for configuration objects and settings, you can search by job ID or job type for manual commits that administrators performed or auto-commits that the firewall or Panorama performed. 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.


- 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**:



For example, click **Global Find** on a zone named **I3-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 I3-vlan-trust:



Search tips:

- If you initiate a search on a firewall that has multiple virtual systems enabled or if custom [Administrative Role Types](#) 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.
- To search for a UUID, you must copy and paste the UUID.

Manage Locks for Restricting Configuration Changes

You can use configuration locks to prevent other administrators from changing the candidate configuration or from committing configuration changes until you manually remove 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](#).

- 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.

- 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 committing changes made to the candidate 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 apply a commit lock 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**.

Manage Configuration Backups

The running configuration on the firewall 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. Saving backup versions of the running or candidate configuration enables you to later restore those versions. For example, if a commit validation shows that the current candidate configuration has more errors than you want to fix, you can restore a previous candidate configuration. You can also revert to the current running configuration without saving a backup first. If you need to export specific parts of the configuration for internal review or audit, you can [Export Configuration Table Data](#).



See [Commit, Validate, and Preview Firewall Configuration Changes](#) for details about commit operations.

- [Save and Export Firewall Configurations](#)
- [Revert Firewall Configuration Changes](#)

Save and Export Firewall Configurations

Saving a backup of the candidate configuration to persistent storage on the firewall enables you to later revert to that backup (see [Revert Firewall Configuration Changes](#)). This is useful for preserving changes that would otherwise be lost if a system event or administrator action causes the firewall to reboot. After rebooting, PAN-OS automatically reverts to the current version of the running configuration, which the firewall stores in a file named `running-config.xml`. Saving backups is also useful if you want to revert to a firewall configuration that is earlier than the current running configuration. The firewall does not automatically save the candidate configuration to persistent storage. You must manually save the candidate configuration as a default snapshot file (`.snapshot.xml`) or as a custom-named snapshot file. The firewall stores the snapshot file locally but you can export it to an external host.



You don't have to save a configuration backup to revert the changes made since the last commit or reboot; just select **Config > Revert Changes** (see [Revert Firewall Configuration Changes](#)).

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](#)).

Palo Alto Networks recommends that you back up any important configuration to a host external to the firewall.

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.

To overwrite the default snapshot file (`.snapshot.xml`) with all the changes that all administrators made, perform one of the following steps:

- Select **Device > Setup > Operations** and **Save candidate configuration**.

-
- Log in to the firewall with an administrative account that is assigned the Superuser role or an [Admin Role profile](#) with the **Save For Other Admins** privilege enabled. Then select **Config > Save Changes** at the top of the web interface, select **Save All Changes** and **Save**.

To create a snapshot that includes all the changes that all administrators made but without overwriting the default snapshot file:

1. Select **Device > Setup > Operations** and **Save named configuration snapshot**.
2. Specify the **Name** of a new or existing configuration file.
3. Click **OK** and **Close**.

To save only specific changes to the candidate configuration without overwriting any part of the default snapshot file:

1. Log in to the firewall with an administrative account that has the [role privileges](#) required to save the desired changes.
2. Select **Config > Save Changes** at the top of the web interface.
3. Select **Save Changes Made By**.
4. To filter the Save Scope by administrator, click **<administrator-name>**, select the administrators, and click **OK**.
5. To filter the Save Scope by location, clear any locations that you want to exclude. The locations can be specific virtual systems, shared policies and objects, or shared device and network settings.
6. Click **Save**, specify the **Name** of a new or existing configuration file, and click **OK**.

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.

Revert Firewall Configuration Changes

Revert operations replace settings in the current candidate configuration with settings from another configuration. Reverting changes is useful when you want to undo changes to multiple settings as a single operation instead of manually reconfiguring each setting.

You can revert pending changes that were made to the firewall configuration since the last commit. The firewall provides the option to filter the pending changes by administrator or *location*. The locations can be specific virtual systems, shared policies and objects, or shared device and network settings. If you saved a snapshot file for a candidate configuration that is earlier than the current running configuration (see [Save and Export Firewall Configurations](#)), you can also revert to that snapshot. Reverting to a snapshot enables you to restore a candidate configuration that existed before the last commit. The firewall automatically saves a new version of the running configuration whenever you commit changes, and you can restore any of those versions.

- Revert to the current running configuration (file named running-config.xml).

This operation undoes changes you made to the candidate configuration since the last commit.

To revert all the changes that all administrators made, perform one of the following steps:

- Select **Device > Setup > Operations, Revert to running configuration**, and click **Yes** to confirm the operation.
- Log in to the firewall with an administrative account that is assigned the Superuser role or an [Admin Role profile](#) with the **Commit For Other Admins** privilege enabled. Then select **Config > Revert Changes** at the top of the web interface, select **Revert All Changes** and **Revert**.

To revert only specific changes to the candidate configuration:

1. Log in to the firewall with an administrative account that has the [role privileges](#) required to revert the desired changes.



The privileges that control commit operations also control revert operations.

2. Select **Config > Revert Changes** at the top of the web interface.
3. Select **Revert Changes Made By**.
4. To filter the Revert Scope by administrator, click **<administrator-name>**, select the administrators, and click **OK**.
5. To filter the Revert Scope by location, clear any locations that you want to exclude.
6. **Revert** the changes.

- Revert to the default snapshot of the candidate configuration.

This is the snapshot that you create or overwrite when you click **Config > Save Changes** at the top 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.

- Revert to 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.

- Revert to one of the following:

- 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.

- Revert to 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, 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. Make sure you are following the [Best Practices for Securing Administrative Access](#) to ensure that you are securing administrative access to your firewalls and other security devices in a way that prevents successful attacks.

- [Administrative Role Types](#)
- [Configure an Admin Role Profile](#)
- [Administrative Authentication](#)
- [Configure Administrative Accounts and Authentication](#)

Administrative Role Types

A *role* defines the type of access that an administrator has to the firewall. The Administrator Types are:

- **Role Based**—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 firewall with multiple virtual systems, you can select whether the role defines access for all virtual systems or specific virtual systems. 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](#).
- **Dynamic**—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.
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).
Virtual system administrator	Access to selected virtual systems on the firewall to create and manage specific aspects of virtual systems. A virtual system

Dynamic Role	Privileges
	administrator doesn't have access to network interfaces, VLANs, virtual wires, virtual routers, IPSec tunnels, GRE tunnels, DHCP, DNS Proxy, QoS, LLDP, or network profiles.
Virtual system administrator (read-only)	Read-only access to selected virtual systems on the firewall and specific aspects of virtual systems. A virtual system administrator with read-only access doesn't have access to network interfaces, VLANs, virtual wires, virtual routers, IPSec tunnels, GRE tunnels, DHCP, DNS Proxy, QoS, LLDP, or network profiles.

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.

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 a Firewall Administrator Account](#).

Administrative Authentication

You can configure the following types of authentication and authorization (role and access domain assignment) for firewall administrators:

Authentication Method	Authorization Method	Description
Local	Local	The administrative account credentials and authentication mechanisms are local to the firewall. You can define the accounts with or without a user database that is local to the firewall—see Local Authentication for the advantages and disadvantages of using a local database. You

Authentication Method	Authorization Method	Description
		use the firewall to manage role assignments but access domains are not supported. For details, see Configure Local or External Authentication for Firewall Administrators .
SSH Keys	Local	The administrative accounts are local to the firewall, but authentication to the CLI is based on SSH keys. You use the firewall to manage role assignments but access domains are not supported. For details, see Configure SSH Key-Based Administrator Authentication to the CLI .
Certificates	Local	The administrative accounts are local to the firewall, but authentication to the web interface is based on client certificates. You use the firewall to manage role assignments but access domains are not supported. For details, see Configure Certificate-Based Administrator Authentication to the Web Interface .
External service	Local	The administrative accounts you define locally on the firewall serve as references to the accounts defined on an external Multi-Factor Authentication , SAML , Kerberos , TACACS+ , RADIUS , or LDAP server. The external server performs authentication. You use the firewall to manage role assignments but access domains are not supported. For details, see Configure Local or External Authentication for Firewall Administrators .
External service	External service	The administrative accounts are defined on an external SAML , TACACS+ , or RADIUS server. The server performs both authentication and authorization. For authorization, you define Vendor-Specific Attributes (VSAs) on the TACACS+ or RADIUS server, or SAML attributes on the SAML server. PAN-OS maps the attributes to administrator roles, access domains, user groups, and virtual systems that you define on the firewall. For details, see: • Configure SAML Authentication • Configure TACACS+ Authentication • Configure RADIUS Authentication

Configure Administrative Accounts and Authentication

If you have already configured an authentication profile (see [Configure an Authentication Profile and Sequence](#)) or you don't require one to authenticate administrators, you are ready to [Configure a Firewall Administrator Account](#). Otherwise, perform one of the other procedures listed below to configure administrative accounts for specific types of authentication.

- [Configure a Firewall Administrator Account](#)
- [Configure Local or External Authentication for Firewall Administrators](#)
- [Configure Certificate-Based Administrator Authentication to the Web Interface](#)
- [Configure SSH Key-Based Administrator Authentication to the CLI](#)
- [Configure API Key Lifetime](#)

Configure a Firewall Administrator Account

Administrative accounts specify [roles](#) and authentication methods for firewall administrators. The service that you use to assign roles and perform authentication determines whether you add the accounts on the firewall, on an external server, or both (see [Administrative Authentication](#)). If the authentication method relies on a local firewall database or an external service, you must configure an authentication profile before adding an administrative account (see [Configure Administrative Accounts and Authentication](#)). If you already configured the authentication profile or you will use [Local Authentication](#) without a firewall database, perform the following steps to add an administrative account on the firewall.



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.

Make sure you are following the [Best Practices for Securing Administrative Access](#) to ensure that you are securing administrative access to your firewalls and other security devices in a way that prevents successful attacks.

STEP 1 | Select **Device > Administrators** and **Add** an account.

STEP 2 | Enter a user **Name**.

If the firewall uses a local user database to authenticate the account, enter the name that you specified for the account in the database (see [Add the user group to the local database](#).)

STEP 3 | Select an **Authentication Profile** or sequence if you [configured either](#) for the administrator.

If the firewall uses [Local Authentication](#) without a local user database for the account, select **None** (default) and enter a **Password**.

STEP 4 | Select the **Administrator Type**.

If you configured a [custom](#) role for the user, select **Role Based** and select the Admin Role **Profile**. Otherwise, select **Dynamic** (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.

STEP 5 | ([Optional](#)) Select a **Password Profile** for administrators that the firewall authenticates locally without a local user database. For details, see [Define a Password Profile](#).

STEP 6 | Click **OK** and **Commit**.

Configure Local or External Authentication for Firewall Administrators

You can use [Local Authentication](#) and [External Authentication Services](#) to authenticate administrators who access the firewall. These authentication methods prompt administrators to respond to one or more authentication challenges, such as a login page for entering a username and password.



If you use an external service to manage both authentication and authorization (role and access domain assignments), see:

- [Configure SAML Authentication](#)
- [Configure TACACS+ Authentication](#)
- [Configure RADIUS Authentication](#)

To authenticate administrators without a challenge-response mechanism, you can [Configure Certificate-Based Administrator Authentication to the Web Interface](#) and [Configure SSH Key-Based Administrator Authentication to the CLI](#).

STEP 1 | ([External authentication only](#)) Enable the firewall to connect to an external server for authenticating administrators.

Configure a server profile:

- [Add a RADIUS server profile.](#)

If the firewall integrates with a [Multi-Factor Authentication](#) (MFA) service through RADIUS, you must add a RADIUS server profile. In this case, the MFA service provides all the authentication factors (challenges). If the firewall integrates with an MFA service through a vendor API, you can still use a RADIUS server profile for the first factor but MFA server profiles are required for additional factors.

- [Add an MFA server profile.](#)
- [Add a TACACS+ server profile.](#)
- [Add a SAML IdP server profile.](#) You cannot combine [Kerberos](#) single sign-on (SSO) with [SAML](#) SSO; you can use only one type of SSO service.
- [Add a Kerberos server profile.](#)
- [Add an LDAP server profile.](#)

STEP 2 | ([Local database authentication only](#)) Configure a user database that is local to the firewall.

1. [Add the user account to the local database.](#)
2. ([Optional](#)) [Add the user group to the local database.](#)

STEP 3 | ([Local authentication only](#)) Define password complexity and expiration settings.

These settings help protect the firewall against unauthorized access by making it harder for attackers to guess passwords.

1. Define global password complexity and expiration settings for all local administrators. The settings don't apply to local database accounts for which you specified a password hash instead of a password (see [Local Authentication](#)).
 1. Select **Device > Setup > Management** and edit the Minimum Password Complexity settings.
 2. Select **Enabled**.
 3. Define the password settings and click **OK**.
2. Define a Password Profile.

You assign the profile to administrator accounts for which you want to override the global password expiration settings. The profiles are available only to accounts that are not associated with a local database (see [Local Authentication](#)).

1. Select **Device > Password Profiles** and **Add** a profile.
2. Enter a **Name** to identify the profile.
3. Define the password expiration settings and click **OK**.

STEP 4 | ([Kerberos SSO only](#)) [Create a Kerberos keytab.](#)

A keytab is a file that contains Kerberos account information for the firewall. To support Kerberos SSO, your network must have a [Kerberos](#) infrastructure.

STEP 5 | Configure an authentication profile.



If your administrative accounts are stored across multiple types of servers, you can create an authentication profile for each type and add all the profiles to an authentication sequence.

Configure an Authentication Profile and Sequence. In the authentication profile, specify the **Type** of authentication service and related settings:

- **External service**—Select the **Type** of external service and select the **Server Profile** you created for it.
- **Local database authentication**—Set the **Type** to **Local Database**.
- **Local authentication without a database**—Set the **Type** to **None**.
- **Kerberos SSO**—Specify the **Kerberos Realm** and **Import** the **Kerberos Keytab**.

STEP 6 | Assign the authentication profile or sequence to an administrator account.

1. **Configure a Firewall Administrator Account.**
 - Assign the **Authentication Profile** or sequence that you configured.
 - (**Local database authentication only**) Specify the **Name** of the user account you added to the local database.
2. **Commit** your changes.
3. (**Optional**) **Test Authentication Server Connectivity** to verify that the firewall can use the authentication profile to authenticate administrators.

Configure Certificate-Based Administrator Authentication to the Web Interface

As a more secure alternative to password-based authentication to the firewall web interface, 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.

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 or a third-party CA.

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 a Firewall Administrator 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.

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 a Firewall Administrator 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 API Key Lifetime

The API keys on the firewall and Panorama enable you to authenticate API calls to the XML API and REST API. Because these keys grant access to the firewall and Panorama that are critical elements of your security posture, as a best practice, specify an API key lifetime to enforce regular key rotation. After you specify the key lifetime, when you regenerate an API key, each key is unique.

In addition to setting a key lifetime that prompts you to regenerate new keys periodically, you can also revoke all currently valid API keys in the event one or more keys are compromised. Revoking keys is a way to expire all currently valid keys.

STEP 1 | Select **Device > Setup > Management**.

STEP 2 | Edit Authentication Settings to specify the **API Key Lifetime (min)**.



Set the API key lifetime to protect against compromise and to reduce the effects of an accidental exposure. By default, the API key lifetime is set to 0, which means that the keys will never expire. To ensure that your keys are frequently rotated and each key is unique when regenerated, you must specify a validity period that ranges between 1–525600 minutes. Refer to the audit and compliance policies for your enterprise to determine how you should specify the lifetime for which your API keys are valid.

STEP 3 | **Commit** the changes.

STEP 4 | (To revoke all API keys) Select **Expire all API Keys** to reset currently valid API keys.

If you have just set a key lifetime and want to reset all API keys to adhere to the new term, you can expire all existing keys.



On confirmation, the keys are revoked and you can view the timestamp for when the **API Keys Last Expired**.

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 a Firewall Administrator 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](#)
- [Provide Granular Access to the Panorama Tab](#)

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.  <i>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.</i>	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
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	Firewall: Yes	Yes	No	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	Match logs are available only if you have a GlobalProtect license (subscription).	Panorama: Yes Device Group/Template: Yes			
User-ID	Specifies whether the administrator can see the User-ID logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
GTP	Specifies whether the mobile network operator can see GTP logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Tunnel Inspection	Specifies whether the administrator can see the Tunnel Inspection logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
SCTP	Specifies whether the mobile network operator can see Stream Control Transmission Protocol (SCTP) logs.  <i>You must enable SCTP on Panorama (Device > Setup > Management) before you can control Administrator access to SCTP logs, custom reports, or predefined reports for Panorama and Device Group/Template.</i>	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	Yes	No	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
		Device Group/Template: No			
Alarms	Specifies whether the administrator can see system-generated alarms.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Authentication	Specifies whether the administrator can see the Authentication logs.	Firewall: Yes Panorama: Yes Device Group/Template: No	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
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

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
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
Block IP List	Specifies whether the administrator can view the block list (Enable or Read Only) and delete entries from the list (Enable). If you disable the setting, the administrator won't be able to view or delete entries from the block list.	Firewall: Yes Panorama: under Context Switch UI: Yes Template: Yes	Yes	Yes	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
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	Firewall: Yes Panorama: Yes	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	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.	Device Group/Template: 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
Email Scheduler	Specifies whether the administrator can schedule report groups for email. Because the generated reports that get emailed may	Firewall: Yes Panorama: Yes	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	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.	Device Group/Template: 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.  <i>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.</i>	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	Yes	No	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
		Device Group/Template: 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
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

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
GTP Log	Specifies whether the mobile network operator can create a custom report that includes data from GTP logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
GTP Summary	Specifies whether the mobile network operator can create a custom report that includes data from GTP logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Tunnel Log	Specifies whether the administrator can create a custom report that includes data from tunnel inspection logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Tunnel Summary	Specifies whether the administrator can create a custom report that includes data from the Tunnel Summary database.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
SCTP Log	Specifies whether the mobile network operator can create a custom report that includes data from SCTP logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
SCTP Summary	Specifies whether the mobile network operator can create a custom report that includes data from the SCTP Summary database.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Userid	Specifies whether the administrator can create a custom report that includes data from the User-ID logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Auth	Specifies whether the administrator can create a custom report that includes data from the Authentication logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
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
View Predefined GTP Reports	Specifies whether the mobile network operator can view GTP 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	Specifies whether the mobile network operator can view SCTP	Firewall: Yes Panorama: Yes	Yes	No	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
SCTP Reports	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.	Device Group/Template: 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 username or IP address) must be explicitly defined, privacy settings that disable the ability to see full IP addresses or usernames 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	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	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.			
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
Tunnel Inspection	Enable this privilege to allow the administrator to view, add, and/or delete Tunnel Inspection 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 Tunnel Inspection rulebase, disable this privilege.	Yes	Yes	Yes
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
Authentication	Enable this privilege to allow the administrator to view, add, and/or delete Authentication 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 Authentication 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
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	Yes	No	Yes

Access Level	Description	Enable	Read Only	Disable
	granular control by enabling the GlobalProtect privilege and restricting HIP Object or HIP Profile access.			
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
Clientless Apps	Specifies whether the administrator can view, add, modify, or delete GlobalProtect VPN Clientless applications.	Yes	Yes	Yes
Clientless App Groups	Specifies whether the administrator can view, add, modify, or delete GlobalProtect VPN Clientless application groups.	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
External Dynamic Lists	Specifies whether the administrator can view, add, or delete external dynamic 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

Access Level	Description	Enable	Read Only	Disable
URL Category	Specifies whether the administrator can view, add, or delete custom URL categories for use in policy.	Yes	Yes	Yes
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
WildFire Analysis	Specifies whether the administrator can view, add, or delete WildFire analysis 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
GTP Protection	Specifies whether the mobile network operator can view, add, or delete GTP Protection profiles.	Yes	Yes	Yes
SCTP Protection	Specifies whether the mobile network operator can view, add, or delete Stream Control Transmission Protocol (SCTP) 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

Access Level	Description	Enable	Read Only	Disable
Authentication	Specifies whether the administrator can view, add, or delete authentication enforcement objects.	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
GRE Tunnels	Specifies whether the administrator can view, add, modify, or delete GRE Tunnel configurations.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
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
Clientless Apps	Specifies whether the administrator can view, add, modify, or delete GlobalProtect Clientless VPN applications.	Yes	Yes	Yes
Clientless App Groups	Specifies whether the administrator can view, add, modify, or delete GlobalProtect Clientless VPN application groups.	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
GlobalProtect IPsec Crypto	Controls access to the Network Profiles > GlobalProtect IPsec Crypto node.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	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.			
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
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	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	profile configuration but cannot add or edit a configuration.			
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

Access Level	Description	Enable	Read Only	Disable
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

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 connections, 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 and Telemetry and Threat Intelligence nodes. If you disable this privilege, the administrator cannot: Load firewall configurations.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	- Save or revert the firewall configuration.  <i>This privilege applies only to the Device > Operations options. The Save and Commit privileges control whether the administrator can save or revert configurations through the Config > Save and Config > Revert options.</i> - Create custom logos. - Configure SNMP monitoring of firewall settings. - Configure the Statistics Service feature. - Configure Telemetry and Threat Intelligence settings. Only administrators with the predefined Superuser role can export or import firewall configurations and shut down the firewall. Only administrators with the predefined Superuser or Device Administrator role can reboot the firewall or restart the dataplane. Administrators with a role that allows access only to specific virtual systems cannot load, save, or revert firewall configurations through the Device > Operations options.			
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.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	If the privilege state is set to read-only, you can view the current configuration but cannot make any changes.			
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
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
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
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	No	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	information for their own administrator account. They will not see any information about other administrator accounts configured on the firewall.			
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
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 RADIUS, TACACS+, LDAP, Kerberos, SAML, multi-factor authentication (MFA), or local database authentication settings. PAN-OS uses authentication profiles to authenticate firewall administrators and Captive Portal or GlobalProtect end users. If you set this privilege to read-only, the administrator can view the Authentication Profile information but cannot create or edit authentication profiles.	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
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	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	virtual systems but cannot add or edit a configuration.			
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, Connection Security, User-ID Agents, 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.  <i>This privilege is not available to Device Group and Template administrators.</i>	Yes	Yes	Yes
Certificate Management	Sets the default state to enable or disable for all of the Certificate settings described below.	Yes	No	Yes
Certificates	Controls access to the Certificates node. If you disable this privilege, the administrator	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	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.			
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	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	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.			
SSL Decryption Exclusion	Controls access to the SSL Decryption Exclusion node. If you disable this privilege, the administrator will not see the node or be able to see the SSL decryption add custom exclusions. If you set this privilege to read-only, the administrator can view existing SSL decryption exceptions but cannot create or edit them.	Yes	Yes	Yes
Response Pages	Controls access to the Response Pages node. If you disable this privilege, the administrator will not see the Response Page node or be able to define a custom HTML message that is downloaded and displayed instead of a requested web page or file. 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 cannot see the Log Settings > System node or specify which System logs the firewall forwards to Panorama or external services (such as a syslog server). If you set this privilege to read-only, the administrator can view the Log Settings > System settings for the firewall but cannot add, edit, or delete the settings.	Yes	Yes	Yes
Configuration	Controls access to the Log Settings > Configuration node. If you disable this privilege, the administrator cannot see the Log Settings > Configuration node or specify which Configuration logs the firewall forwards to Panorama or external services (such as a syslog server).	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	If you set this privilege to read-only, the administrator can view the Log Settings > Configuration settings for the firewall but cannot add, edit, or delete the settings.			
User-ID	Controls access to the Log Settings > User-ID node. If you disable this privilege, the administrator cannot see the Log Settings > User-ID node or specify which User-ID logs the firewall forwards to Panorama or external services (such as a syslog server). If you set this privilege to read-only, the administrator can view the Log Settings > User-ID settings for the firewall but cannot add, edit, or delete the settings.	Yes	Yes	Yes
HIP Match	Controls access to the Log Settings > HIP Match node. If you disable this privilege, the administrator cannot see the Log Settings > HIP Match node or specify which Host Information Profile (HIP) match logs the firewall forwards to Panorama or external services (such as a syslog server). HIP match logs provide information on Security policy rules that apply to GlobalProtect endpoints If you set this privilege to read-only, the administrator can view the Log Settings > HIP settings for the firewall but cannot add, edit, or delete the settings.	Yes	Yes	Yes
Correlation	Controls access to the Log Settings > Correlation node. If you disable this privilege, the administrator cannot see the Log Settings > Correlation node or add, delete, or modify correlation log forwarding settings or tag source or destination IP addresses. If you set this privilege to read-only, the administrator can view the Log Settings > Correlation settings for the firewall but cannot add, edit, or delete the settings.	Yes	Yes	Yes
Alarm Settings	Controls access to the Log Settings > Alarm Settings node. If you disable this privilege, the administrator cannot see the Log Settings > Alarm Settings node or configure notifications that the firewall generates when a Security policy rule (or group of rules) is hit repeatedly within a configurable time period. If you set this privilege to read-only, the administrator can view the Log Settings >	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	Alarm Settings for the firewall but cannot edit the settings.			
Manage Logs	Controls access to the Log Settings > Manage Logs node. If you disable this privilege, the administrator cannot see the Log Settings > Manage Logs node or 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 an email server profile.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
HTTP	Controls access to the Server Profiles > HTTP node. If you disable this privilege, the administrator will not see the Server Profiles > HTTP node or be able to configure an HTTP server profile that can be used to enable log forwarding to HTTP destinations any log entries. If you set this privilege to read-only, the administrator can view the Server Profiles > HTTP information but cannot configure an HTTP server profile.	Yes	Yes	Yes
Netflow	Controls access to the Server Profiles > Netflow node. If you disable this privilege, the administrator will not see the Server Profiles > Netflow node or be able to define a NetFlow server profile, which specifies the frequency of the export along with the NetFlow servers that will receive the exported data. 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	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	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.			
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
SAML Identity Provider	Controls access to the Server Profiles > SAML Identity Provider node. If you disable this privilege, the administrator cannot see the node or configure SAML identity provider (IdP) server profiles. If you set this privilege to read-only, the administrator can view the Server Profiles > SAML Identity Provider information but cannot configure SAML IdP server profiles.	Yes	Yes	Yes
Multi Factor Authentication	Controls access to the Server Profiles > Multi Factor Authentication node. If you disable this privilege, the administrator cannot see the node or configure multi-factor authentication (MFA) server profiles. If you set this privilege to read-only, the administrator can view the Server Profiles > SAML Identity Provider information but cannot configure MFA server profiles.			
Local User Database	Sets the default state to enable or disable for all of the Local User Database settings described below.	Yes	No	Yes
Users	Controls access to the Local User Database > Users node. If you disable this privilege, the administrator will not see the Local User Database > Users node or set up a local database on the firewall to store authentication information for remote access users, firewall administrators, and Captive Portal users.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	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.			
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
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
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

Access Level	Description	Enable	Read Only	Disable
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 app. If you set this privilege to read-only, the administrator can view the available GlobalProtect Client releases but cannot download or install the app 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
Support	Controls access to the Support node. If you disable this privilege, the administrator cannot see the Support node, activate support, or access production and security alerts from Palo Alto Networks. If you set this privilege to read-only, the administrator can see the Support node and access production and security alerts but cannot activate support. Only administrators with the predefined Superuser role can use the Support node to generate tech support files or generate and download stats dump and core 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	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
	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.			

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.  <i>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.</i>	Yes	N/A	Yes
Show User Names in Logs and Reports	When disabled, usernames obtained by traffic running through the Palo Alto Networks firewall are not shown in logs or reports. Columns where the usernames would normally be displayed are empty.	Yes	N/A	Yes

Access Level	Description	Enable	Read Only	Disable
	 <i>Scheduled reports that are displayed in the interface through Monitor > Reports or reports that are sent via the email scheduler will still display usernames. 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.</i>			
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 revert), save, and validate functions when creating or editing an Admin Role profile (**Device > Admin Roles**), scroll down to the **Commit**, **Save**, and **Validate** options on the **WebUI** tab.

Access Level	Description	Enable	Read Only	Disable
Commit	Sets the default state to enabled or disabled for all of the commit and revert privileges described below.	Yes	N/A	Yes
Device	When disabled, an administrator cannot commit or revert changes that any administrator made to the firewall configuration, including his or her own changes.	Yes	N/A	Yes
Commit For Other Admins	When disabled, an administrator cannot commit or revert changes that other administrators made to the firewall configuration.	Yes	N/A	Yes
Save	Sets the default state to enabled or disabled for all of the save operation privileges described below.	Yes	N/A	Yes
Partial save	When disabled, an administrator cannot save changes that any administrator made to the	Yes	N/A	Yes

Access Level	Description	Enable	Read Only	Disable
	firewall configuration, including his or her own changes.			
Save For Other Admins	When disabled, an administrator cannot save changes that other administrators made to the firewall 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, including Management, Operations and Telemetry, Services, Content-ID, WildFire, Session, or HSM. If you set the privilege to: read-only, the administrator can see the information but cannot edit it.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	disable this privilege, the administrator cannot see or edit the information.				
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
Admin Roles	Specifies whether the administrator can view Panorama administrator roles.	Panorama: Yes Device Group/Template: No	No	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	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.				
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  <i>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.</i>	Yes	Yes	Yes
Authentication Profile	Specifies whether the administrator can view, add, edit, delete, or clone	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	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.				
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
User Identification	Specifies whether the administrator can configure User-ID connection security and view, add, edit, or delete User-ID redistribution points (such as User-ID agents). If you set this privilege to read-only, the administrator can view settings for User-ID connection security and redistribution points but can't manage the settings. If you disable this privilege, the administrator can't see or manage settings for User-ID connection security or redistribution points.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Managed Devices	Specifies whether the administrator can view, add, edit, 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,	Panorama: Yes Device Group/Template: Yes	Yes (No for Device Group and Template roles)	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	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.  <i>An administrator with Device Deployment privileges can still select Panorama > Device Deployment to install updates on managed firewalls.</i>				
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  <i>Device Group and Template administrators can see only the templates and stacks that are within the access domains assigned to those administrators.</i>	Yes (No for Device Group and Template admins)	Yes	Yes
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  <i>Device Group and Template administrators can access only the</i>	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
		<i>device groups that are within the access domains assigned to those administrators.</i>			
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.  <i>An administrator with Device Deployment privileges can still use the Panorama > Device Deployment options to install updates on managed collectors.</i>	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

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
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
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

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
SSL/TLS Service Profile	Specifies whether the administrator can view, add, edit, delete or clone SSL/TLS Service profiles. 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
System	Specifies whether the administrator can see and configure the settings that control the forwarding of System logs to external services (syslog, email, SNMP trap, or HTTP 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.  <i>This privilege pertains only to System logs that Panorama and Log Collectors generate. The Collector Groups privilege (Panorama > Collector Groups) controls forwarding for System logs that Log Collectors receive from firewalls. The Device > Log Settings > System privilege controls</i>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	<i>log forwarding from firewalls directly to external services (without aggregation on Log Collectors).</i>				
Config	Specifies whether the administrator can see and configure the settings that control the forwarding of Config logs to external services (syslog, email, SNMP trap, or HTTP 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.  <i>This privilege pertains only to Config logs that Panorama and Log Collectors generate. The Collector Groups privilege (Panorama > Collector Groups) controls forwarding for Config logs that Log Collectors receive from firewalls. The Device > Log Settings > Configuration privilege controls log forwarding from firewalls directly to external services (without aggregation on Log Collectors).</i>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
User-ID	Specifies whether the administrator can see and configure the settings that control	Panorama: Yes	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	the forwarding of User-ID logs to external services (syslog, email, SNMP trap, or HTTP 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.  <i>This privilege pertains only to User-ID logs that Panorama generates. The Collector Groups privilege (Panorama > Collector Groups) controls forwarding for User-ID logs that Log Collectors receive from firewalls. The Device > Log Settings > User-ID privilege controls log forwarding from firewalls directly to external services (without aggregation on Log Collectors).</i>	Device Group/Template: No			
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 in Legacy mode to external services (syslog, email, SNMP trap, or HTTP 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.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	If you disable this privilege, the administrator can't see or manage the settings.  <i>The Collector Groups privilege (Panorama > Collector Groups) controls forwarding for HIP Match logs that Log Collectors receive from firewalls. The Device > Log Settings > HIP Match privilege controls log forwarding from firewalls directly to external services (without aggregation on Log Collectors).</i>				
Correlation	Specifies whether the administrator can see and configure the settings that control the forwarding of Correlation logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP 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.  <i>The Collector Groups privilege (Panorama > Collector Groups) controls forwarding of Correlation logs from a Panorama M-Series appliance or Panorama</i>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	<i>virtual appliance in Panorama mode.</i>				
Traffic	Specifies whether the administrator can see and configure the settings that control the forwarding of Traffic logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP 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 Collector Groups privilege (<i>Panorama > Collector Groups</i>) controls forwarding for Traffic logs that Log Collectors receive from firewalls. The Log Forwarding privilege (<i>Objects > Log Forwarding</i>) controls forwarding from firewalls directly to external services (without aggregation on Log Collectors).	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 in Legacy mode to external services (syslog, email, SNMP trap, or HTTP servers). If you set this privilege to read-only, the administrator can see the	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	forwarding settings of Threat logs but can't manage them. If you disable this privilege, the administrator can't see or manage the settings.  <i>The Collector Groups privilege (Panorama > Collector Groups) controls forwarding for Threat logs that Log Collectors receive from firewalls. The Log Forwarding privilege (Objects > Log Forwarding) controls forwarding from firewalls directly to external services (without aggregation on Log Collectors).</i>				
WildFire	Specifies whether the administrator can see and configure the settings that control the forwarding of WildFire logs from a Panorama virtual appliance in Legacy mode to external services (syslog, email, SNMP trap, or HTTP 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.  <i>The Collector Groups privilege (Panorama > Collector Groups) controls the forwarding for WildFire logs that</i>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	<i>Log Collectors receive from firewalls. The Log Forwarding privilege (Objects > Log Forwarding) controls forwarding from firewalls directly to external services (without aggregation on Log Collectors).</i>				
Server Profiles	Sets the default state, enabled or disabled, for all the server profile privileges.  <i>These privileges pertain only to the server profiles that are used for forwarding logs from Panorama or Log Collectors and the server profiles that are used for authenticating Panorama administrators. The Device > Server Profiles privileges control access to the server profiles that are used for forwarding logs directly from firewalls to external services and for authenticating firewall administrators.</i>	Panorama: Yes Device Group/Template: No	Yes	No	Yes
SNMP Trap	Specifies whether the administrator can see and configure SNMP trap server profiles.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	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.				
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	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	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.				
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
SAML Identity Provider	Specifies whether the administrator can see and configure the SAML Identity Provider (IdP) server profiles that are used to authenticate Panorama administrators. If you set this privilege to read-only, the administrator can see the	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	SAML IdP server profiles but can't manage them. If you disable this privilege, the administrator can't see or manage the SAML IdP server profiles.				
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
Software	Specifies whether the administrator can: view information about software updates installed on the Panorama management server; 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.  <i>The Panorama > Device Deployment > Software privilege controls access to PAN-OS software deployed on firewalls and Panorama software deployed</i>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	<i>on Dedicated Log Collectors.</i>				
Dynamic Updates	Specifies whether the administrator can: view information about content updates installed on the Panorama management server (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.  <i>The Panorama > Device Deployment > Dynamic Updates privilege controls access to content updates deployed on firewalls and Dedicated Log Collectors.</i>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Support	Specifies whether the administrator can: view Panorama support license information, product alerts, and security alerts; activate a support license, and manage cases. Only a superuser admin can generate Tech Support files. 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	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	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.				
Device Deployment	Sets the default state, enabled or disabled, for all the privileges associated with deploying licenses and software or content updates to firewalls and Log Collectors.  <i>The Panorama > Software and Panorama > Dynamic Updates privileges control the software and content updates installed on a Panorama management server.</i>	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
GlobalProtect Client	Specifies whether the administrator can: view information about GlobalProtect 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 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 app software updates, see the associated release notes, or activate the updates on firewalls.	Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes
Dynamic Updates	Specifies whether the administrator can: view information about the content updates (for example, Applications updates) installed on firewalls and Dedicated Log Collectors; download, upload, or install the updates; and view the associated release notes. 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.	Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes

Access Level	Description	Administrator Role Availability	Enable	Read Only	Disable
	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.				
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	Yes	No	Yes

Access Level	Description	Enable	Read Only	Disable
	and/or the Show User Names In Logs And Reports option.			
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
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

Access Level	Description	Enable	Read Only	Disable
Device	Controls access to the Device tab. If you disable this privilege, the administrator will not see the Device tab and will not have access to any 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.  <i>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.</i>	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
Save	Sets the default state (enabled or disabled) for all the save privileges described below (Partial Save and Save For Other Admins).	Yes	No	Yes
• Partial Save	When disabled, an administrator cannot save changes that any administrator made to the Panorama configuration.	Yes	No	Yes
• Save For Other Admins	When disabled, an administrator cannot save changes that other administrators made to the Panorama configuration.	Yes	No	Yes

Access Level	Description	Enable	Read Only	Disable
Commit	Sets the default state (enabled or disabled) for all the commit, push, and revert privileges described below (Panorama, Device Groups, Templates, Force Template Values, Collector Groups, WildFire Appliance Clusters).	Yes	No	Yes
Panorama	When disabled, an administrator cannot commit or revert configuration changes that any administrators made, including his or her own changes.	Yes	No	Yes
Commit for Other Admins	When disabled, an administrator cannot commit or revert configuration changes that other administrators made.	Yes	No	Yes
Device Groups	When disabled, an administrator cannot push changes to device groups.	Yes	No	Yes
Templates	When disabled, an administrator cannot push changes to templates.	Yes	No	Yes
Force Template Values	This privilege controls access to the Force Template Values option in the Push Scope Selection dialog. When disabled, an administrator cannot replace overridden settings in local firewall configurations with settings that Panorama pushes from a template.  <i>If you push a configuration with Force Template Values enabled, all overridden values on the firewall are replaced with values from the template. Before you use this option, check for overridden values on the firewalls to ensure your commit does not result in any unexpected network outages or issues caused by replacing those overridden values.</i>	Yes	No	Yes
Collector Groups	When disabled, an administrator cannot push changes to Collector Groups.	Yes	No	Yes
WildFire Appliance Clusters	When disabled, an administrator cannot push changes to WildFire appliance clusters.	Yes	No	Yes

Access Level	Description	Enable	Read Only	Disable
Tasks	When disabled, an administrator cannot access the Task Manager.	Yes	No	Yes
Global	Controls access to the global settings (system alarms) described in Provide Granular Access to Global Settings .	Yes	No	Yes

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 updates when you Enable VM Monitoring to Track Changes on the Virtual Network. 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.  <i>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.</i>
161	UDP	Port the firewall listens on for polling requests (GET messages) from the SNMP manager.
514	TCP	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.
514	UDP	
6514	SSL	

Destination Port	Protocol	Description
2055	UDP	Default port the firewall uses to send NetFlow records to a NetFlow collector if you Configure NetFlow Exports , but this is configurable.
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 TLS 1.2 TCP	Ports used for User-ID™ Captive Portal: 6080 for NT LAN Manager (NTLM) authentication, 6081 for Captive Portal without an SSL/TLS Server Profile, and 6082 for Captive Portal with an SSL/TLS Server Profile.
10443	SSL	Port that the firewall and Panorama use to provide contextual information about a threat or to seamlessly shift your threat investigation to the Threat Vault and AutoFocus.

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.

Destination Port	Protocol	Description
		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.
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.
28443	TCP	Used for managed devices (firewalls and Log Collectors) to retrieve software and content updates from Panorama.  <i>Only devices that run PAN-OS 8.x and later releases retrieve updates from Panorama over this port. For devices running earlier releases, Panorama pushes the update packages over port 3978.</i>
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.

Destination Port	Protocol	Description
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. Used for communication between Log Collectors in a Collector Group for log distribution.
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.
10443	SSL	Port that Panorama uses to provide contextual information about a threat or to seamlessly shift your threat investigation to the Threat Vault and AutoFocus.
23000 to 23999	TCP, UDP, or SSL	Used for Syslog communication between Panorama and the Traps ESM components.

Ports Used for GlobalProtect

GlobalProtect uses the following ports.

Destination Port	Protocol	Description
443	TCP	Used for communication between GlobalProtect apps and portals, or GlobalProtect apps and gateways and for SSL tunnel connections. GlobalProtect gateways also use this port to collect host information from GlobalProtect apps and perform host information profile (HIP) checks.
4501	UDP	Used for IPSec tunnel connections between GlobalProtect apps 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 6514	TCP UDP SSL	Port the User-ID agent listens on for authentication syslog messages if you Configure User-ID to Monitor Syslog Senders for User Mapping . The port depends on the type of agent and protocol: - PAN-OS integrated User-ID agent—Port 6514 for SSL and port 514 for UDP. - Windows-based User-ID agent—Port 514 for both TCP and UDP.
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 .

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).
5985	HTTP	Port the User-ID agent uses to monitor security logs and session information with the WinRM protocol over HTTP.
5986	HTTPS	Port the User-ID agent uses to monitor security logs and session information with the WinRM protocol over HTTPS.

Reset the Firewall to Factory Default Settings

Resetting the firewall to factory defaults will result in the loss of all configuration settings and logs.

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](#).

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:

Supported USB Flash Drives

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. To create this file, see [Create the init-cfg.txt file](#). 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)
<pre>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-vsyst, jumbo-frame dhcp-send-hostname=no dhcp-send-client-id=no dhcp-accept-server-hostname=no dhcp-accept-server-domain=no</pre>	<pre>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-vsyst, jumbo-frame dhcp-send-hostname=yes dhcp-send-client-id=yes dhcp-accept-server-hostname=yes dhcp-accept-server-domain=yes</pre>

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.

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.

Field	Description
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.
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-vsyt, 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 be running a PAN-OS 7.1.0 or later image 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.

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 > Devices > Register New Device > Register device using Serial Number or Authorization Code**.
2. Follow the steps to [Register the Firewall](#).
3. Click **Submit**.

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 > Devices** on the left-hand navigation pane.
2. For each device S/N you just registered, click the **Action** link (the pencil icon).
3. Under Activate Licenses, 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**.

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 firewall running a PAN-OS 7.1.0 or later image) 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
```

STEP 9 | Prepare the USB flash drive.

1. Insert the USB flash drive into the firewall that you used in the prior step.
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.

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.

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 dname) 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.

Authentication

Authentication is a method for protecting services and applications by verifying the identities of users so that only legitimate users have access. Several firewall and Panorama features require authentication. Administrators authenticate to access the web interface, CLI, or XML API of the firewall and Panorama. End users authenticate through Captive Portal or GlobalProtect to access various services and applications. You can choose from several authentication services to protect your network and to accommodate your existing security infrastructure while ensuring a smooth user experience.

If you have a public key infrastructure, you can deploy certificates to enable authentication without users having to manually respond to login challenges (see Certificate Management). Alternatively, or in addition to certificates, you can implement interactive authentication, which requires users to authenticate using one or more methods. The following topics describe how to implement, test, and troubleshoot the different types of interactive authentication:

- > Authentication Types
- > Plan Your Authentication Deployment
- > Configure Multi-Factor Authentication
- > Configure SAML Authentication
- > Configure Kerberos Single Sign-On
- > Configure Kerberos Server Authentication
- > Configure TACACS+ Authentication
- > Configure RADIUS Authentication
- > Configure LDAP Authentication
- > Connection Timeouts for Authentication Servers
- > Configure Local Database Authentication
- > Configure an Authentication Profile and Sequence
- > Test Authentication Server Connectivity
- > Authentication Policy
- > Troubleshoot Authentication Issues

Authentication Types

- [External Authentication Services](#)
- [Multi-Factor Authentication](#)
- [SAML](#)
- [Kerberos](#)
- [TACACS+](#)
- [RADIUS](#)
- [LDAP](#)
- [Local Authentication](#)

External Authentication Services

The firewall and Panorama can use external servers to control administrative access to the web interface and end user access to services or applications through Captive Portal and GlobalProtect. In this context, any authentication service that is not local to the firewall or Panorama is considered external, regardless of whether the service is internal (such as Kerberos) or external (such as a SAML identity provider) relative to your network. The server types that the firewall and Panorama can integrate with include [Multi-Factor Authentication](#) (MFA), [SAML](#), [Kerberos](#), [TACACS+](#), [RADIUS](#), and [LDAP](#). Although you can also use the [Local Authentication](#) services that the firewall and Panorama support, usually external services are preferable because they provide:

- Central management of all user accounts in an external identity store. All the supported external services provide this option for end users and administrators.
- Central management of account authorization (role and access domain assignments). SAML, TACACS+, and RADIUS support this option for administrators.
- Single sign-on (SSO), which enables users to authenticate only once for access to multiple services and applications. SAML and Kerberos support SSO.
- Multiple authentication challenges of different types (factors) to protect your most sensitive services and applications. MFA services support this option.

Authentication through an external service requires a server profile that defines how the firewall connects to the service. You assign the server profile to authentication profiles, which define settings that you customize for each application and set of users. For example, you can configure one authentication profile for administrators who access the web interface and another profile for end users who access a GlobalProtect portal. For details, see [Configure an Authentication Profile and Sequence](#).

Multi-Factor Authentication

You can [Configure Multi-Factor Authentication](#) (MFA) to ensure that each user authenticates using multiple methods (factors) when accessing highly sensitive services and applications. For example, you can force users to enter a login password and then enter a verification code that they receive by phone before allowing access to important financial documents. This approach helps to prevent attackers from accessing every service and application in your network just by stealing passwords. Of course, not every service and application requires the same degree of protection, and MFA might not be necessary for less sensitive services and applications that users access frequently. To accommodate a variety of security needs, you can [Configure Authentication Policy](#) rules that trigger MFA or a single authentication factor (such as login credentials or certificates) based on specific services, applications, and end users.

When choosing how many and which types of authentication factors to enforce, it's important to understand how policy evaluation affects the user experience. When a user requests a service or application, the firewall first evaluates Authentication policy. If the request matches an Authentication policy rule with MFA enabled, the firewall displays a Captive Portal web form so that users can authenticate

for the first factor. If authentication succeeds, the firewall displays an MFA login page for each additional factor. Some MFA services prompt the user to choose one factor out of two to four, which is useful when some factors are unavailable. If authentication succeeds for all factors, the firewall evaluates [Security policy](#) for the requested service or application.



To reduce the frequency of authentication challenges that interrupt the user workflow, you can configure the first factor to use [Kerberos](#) or [SAML single sign-on \(SSO\)](#) but not [NT LAN Manager \(NTLM\)](#) authentication.

To implement MFA for GlobalProtect, refer to [Configure GlobalProtect](#) to facilitate multi-factor authentication notifications.

You cannot use MFA authentication profiles in authentication sequences.

For end-user authentication via [Authentication Policy](#), the firewall directly [integrates](#) with several MFA platforms (Duo v2, [Okta Adaptive](#), PingID, and [RSA SecurID](#)), as well as integrating through RADIUS or SAML for all other MFA platforms. For remote user authentication to GlobalProtect portals and gateways and for administrator authentication to the Panorama and PAN-OS web interface, the firewall integrates with MFA vendors using RADIUS and SAML only.

The firewall supports the following MFA factors:

Factor	Description
Push	An endpoint device (such as a phone or tablet) prompts the user to allow or deny authentication.
Short message service (SMS)	An SMS message on the endpoint device prompts the user to allow or deny authentication. In some cases, the endpoint device provides a code that the user must enter in the MFA login page.
Voice	An automated phone call prompts the user to authenticate by pressing a key on the phone or entering a code in the MFA login page.
One-time password (OTP)	An endpoint device provides an automatically generated alphanumeric string, which the user enters in the MFA login page to enable authentication for a single transaction or session.

SAML

You can use Security Assertion Markup Language (SAML) 2.0 to authenticate administrators who access the firewall or Panorama web interface and end users who access web applications that are internal or external to your organization. In environments where each user accesses many applications and authenticating for each one would impede user productivity, you can configure SAML single sign-on (SSO) to enable one login to access multiple applications. Likewise, SAML single logout (SLO) enables a user to end sessions for multiple applications by logging out of just one session. SSO is available to administrators who access the web interface and to end users who access applications through GlobalProtect or Captive Portal. SLO is available to administrators and GlobalProtect end users, but not to Captive Portal end users. When you configure SAML authentication [on the firewall](#) or [on Panorama](#), you can specify SAML attributes for administrator authorization. SAML attributes enable you to quickly change the roles, access domains, and user groups of administrators through your directory service, which is often easier than reconfiguring settings on the firewall or Panorama.



Administrators cannot use SAML to authenticate to the CLI on the firewall or Panorama.

You cannot use SAML authentication profiles in authentication sequences.

SAML authentication requires a *service provider* (the firewall or Panorama), which controls access to applications, and an *identity provider* (IdP) such as PingFederate, which authenticates users. When a user requests a service or application, the firewall or Panorama intercepts the request and redirects the user to the IdP for authentication. The IdP then authenticates the user and returns a *SAML assertion*, which indicates authentication succeeded or failed. [SAML Authentication for Captive Portal End Users](#) illustrates SAML authentication for an end user who accesses applications through Captive Portal.




Figure 1: SAML Authentication for Captive Portal End Users

Kerberos

Kerberos is an authentication protocol that enables a secure exchange of information between parties over an insecure network using unique keys (called tickets) to identify the parties. The firewall and Panorama support two types of Kerberos authentication for administrators and end users:

- **Kerberos server authentication**—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. For the configuration steps, see [Configure Kerberos Server Authentication](#).
- **Kerberos single sign-on (SSO)**—A network that supports Kerberos V5 SSO prompts a user to log in only for initial access to the network (such as logging in to Microsoft Windows). After this initial login, the user can access any browser-based service in the network (such as 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 another external authentication service (such as a TACACS+ server), the firewall 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.

For the configuration steps, see [Configure Kerberos Single Sign-On](#).



Kerberos SSO is available only for services and applications that are internal to your Kerberos environment. To enable SSO for external services and applications, use [SAML](#).

TACACS+

Terminal Access Controller Access-Control System Plus (TACACS+) is a family of protocols that enable authentication and authorization through a centralized server. TACACS+ encrypts usernames and passwords, making it more secure than RADIUS, which encrypts only passwords. TACACS+ is also more reliable because it uses TCP, whereas RADIUS uses UDP. You can configure TACACS+ authentication for end users or administrators [on the firewall](#) and for administrators [on Panorama](#). Optionally, you can use TACACS+ Vendor-Specific Attributes (VSAs) to manage administrator authorization. TACACS+ VSAs enable you to quickly change the roles, access domains, and user groups of administrators through your directory service instead of reconfiguring settings on the firewall and Panorama.

The firewall and Panorama support the following TACACS+ attributes and VSAs. Refer to your TACACS+ server documentation for the steps to define these VSAs on the TACACS+ server.

Name	Value
service	This attribute is required to identify the VSAs as specific to Palo Alto Networks. You must set the value to PaloAlto .
protocol	This attribute is required to identify the VSAs as specific to Palo Alto Networks devices. You must set the value to firewall .
PaloAlto-Admin-Role	A default (dynamic) administrative role name or a custom administrative role name on the firewall.
PaloAlto-Admin-Access-Domain	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	A default (dynamic) administrative role name or a custom administrative role name on Panorama.
PaloAlto-Panorama-Admin-Access-Domain	The name of an access domain for Device Group and Template administrators (configured in the Panorama > Access Domains page).
PaloAlto-User-Group	The name of a user group in the Allow List of an authentication profile.

RADIUS

Remote Authentication Dial-In User Service (RADIUS) is a broadly supported networking protocol that provides centralized authentication and authorization. You can configure RADIUS authentication for end users or administrators [on the firewall](#) and for administrators [on Panorama](#). Optionally, you can use RADIUS Vendor-Specific Attributes (VSAs) to manage administrator authorization. RADIUS VSAs enable you to quickly change the roles, access domains, and user groups of administrators through your directory service instead of reconfiguring settings on the firewall and Panorama. You can also configure the firewall to use a RADIUS server for:

- [Collecting VSAs from GlobalProtect endpoints.](#)
- Implementing [Multi-Factor Authentication.](#)

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 (such as administrative access to the web interface) that initiates the authentication process.

The firewall and Panorama support the following RADIUS 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. Refer to your RADIUS server documentation for the steps to define these VSAs.

Alternatively, you can download the [Palo Alto Networks RADIUS dictionary](#), which defines the authentication attributes that the Palo Alto Networks firewall and a RADIUS server use to communicate with each other, and install it on your RADIUS server to map the attributes to the RADIUS binary data.



*When you predefine dynamic administrator roles for users on the server, use lower-case to specify the role (for example, enter **superuser**, not **SuperUser**).*



When configuring the advanced vendor options on a Cisco Secure Access Control Server (ACS), you must set both the Vendor Length Field Size and Vendor Type Field Size to 1. Otherwise, authentication will fail.

Name	Number	Value

VSAs for administrator account management and authentication

PaloAlto-Admin-Role	1	A default (dynamic) administrative role name or a custom administrative role name on the firewall.
PaloAlto-Admin-Access-Domain	2	The name of an access domain for firewall administrators (configured in the Device > Access Domains page). Define this VSA if the firewall has multiple virtual systems.
PaloAlto-Panorama-Admin-Role	3	A default (dynamic) administrative role name or a custom administrative role name on Panorama.
PaloAlto-Panorama-Admin-Access-Domain	4	The name of an access domain for Device Group and Template administrators (configured in the Panorama > Access Domains page).
PaloAlto-User-Group	5	The name of a user group that an authentication profile references.

VSAs forwarded from GlobalProtect endpoints 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	

Name	Number	Value
PaloAlto-Client-Hostname	9	
PaloAlto-GlobalProtect-Client-Version	10	

LDAP

Lightweight Directory Access Protocol (LDAP) is a standard protocol for accessing information directories. You can [Configure LDAP Authentication](#) for end users and for firewall and Panorama administrators.

Configuring the firewall to connect to an LDAP server also enables you to define policy rules based on users and user groups instead of just IP addresses. For the steps, see [Map Users to Groups](#) and [Enable User- and Group-Based Policy](#).

Local Authentication

Although the firewall and Panorama provide local authentication for administrators and end users, [External Authentication Services](#) are preferable in most cases because they provide central account management. However, you might require special user accounts that you don't manage through the directory servers that your organization reserves for regular accounts. For example, you might define a superuser account that is local to the firewall so that you can access the firewall even if the directory server is down. In such cases, you can use the following local authentication methods:

- **(Firewall only) Local database authentication**—To [Configure Local Database Authentication](#), you create a database that runs locally on the firewall and contains user accounts (usernames and passwords or hashed passwords) and user groups. This type of authentication is useful for creating user accounts that reuse the credentials of existing Unix accounts in cases where you know only the hashed passwords, not the plaintext passwords. Because local database authentication is associated with authentication profiles, you can accommodate deployments where different sets of users require different authentication settings, such as [Kerberos](#) single sign-on (SSO) or [Multi-Factor Authentication](#) (MFA). (For details, see [Configure an Authentication Profile and Sequence](#)). For accounts that use plaintext passwords, you can also [define password complexity and expiration settings](#). This authentication method is available to administrators who access the firewall (but not Panorama) and end users who access services and applications through Captive Portal or GlobalProtect.
- **Local authentication without a database**—You can configure [firewall administrative accounts](#) or [Panorama administrative accounts](#) without creating a database of users and user groups that runs locally on the firewall or Panorama. Because this method is not associated with authentication profiles, you cannot combine it with Kerberos SSO or MFA. However, this is the only authentication method that allows password profiles, which enable you to associate individual accounts with password expiration settings that differ from the global settings. (For details, see [Define password complexity and expiration settings](#))

Plan Your Authentication Deployment

The following are key questions to consider before you implement an authentication solution for administrators who access the firewall and end users who access services and applications through Captive Portal.

For both end users and administrators, consider:

- ❑ How can you leverage your existing security infrastructure? Usually, integrating the firewall with an existing infrastructure is faster and cheaper than setting up a new, separate solution just for firewall services. The firewall can integrate with [Multi-Factor Authentication](#), [SAML](#), [Kerberos](#), [TACACS+](#), [RADIUS](#), and [LDAP](#) servers. If your users access services and applications that are external to your network, you can use SAML to integrate the firewall with an identity provider (IdP) that controls access to both external and internal services and applications.
- ❑ How can you optimize the user experience? If you don't want users to authenticate manually and you have a public key infrastructure, you can implement certificate authentication. Another option is to implement [Kerberos](#) or [SAML](#) single sign-on (SSO) so that users can access multiple services and applications after logging in to just one. If your network requires additional security, you can combine certificate authentication with interactive (challenge-response) authentication.
- ❑ Do you require special user accounts that you don't manage through the directory servers that your organization reserves for regular accounts? For example, you might define a superuser account that is local to the firewall so that you can access the firewall even if the directory server is down. You can configure [Local Authentication](#) for these special-purpose accounts.



[External Authentication Services](#) are usually preferable to local authentication because they provide central account management, reliable authentication services, and usually logging and troubleshooting features.

For end users only, consider:

- ❑ Which services and applications are more sensitive than others? For example, you might want stronger authentication for key financial documents than for search engines. To protect your most sensitive services and applications, you can configure [Multi-Factor Authentication](#) (MFA) to ensure that each user authenticates using multiple methods (factors) when accessing those services and applications. To accommodate a variety of security needs, [Configure Authentication Policy](#) rules that trigger MFA or single factor authentication (such as login credentials or certificates) based on specific services, applications, and end users. Other ways to reduce your attack surface include [network segmentation](#) and [user groups for allowed applications](#).

For administrators only, consider:

- ❑ Do you use an external server to centrally manage authorization for all administrative accounts? By defining Vendor-Specific Attributes (VSAs) on the external server, you can quickly change administrative role assignments through your directory service instead of reconfiguring settings on the firewall. VSAs also enable you to specify access domains for administrators of firewalls with multiple virtual systems. [SAML](#), [TACACS+](#), and [RADIUS](#) support external authorization.

Configure Multi-Factor Authentication

To use [Multi-Factor Authentication](#) (MFA) for protecting sensitive services and applications, you must configure Captive Portal to display a web form for the first authentication factor and to record [Authentication Timestamps](#). The firewall uses the timestamps to evaluate the timeouts for [Authentication Policy](#) rules. To enable additional authentication factors, you can integrate the firewall with MFA vendors through RADIUS or vendor APIs. After evaluating Authentication policy, the firewall evaluates Security policy, so you must configure rules for both policy types.



Palo Alto Networks provides support for [MFA vendors](#) through Applications content updates. This means that if you use Panorama to push device group configurations to firewalls, you must [install the same Applications updates](#) on the firewalls as on Panorama to avoid mismatches in vendor support.

MFA vendor API integrations are supported for end-user authentication through Authentication Policy only. For remote user authentication to GlobalProtect portals or gateways or for administrator authentication to the PAN-OS or Panorama web interface, you can only use MFA vendors supported through RADIUS or SAML; MFA services through vendor APIs are not supported in these use cases.

STEP 1 | [Configure Captive Portal](#) in **Redirect** mode to display a web form for the first authentication factor, to record authentication timestamps, and to update user mappings.

STEP 2 | Configure one of the following server profiles to define how the firewall will connect to the service that authenticates users for the first authentication factor.

- [Add a RADIUS server profile](#). This is required if the firewall integrates with an MFA vendor through RADIUS. In this case, the MFA vendor provides the first and all additional authentication factors, so you can skip the next step (configuring an MFA server profile). If the firewall integrates with an MFA vendor through an API, you can still use a RADIUS server profile for the first factor but MFA server profiles are required for the additional factors.
- [Add a SAML IdP server profile](#).
- [Add a Kerberos server profile](#).
- [Add a TACACS+ server profile](#).
- [Add an LDAP server profile](#).



In most cases, an external service is recommended for the first authentication factor. However, you can configure [Configure Local Database Authentication](#) as an alternative.

STEP 3 | Add an MFA server profile.

The profile defines how the firewall connects to the MFA server. Add a separate profile for each authentication factor after the first factor. The firewall integrates with these MFA servers through vendor APIs. You can specify up to three additional factors. Each MFA vendor provides one factor, though some vendors let users choose one factor out of several.

1. Select **Device > Server Profiles > Multi Factor Authentication** and **Add** a profile.
2. Enter a **Name** to identify the MFA server.
3. Select the **Certificate Profile** that the firewall will use to [validate the MFA server certificate](#) when establishing a secure connection to the MFA server.
4. Select the **MFA Vendor** you deployed.
5. Configure the **Value** of each vendor attribute.

The attributes define how the firewall connects to the MFA server. Each vendor **Type** requires different attributes and values; refer to your vendor documentation for details.

6. Click **OK** to save the profile.

STEP 4 | Configure an authentication profile.

The profile defines the order of the authentication factors that users must respond to.

1. Select **Device > Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the authentication profile.
3. Select the **Type** for the first authentication factor and select the corresponding **Server Profile**.
4. Select **Factors**, **Enable Additional Authentication Factors**, and **Add** the MFA server profiles you configured.

The firewall will invoke each MFA service in the listed order, from top to bottom.

5. Click **OK** to save the authentication profile.

STEP 5 | Configure an authentication enforcement object.

The object associates each authentication profile with a Captive Portal method. The method determines whether the first authentication challenge (factor) is transparent or requires a user response.

Select the **Authentication Profile** you configured and enter a **Message** that tells users how to authenticate for the first factor. The message displays in the Captive Portal web form.



If you set the Authentication Method to browser-challenge, the Captive Portal web form displays only if Kerberos SSO authentication fails. Otherwise, authentication for the first factor is automatic; users won't see the web form.

STEP 6 | Configure an Authentication policy rule.

The rule must match the services and applications you want to protect and the users who must authenticate.

1. Select **Policies > Authentication** and **Add** a rule.
2. Enter a **Name** to identify the rule.
3. Select **Source** and **Add** specific zones and IP addresses or select **Any** zones or IP addresses.

The rule applies only to traffic coming from the specified IP addresses or from [interfaces in the specified zones](#).

4. Select **User** and select or **Add** the source users and user groups to which the rule applies (default is **any**).
5. Select **Destination** and **Add** specific zones and IP addresses or select **any** zones or IP addresses.

The IP addresses can be resources (such as servers) for which you want to control access.

6. Select **Service/URL Category** and select or **Add** the [services and service groups](#) for which the rule controls access (default is **service-http**).
7. Select or **Add** the [URL Categories](#) for which the rule controls access (default is **any**). For example, you can create a custom URL category that specifies your most sensitive internal sites.
8. Select **Actions** and select the **Authentication Enforcement** object you created.
9. Specify the **Timeout** period in minutes (default 60) during which the firewall prompts the user to authenticate only once for repeated access to services and applications.



Timeout is a tradeoff between tighter security (less time between authentication prompts) and the user experience (more time between authentication prompts). More frequent authentication is often the right choice for access to critical systems and sensitive areas such as a data center. Less frequent authentication is often the right

choice at the network perimeter and for businesses for which the user experience is key.

10. Click **OK** to save the rule.

STEP 7 | Customize the MFA login page.

The firewall displays this page to tell users how to authenticate for MFA factors and to indicate the authentication status (in progress, succeeded, or failed).

1. Select **Device > Response Pages** and select **MFA Login Page**.
2. Select the **Predefined** response page and **Export** the page to your client system.
3. On your client system, use an HTML editor to customize the downloaded response page and save it with a unique filename.
4. Return to the MFA Login Page dialog on the firewall, **Import** your customized page, **Browse** to select the **Import File**, select the **Destination** (virtual system or **shared** location), click **OK**, and click **Close**.

STEP 8 | Configure a Security policy rule that allows users to access the services and applications that require authentication.

1. [Create a Security Policy Rule](#).
2. **Commit** your changes.



*The [automated correlation engine](#) on the firewall uses several correlation objects to detect events on your network that could indicate credential abuse relating to MFA. To review the events, select **Monitor > Automated Correlation Engine > Correlated Events**.*

STEP 9 | Verify that the firewall enforces MFA.

1. Log in to your network as one of the source users specified in the Authentication rule.
2. Request a service or application that matches one of the services or applications specified in the rule.

The firewall displays the Captive Portal web form for the first authentication factor. The page contains the message you entered in the authentication enforcement object. For example:



3. Enter your user credentials for the first authentication challenge.

The firewall then displays an MFA login page for the next authentication factor. For example, the MFA service might prompt you to select the Voice, SMS, push, or PIN code (OTP) authentication method. If you select push, your phone prompts you to approve the authentication.



4. Authenticate for the next factor.

The firewall displays an authentication success or failure message. If authentication succeeded, the firewall displays an MFA login page for the next authentication factor, if any.

Repeat this step for each MFA factor. After you authenticate for all the factors, the firewall evaluates Security policy to determine whether to allow access to the service or application.

5. End the session for the service or application you just accessed.
6. Start a new session for the same service or application. Be sure to perform this step within the **Timeout** period you configured in the Authentication rule.

The firewall allows access without re-authenticating.

7. Wait until the **Timeout** period expires and request the same service or application.

The firewall prompts you to re-authenticate.

Configure MFA Between RSA SecurID and the Firewall

Multi-factor authentication allows you to protect company assets by using multiple factors to verify a user's identity before allowing them to access network resources. To enable multi-factor authentication (MFA) between the firewall and the RSA SecurID Access Cloud Authentication Service, you must first configure the RSA SecurID Service so that you have the details that you need to configure the firewall to authenticate users using multiple factors. After you have performed the required configuration on the RSA SecurID Access Console, you can configure the firewall to integrate with RSA SecurID.



The Palo Alto Networks next-generation firewall integrates with the RSA SecurID Access Cloud Authentication Service. The MFA API integration with RSA SecurID is supported for cloud-based services only and does not support two-factor authentication for the on-premise Authentication Manager when the second factor uses the Vendor Specific API. The minimum content version required for this integration is 752 and PAN-OS 8.0.2.

- [Get the RSA SecurID Access Cloud Authentication Service Details](#)
- [Configure the Firewall for MFA with RSA SecurID](#)


Get the RSA SecurID Access Cloud Authentication Service Details

In order to securely pass user authentication requests to and from the firewall and the RSA SecurID Access Cloud Authentication Service, you must first go to the RSA SecurID Access Console and configure the RSA Access ID, the authentication service URL, and the client API key that the firewall needs to authenticate to and interact with the service. The firewall also needs the Access Policy ID that uses either the RSA Approve or RSA Tokencode authentication method to authenticate to the identity source.


- **Generate the RSA SecurID API key**—Log on to RSA SecurID Access Console and select **My Account > Company Settings > Authentication API Keys**. Add a new key and then **Save Settings** and **Publish Changes**.



- **Get the RSA SecurID Access endpoint API (Authentication Service Domain) to which the firewall must connect**—Select **Platform > Identity Routers**, pick an Identity Router to **Edit** and jot down the **Authentication Service Domain**. In this example it is `https://rsaready.auth-demo.auth`.



- **Get the Access Policy ID**—Select **Access > Policies** and jot down the name of the access policy that will allow the firewall to act as an authentication client to the RSA SecurID service. The policy must be configured to use either the RSA Approve or the RSA Tokencode authentication methods only.



Configure the Firewall for MFA with RSA SecurID

After you [Get the RSA SecurID Access Cloud Authentication Service Details](#), you can configure the firewall to prompt users for an RSA SecurID token when MFA is invoked.

STEP 1 | Configure the firewall to trust the SSL certificate provided by the RSA SecurID Access endpoint API.

1. Export the SSL certificate from the RSA SecurID Access endpoint and [import it into the firewall](#).

To enable trust between the firewall and the RSA SecurID Access endpoint API, you must either import a self-signed certificate, or the CA certificate used to sign the certificate.

2. [Configure a Certificate Profile](#) (**Device > Certificate Management > Certificate Profile** and click **Add**).

STEP 2 | Configure Captive Portal (**Device > User Identification > Captive Portal Settings**) in Redirect mode to display a web form for authenticating to RSA SecureID. Make sure to specify the Redirect Host as an IP address or a hostname (with no periods in its name) that resolves to the IP address of the Layer 3 interface on the firewall to which web requests are redirected.

STEP 3 | Configure a multi-factor authentication server profile to specify how the firewall must connect with the RSA SecurID cloud service (**Device > Server Profiles > Multi Factor Authentication** and click **Add**).

1. Enter a **Name** to identify the MFA server profile.
2. Select the **Certificate Profile** that you created earlier, **rsa-cert-profile** in this example. The firewall will use this certificate when establishing a secure connection with RSA SecurID cloud service.
3. In the **MFA Vendor** drop-down, select **RSA SecurID Access**.
4. Configure the **Value** for each attribute that you noted in [Get the RSA SecurID Access Cloud Authentication Service Details](#):
 - **API Host**—Enter the hostname or IP address of the RSA SecurID Access API endpoint to which the firewall must connect, **rsaready.auth-demo.auth** in this example.
 - **Base URI**—Do not modify the default value (**/mfa/v1_1**)
 - **Client Key**—Enter the RSA SecurID Client Key.

- **Access ID**—Enter the RSA SecurID Access ID.
- **Assurance Policy**—Enter the RSA SecurID Access Policy name, mfa-policy in this example.
- **Timeout**—The default is 30 seconds.

The screenshot shows the 'Multi Factor Authentication Server Profile' configuration window. The 'Profile Name' is 'rsa-mfa' and the 'Certificate Profile' is 'mfa_cert_prof'. Under 'Server Settings', the 'MFA Vendor' is 'RSA SecurID Access'. A table lists the following settings:

Name	Value
API Host	rsaready.auth-demo.auth
Base URI	/mfa/v1_1
Client Key	*****
Access ID	*****
Assurance Policy	mfa-policy
Timeout (sec)	30 [5 - 600]

At the bottom right are 'OK' and 'Cancel' buttons.

5. Save the profile.

STEP 4 | Configure an authentication profile (Device > Authentication Profile and click Add).

The profile defines the order of the authentication factors that users must respond to.

1. Select the **Type** for the first authentication factor and select the corresponding **Server Profile**.
2. Select **Factors**, **Enable Additional Authentication Factors**, and **Add** the rsa-mfa server profile you created earlier in this example.

The screenshot shows the 'Authentication Profile' configuration window. The 'Profile Name' is 'RSA'. The 'Authentication' tab is selected. The 'Enable Additional Authentication Factors' checkbox is checked. Below, in the 'Factors' section, the 'rsa-mfa' server profile is listed and selected. At the bottom right are 'OK' and 'Cancel' buttons.

3. Click **OK** to save the authentication profile.

STEP 5 | Configure an authentication enforcement object. (Objects > Authentication and click Add).

Make sure to select the authentication profile you just defined called RSA in this example.



STEP 6 | Configure an Authentication policy rule. (Policies > Authentication and click Add)

Your authentication policy rule must match the services and applications you want to protect, specify the users who must authenticate, and include the authentication enforcement object that triggers the authentication profile. In this example, RSA SecurID Access authenticates all users who accessing HTTP, HTTPS, SSH, and VNC traffic with the authentication enforcement object called RSA Auth Enforcement (in **Actions**, select the **Authentication Enforcement** object).




STEP 7 | Commit your changes on the firewall.


STEP 8 | Verify that users on your network are being secured using RSA SecurID using the Push or PIN Code authentication method you enabled.

1. Push authentication

1. Ask a user on your network to launch a web browser and access a website. The Captive Portal page with the IP address or hostname for the Redirect Host you defined earlier should display.
2. Verify that the user enters the credentials for the first authentication factor and then continues to the secondary authentication factor, and selects **Push**.




3. Check for a **Sign-In request** on the RSA SecurID Access application on the user's mobile device.
4. Ask the user to **Accept** the Sign-In Request on the mobile device, and wait for a few seconds for the firewall to receive the notification of successful authentication. The user should be able to access the requested website.



To test an authentication failure, Decline the sign-in request on the mobile device.

2. PIN Code authentication

1. Ask a user on your network to launch a web browser and access a website. The Captive Portal page with the IP address or hostname for the redirect host you defined earlier should display.
2. Verify that the user enters the credentials for the first authentication factor and then continues to the secondary authentication factor, and selects **PIN Code**.



3. Check that a **PIN Code** displays on the RSA SecurID Access application on the user's mobile device.



4. Ask the user to copy the PIN code in the **Enter the PIN...** prompt of the web browser and click **Submit**. Wait for a few seconds for the firewall to receive the notification of successful authentication. The user should be able to access the requested website.

Configure MFA Between Okta and the Firewall

Multi-factor authentication allows you to protect company assets by using multiple factors to verify the identity of users before allowing them to access network resources.

To enable multi-factor authentication (MFA) between the firewall and the Okta identity management service:


- [Configure Okta](#)
- [Configure the firewall to integrate with Okta](#)
- [Verify MFA with Okta](#)

Configure Okta

Log in to the Okta Admin Portal to create your user accounts, define your Okta MFA policy, and obtain the token information required to configure MFA with Okta on the firewall.

STEP 1 | Create your Okta Admin user account.

1. Submit your email address and name, then click **Get Started**.
2. Click the link in the confirmation email and use the included temporary password to log in to the Okta Admin Portal.




3. Create a new password that includes at least 8 characters, one lowercase letter, one uppercase letter, a number, and does not include any part of your username.
4. Select a password reminder question and enter the answer.
5. Select a security image, then **Create My Account**.


STEP 2 | Configure your Okta service.



If you log in and are not redirected to the Okta Admin Portal, select Admin at the upper right.



1. From the Okta Dashboard, log in with your Okta Admin credentials, then select **Applications > Applications**.




2. Select **Add Application**.
3. Search for **Okta Verify**.
4. Select **Add**, then **Done**.




STEP 3 | Create one or more user groups to categorize your users (for example, by device, by policy, or by department) and assign the Okta Verify application.

1. Select **Directory > Groups**.



2. Click **Add Group**.



3. Enter a group **Name** and optionally a **Group Description**, then **Add Group**.



The default group **Everyone** includes all users configured for your organization during the first step in [Configure Okta](#).

4. Select the group you created, then select **Manage Apps**.
5. **Assign** the Okta Verify application you added in Step 2.

6. After the application has been **Assigned**, click **Done**.
7. Repeat this process for all groups that will use the Okta Verify application for MFA.

STEP 4 | Add users and assign them to a group.

1. From the Okta Dashboard, select **Directory > People > Add Person**.

2. Enter the user's **First Name**, **Last Name**, and **Username**. The username must match the **Primary email**, which populates automatically, and the username entered on the firewall. You can optionally enter an alternate email address for the user as the **Secondary Email**.

The screenshot shows the 'Add Person' form with the following fields and values:

- First name: Example
- Last name: User
- Username: exampleuser@paloaltonetworks.com
- Primary email: exampleuser@paloaltonetworks.com
- Secondary email (optional): it_email@paloaltonetworks.com
- Groups (optional): MFA_Obs
- Password: Set by user
- ☒ Send user activation email now

Buttons at the bottom: Save, Save and Add Another, Cancel.

3. Enter the name of the group or **Groups** to associate with this user. When you start typing, the group name populates automatically.
4. Check **Send user activation email now**, then **Save** to add a single user or **Save and Add Another** to continue adding users.

STEP 5 | Assign a test policy to users.


1. Select **Security > Authentication > Sign On**.

There is a **Default Policy** with a **Default Rule** that does not prompt users to log in with MFA.

2. Enter the **Rule Name** and check **Prompt for Factor** to enforce the MFA prompt, and select the type of prompt (**Per Device**, **Every Time**, or **Per Session**), then **Create Rule**.

STEP 6 | Record the Okta authentication token information in a safe place because it is only displayed once.


1. Select **Security > API > Tokens**.
2. Select **Create Token**.




3. Enter a name for the token, then **Create Token**.

4. Copy the **Token Value**.

You can click the **Copy to clipboard** button to copy the Token Value to your clipboard.



5. In the URL for the Okta Admin Dashboard, copy the portion of the URL after **https://** up to **/admin** to use as the **API host**.



6. Omit the domain **okta.com** from this URL to use as the **Organization**.



For example, in the example Okta Admin Dashboard URL above, **https://paloaltonetworks-doc-admin.okta.com/admin/dashboard**:

- The API hostname is **paloaltonetworks-doc-admin.okta.com**.
- The Organization is **paloaltonetworks-doc-admin**.


STEP 7 | Export all certificates in the certificate chain using Base-64 encoding:

1. Depending on your browser, use one of the following methods to export all certificates in the chain.
 - **Chrome**—Press **F12**, then select **Security > View Certificate > Details > Copy to File**.
 - **Firefox**—Select **Options > Privacy & Security > View Certificates > Export**.
 - **Internet Explorer**—Select **Settings > Internet Options > Content > Certificates > Export**.
2. Use the Certificate Export Wizard to export all certificates in the chain and select **Base-64 encoded X.509** as the format.

Configure the firewall to integrate with Okta


As a prerequisite, confirm that you have **mapped** all users that you want to authenticate using Okta.

STEP 1 | Import all certificates in the certificate chain on the firewall and add the imported CA certificates (root and intermediate) to a **Certificate Profile**.



STEP 2 | Add a **Multi Factor Authentication Server Profile** for Okta.


1. Select **Device > Server Profiles > Multi Factor Authentication**.
2. **Add** an MFA server profile.



3. Enter a **Profile Name**.
4. Select the **Certificate Profile** you created in Step 1 in [Configure the firewall to integrate with Okta](#).
5. Select **Okta Adaptive** as the **MFA Vendor**.
6. Enter the **API Host**, **Token**, and **Organization** from Step 4 in [Configure the firewall to integrate with Okta](#).

STEP 3 | [Configure Captive Portal](#) using **Redirect Mode** to redirect users to the MFA vendor's challenge.

STEP 4 | Enable response pages on the [Interface Management Profile](#) to redirect users to the response page challenge.



STEP 5 | Create an [Authentication Profile](#) and add the MFA vendor as a **Factor** (see [Configure Multi-Factor Authentication](#), Step 3.)



STEP 6 | [Enable User-ID](#) on the source zone to require identified users to respond to the challenge using your MFA vendor.

STEP 7 | Create an Authentication Enforcement Object to use the MFA vendor and create an Authentication policy rule (see [Configure Authentication Policy](#), Steps 4 and 5).

STEP 8 | **Commit** your changes.

Verify MFA with Okta

STEP 1 | Verify your users received their enrollment emails, have activated their accounts, and have downloaded the Okta Verify app on their devices.

STEP 2 | Go to a website that will prompt the response page challenge.



If you are using a self-signed certificate instead of a PKI-assigned certificate from your organization, a security warning displays that users must click through to access the challenge.

STEP 3 | Log in to the response page using your Okta credentials.

STEP 4 | Confirm the device receives the challenge push notification.

STEP 5 | Confirm users can successfully access the page after authenticating the challenge by accepting the push notification on their devices.

Configure MFA Between Duo and the Firewall

Multi-factor authentication (MFA) allows you to protect company assets by using multiple factors to verify the identity of users before allowing them to access network resources. There are multiple ways to use the Duo identity management service to authenticate with the firewall:

- Two-factor authentication for VPN logins using the [GlobalProtect Gateway](#) and a [RADIUS](#) server profile (supported on PAN-OS 7.0 and later).
- API-based integration using [Captive Portal](#) and an [MFA server profile](#) (does not require a Duo Authentication Proxy or SAML IdP - supported on PAN-OS 8.0 and later).
- SAML integration for on-premise servers (supported on PAN-OS 8.0 and later).

To enable SAML MFA between the firewall and Duo to secure administrative access to the firewall:

- [Configure Duo for SAML MFA with Duo Access Gateway](#)
- [Configure the Firewall to Integrate with Duo](#)
- [Verify MFA with Duo](#)

Configure Duo for SAML MFA with Duo Access Gateway

Before you begin, verify that you have deployed the [DuoAccessGateway](#) (DAG) on an on-premise server in your DMZ zone.

Create your Duo administrator account and configure the Duo Access Gateway to authenticate your users before they can access resources.

STEP 1 | Create your Duo administrator account.

1. On the Duo account creation page, enter your **First Name, Last Name, Email Address, Cell Phone Number, Company / Account Name**, and select the number of employees in the organization.
2. Agree to the Terms and Privacy Policy and respond to the reCAPTCHA challenge to **Create My Account**.



DUO

Get Your Free Duo Account
Current customers can upgrade now to try more features.

First Name: Last Name:

Email Address: Phone:

Company / Account Name: Select an Option:


☐ I'm an MSP, Reseller, or Partner

☒ By signing up I agree to the Terms and Privacy Policy

☐ I'm not a robot 

STEP 2 | Verify your Duo administrator account.

1. Select the authentication verification method (**Duo Push**, **Text Me**, or **Calling...**).
2. Enter the **Passcode** you receive and **Submit** it to verify your account.



DUO

Setup complete. Click "Text Me" or "Call Me" to complete authentication using your phone as a verification method.

1. Log In

2. Confirm Your Identity

We'll contact you at XXX-XXX-XXXX


☒ Duo Push
 ☐ Text Me
 ☐ Calling...

Passcode:


STEP 3 | Configure your Duo service for SAML.

After creating your configuration, download the configuration file at the top of the page.

1. In the Duo Admin Panel, select **Applications > Protect an Application**.
2. Enter **Palo Alto Networks** to search the applications.
3. Locate **SAML - Palo Alto Networks** in the list of results, then **Protect this Application**.



4. Enter the **Domain**.
5. Select **Admin UI** as the **Palo Alto Networks Service**.
6. Configure your **Policy** and other **Settings**, and **Save Configuration**.



7. Download your configuration file.

The link to download the file is at the top of the page.



STEP 4 | Upload the configuration file to the Duo Access Gateway (DAG).

1. In the DAG admin console, select **Applications**.
2. Click **Choose File** and select the configuration file you downloaded, then **Upload** it.
3. In **Settings > Session Management**, disable **User agent binding**, then **Save Settings**.

STEP 5 | In the DAG admin console, configure your Active Directory or OpenLDAP server as the authentication source and download the metadata file.

1. Log in to the DAG admin console.
2. In **Authentication Source > Set Active Source**, select your **Source type** (Active Directory or OpenLDAP) and **Set Active Source**.


3. In **Configure Sources**, enter the **Attributes**.
 - For Active Directory, enter **mail,sAMAccountName,userPrincipalName,objectGUID**.
 - For OpenLDAP, enter **mail,uid**.
 - For any custom attributes, append them to the end of the list and separate each attribute with a comma. Do not delete any existing attributes.
4. **Save Settings** to save the configuration.
5. Select **Applications > Metadata**, then click **Download XML metadata** to download the XML metadata you will need to import into the firewall.

The file will be named dag.xml. Because this file includes sensitive information to authenticate your Duo account with the firewall, make sure to keep the file in a secure location to avoid the risk of compromising this information.

Configure the Firewall to Integrate with Duo

STEP 1 | Import the Duo metadata.

1. Log on to the firewall web interface.
2. On the firewall, select **Device > Server Profiles > SAML Identity Provider > Import**.
3. Enter the **Profile Name**.
4. **Browse** to the **Identity Provider Metadata** file (dag.xml).
5. If the Duo Access Gateway provides a self-signed certificate as the signing certificate for the IdP, you cannot **Validate Identity Provider Certificate**. In this case, ensure that you are using PAN-OS 9.0.9 or a later 9.0 version to mitigate exposure to [CVE-2020-2021](#).



The screenshot shows the 'SAML Identity Provider Server Profile Import' dialog box. The 'Profile Name' field is set to 'Duo Access Gateway Profile'. The 'Administrator Use Only' checkbox is unchecked. Under the 'Identity Provider Configuration' section, the 'Identity Provider Metadata' field is set to 'C:\fakepath\dag.xml' with a 'Browse...' button to its right. The 'Validate Identity Provider Certificate' and 'Validate Metadata Signature' checkboxes are both unchecked. The 'Maximum Clock Skew (sec)' field is set to '60'. At the bottom right, there are 'OK' and 'Cancel' buttons.

STEP 2 | Add an authentication profile.

The authentication profile allows Duo as the identity provider that validates administrator login credentials.

1. **Add an Authentication Profile**.
2. Enter the profile **Name**.

3. Select **SAML** as the authentication **Type**.
4. Select **Duo Access Gateway Profile** as the **IdP Server Profile**.
5. Select the certificate you want to use for SAML communication with the Duo Access Gateway for the **Certificate for Signing Requests**.
6. Enter **user.username** as the **Username Attribute**.

The screenshot shows the 'Authentication Profile' configuration window with the 'Advanced' tab selected. The 'Name' field is 'Duo Access Gateway'. The 'Type' is 'SAML'. The 'IdP Server Profile' is 'Duo Access Gateway Profile'. The 'Certificate for Signing Requests' is 'cert_admin'. There is a checkbox for 'Enable Single Logout' which is unchecked. The 'Certificate Profile' is 'None'. Below these fields is a section titled 'User Attributes in SAML Messages from IDP' with four input fields: 'Username Attribute' (containing 'User.Username'), 'User Group Attribute', 'Admin Role Attribute', and 'Access Domain Attribute'. At the bottom right are 'OK' and 'Cancel' buttons.

Authentication Profile

Name: Duo Access Gateway

Authentication Factors Advanced

Type: SAML

IdP Server Profile: Duo Access Gateway Profile

Certificate for Signing Requests: cert_admin
Select the certificate to sign SAML messages to IDP

☐ Enable Single Logout

Certificate Profile: None

User Attributes in SAML Messages from IDP

Username Attribute: User.Username


User Group Attribute:

Admin Role Attribute:

Access Domain Attribute:


OK Cancel

7. Select **Advanced** to **Add** an allow list.
8. Select **all**, then click **OK**.
9. **Commit** the changes.



STEP 3 | Specify the authentication settings that the firewall uses for SAML authentication with Duo.

1. Select **Device > Setup > Management** and edit the **Authentication Settings**.
2. Select **Duo Access Gateway** as the **Authentication Profile**, then click **OK**.




3. **Commit** your changes.

STEP 4 | Add accounts for administrators who will authenticate to the firewall using Duo.

1. Select **Device > Administrators** and **Add** an account.
2. Enter a user **Name**.
3. Select **Duo Access Gateway** as the **Authentication Profile**.

4. Select the **Administrator Type**, then click **OK**.

Select **Role Based** if you want to use a custom role for the user. Otherwise, select **Dynamic**. To require administrators to log in using SSO with Duo, assign the authentication profile to all current administrators.

A screenshot of the 'Administrator' configuration dialog box. The dialog has a title bar 'Administrator' with a help icon. It contains several fields: 'Name' with the value 'Admin_User', 'Authentication Profile' with a dropdown menu showing 'Duo Access Gateway', two unchecked checkboxes for 'Use only client certificate authentication (Web)' and 'Use Public Key Authentication (SSH)', 'Administrator Type' with radio buttons for 'Dynamic' (selected) and 'Role Based', and a dropdown menu at the bottom showing 'Superuser'. At the bottom right are 'OK' and 'Cancel' buttons.

Administrator

Name: Admin_User

Authentication Profile: Duo Access Gateway

☐ Use only client certificate authentication (Web)

☐ Use Public Key Authentication (SSH)

Administrator Type: ☒ Dynamic ☐ Role Based

Superuser

OK Cancel

Verify MFA with Duo

- STEP 1** | Log in to the web interface on the firewall.
- STEP 2** | Select **Use Single Sign-On** and **Continue**.
- STEP 3** | Enter your login credentials on the Duo Access Gateway login page.
- STEP 4** | Select an authentication method (push notification, phone call, or passcode entry).
When you authenticate successfully, you will be redirected to the firewall web interface.

Configure SAML Authentication

To configure [SAML](#) single sign-on (SSO) and single logout (SLO), you must register the firewall and the IdP with each other to enable communication between them. If the IdP provides a metadata file containing registration information, you can import it onto the firewall to register the IdP and to create an IdP server profile. The server profile defines how to connect to the IdP and specifies the certificate that the IdP uses to sign SAML messages. You can also use a certificate for the firewall to sign SAML messages. Using certificates is a requirement to secure communications between the firewall and the IdP.

Palo Alto Networks requires HTTPS to ensure the confidentiality of all SAML transactions instead of alternative approaches such as encrypted SAML assertions. To ensure the integrity of all messages processed in a SAML transaction, Palo Alto Networks requires digital certificates to cryptographically sign all messages.

The following procedure describes how to configure SAML authentication for end users and firewall administrators. You can also [configure SAML authentication for Panorama administrators](#).



SSO is available to administrators and to GlobalProtect and Captive Portal end users. SLO is available to administrators and GlobalProtect end users, but not to Captive Portal end users.

Administrators can use SAML to authenticate to the firewall web interface, but not to the CLI.

STEP 1 | Obtain the certificates that the IdP and firewall will use to sign SAML messages.

If the certificates don't specify key usage attributes, all usages are allowed by default, including signing messages. In this case, you can [Obtain Certificates](#) by any method.

If the certificates do specify key usage attributes, one of the attributes must be Digital Signature, which is not available on certificates that you generate on the firewall or Panorama. In this case, you must [import the certificates](#):

- **Certificate the firewall uses to sign SAML messages**—Import the certificate from your enterprise certificate authority (CA) or a third-party CA.
- **Certificate the IdP uses to sign SAML messages** ([Required for all deployments](#))—Import a metadata file containing the certificate from the IdP (see the next step). The IdP certificate is limited to the following algorithms:

Public key algorithms—RSA (1,024 bits or larger) and ECDSA (all sizes). A firewall in FIPS/CC mode supports RSA (2,048 bits or larger) and ECDSA (all sizes).

Signature algorithms—SHA1, SHA256, SHA384, and SHA512. A firewall in FIPS/CC mode supports SHA256, SHA384, and SHA512.

STEP 2 | Add a SAML IdP server profile.

The server profile registers the IdP with the firewall and defines how they connect.

In this example, you import a SAML metadata file from the IdP so that the firewall can automatically create a server profile and populate the connection, registration, and IdP certificate information.



If the IdP doesn't provide a metadata file, select Device > Server Profiles > SAML Identity Provider, Add the server profile, and manually enter the information (consult your IdP administrator for the values).

1. Export the SAML metadata file from the IdP to a client system from which you can upload the metadata to the firewall.

The certificate specified in the file must meet the requirements listed in the preceding step. Refer to your IdP documentation for instructions on exporting the file.

2. Select **Device > Server Profiles > SAML Identity Provider** and **Import** the metadata file onto the firewall.
3. Enter a **Profile Name** to identify the server profile.
4. **Browse** to the **Identity Provider Metadata** file.
5. Select **Validate Identity Provider Certificate** (default) to validate the chain of trust and optionally the revocation status of the IdP certificate.

To enable this option, a Certificate Authority (CA) must issue your IdP's signing certificate. You must create a Certificate Profile that has the CA that issued the IdP's signing certificate. In the Authentication Profile, select the SAML Server profile and Certificate Profile to validate the IdP certificate.

If your IdP signing certificate is a self-signed certificate, there is no chain of trust; as a result, you cannot enable this option. The firewall always validates the signature of the SAML Responses or Assertions against the Identity Provider certificate that you configure whether or not you enable the **Validate Identity Provider Certificate** option. If your IdP provides a self-signed certificate, ensure that you are using PAN-OS 9.0.9 or a later 9.0 version to mitigate exposure to [CVE-2020-2021](#).



Validate the certificate to ensure it hasn't been compromised and to improve security.

6. Enter the **Maximum Clock Skew**, which is the allowed difference in seconds between the system times of the IdP and the firewall at the moment when the firewall validates IdP messages (default is 60; range is 1 to 900). If the difference exceeds this value, authentication fails.
7. Click **OK** to save the server profile.
8. Click the server profile Name to display the profile settings. Verify that the imported information is correct and edit it if necessary.
9. Whether you import the IdP metadata or manually enter the IdP information, always ensure that the signing certificate of your SAML identity provider is the **Identity Provider Certificate** for your server profile and your IdP sends signed SAML Responses, Assertions, or both.

STEP 3 | Configure an authentication profile.

The profile defines authentication settings that are common to a set of users.

1. Select **Device > Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the profile.
3. Set the **Type** to **SAML**.
4. Select the **IdP Server Profile** you configured.
5. Select the **Certificate for Signing Requests**.

The firewall uses this certificate to sign messages it sends to the IdP. You can import a certificate generated by your enterprise CA or you can generate a certificate using the root CA that was generated on the firewall or Panorama.

6. (Optional) **Enable Single Logout** (disabled by default).
7. Select the **Certificate Profile** that the firewall will use to validate the **Identity Provider Certificate**.
8. Enter the **Username Attribute** that IdP messages use to identify users (default **username**).



*When you predefine dynamic administrator roles for users, use lower-case to specify the role (for example, enter **superreader**, not **SuperReader**). If you manage administrator authorization in the IdP identity store, specify the Admin Role Attribute and Access Domain Attribute also.*

9. Select **Advanced** and **Add** the users and user groups that are allowed to authenticate with this authentication profile.
10. Click **OK** to save the authentication profile.

STEP 4 | Assign the authentication profile to firewall applications that require authentication.

1. Assign the authentication profile to:
 - Administrator accounts that you manage locally on the firewall. In this example, [Configure a Firewall Administrator Account](#) before you verify the SAML configuration later in this procedure.
 - Administrator accounts that you manage externally in the IdP identity store. Select **Device > Setup > Management**, edit the Authentication Settings, and select the **Authentication Profile** you configured.
 - Authentication policy rules that secure the services and applications that end users access through Captive Portal. See [Configure Authentication Policy](#).
 - [GlobalProtect](#) portals and gateways that end users access.
2. **Commit** your changes.

The firewall validates the **Identity Provider Certificate** that you assigned to the SAML IdP server profile.

STEP 5 | Create a SAML metadata file to register the firewall application (management access, Captive Portal, or GlobalProtect) on the IdP.

1. Select **Device > Authentication Profile** and, in the Authentication column for the authentication profile you configured, click **Metadata**.
2. In the **Commands** drop-down, select the application you want to register:
 - **management** (default)—Administrative access to the web interface.
 - **captive-portal**—End user access to services and applications through Captive Portal.
 - **global-protect**—End user access to services and applications through GlobalProtect.
3. ([Captive Portal or GlobalProtect only](#)) for the **Vsysname Combo**, select the virtual system in which the Captive Portal settings or GlobalProtect portal are defined.
4. Enter the interface, IP address, or hostname based on the application you will register:
 - **management**—For the **Management Choice**, select **Interface** (default) and select an interface that is enabled for management access to the web interface. The default selection is the IP address of the MGT interface.
 - **captive-portal**—For the **IP Hostname**, enter the IP address or hostname of the **Redirect Host** (see **Device > User Identification > Captive Portal Settings**).
 - **global-protect**—For the **IP Hostname**, enter the hostname or IP address of the GlobalProtect portal or gateway.
5. Click **OK** and save the metadata file to your client system.
6. Import the metadata file into the IdP server to register the firewall application. Refer to your IdP documentation for instructions.

STEP 6 | Verify that users can authenticate using SAML SSO.

For example, to verify that SAML is working for access to the web interface using a local administrator account:

1. Go to the URL of the firewall web interface.
2. Click **Use Single Sign-On**.
3. Enter the username of the administrator.
4. Click **Continue**.

The firewall redirects you to authenticate to the IdP, which displays a login page. For example:



5. Log in using your SSO username and password.

After you successfully authenticate on the IdP, it redirects you back to the firewall, which displays the web interface.

6. Use your firewall administrator account to request access to another SSO application.

Successful access indicates SAML SSO authentication succeeded.

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. With Kerberos SSO enabled, the user needs to log in only for initial access to your network (such as logging in to Microsoft Windows). After this initial login, the user can access any browser-based service in the network (such as the firewall web interface) without having to log in again until the SSO session expires.

STEP 1 | Create a Kerberos keytab.

The keytab is a file that contains the principal name and password of the firewall, and is required for the SSO process. When you configure Kerberos in your [Authentication Profile and Sequence](#), the firewall first checks for a Kerberos SSO hostname. If you provide a hostname, the firewall searches the keytabs for a service principal name that matches the hostname and uses only that keytab for decryption. If you do not provide a hostname, the firewall tries each keytab in the authentication sequence until it is able to successfully authenticate using Kerberos.



If the Kerberos SSO hostname is included in the request sent to the firewall, then the hostname must match the service principal name of the keytab; otherwise, the Kerberos authentication request is not sent.

1. Log in to the Active Directory server and open a command prompt.
2. Enter the following command to register the service principal name (SPN) for GlobalProtect or Captive Portal, where `<portal_fqdn>` and `<service_account_username>` are variables.
setspn -s HTTP/<portal_fqdn> <service_account_username>
3. Create Kerberos account for the firewall. Refer to your Kerberos documentation for the steps.
4. Log in to the KDC and open a command prompt.
5. Enter the following command, where `<portal_fqdn>`, `<kerberos_realm>`, `<netbios_name>`, `<service_account_username>`, `<password>`, `<filename>`, and `<algorithm>` are variables.

```
ktpass /princ HTTP <portal_fqdn>@<kerberos_realm> /mapuser  
<netbios_name>\<service_account_username> /pass <password>/out  
<filename>.keytab /ptype KRB5_NT_PRINCIPAL /crypto <algorithm>
```



*The `<kerberos_realm>` value must be in all uppercase characters (for example, enter **AD1.EXAMPLE.COM**, not **ad1.example.com**).*



*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 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 | [Configure an Authentication Profile and Sequence](#) to define Kerberos settings and other authentication options that are common to a set of users.

- Enter the **Kerberos Realm** (usually the DNS domain of the users, except that the realm is uppercase).
- **Import** the **Kerberos Keytab** that you created for the firewall.

STEP 3 | Assign the authentication profile to the firewall application that requires authentication.

-
- Administrative access to the web interface—[Configure a Firewall Administrator Account](#) and assign the authentication profile you configured.
 - End user access to services and applications—Assign the authentication profile you configured to an authentication enforcement object. When configuring the object, set the **Authentication Method** to **browser-challenge**. Assign the object to Authentication policy rules. For the full procedure to configure authentication for end users, see [Configure Authentication Policy](#).

Configure Kerberos Server Authentication

You can use [Kerberos](#) to natively authenticate end users and firewall or Panorama administrators 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.



To use a Kerberos server for authentication, the server must be accessible over an IPv4 address. IPv6 addresses are not supported.

STEP 1 | Add a Kerberos server profile.

The profile defines how the firewall connects to the Kerberos server.

1. Select **Device** > **Server Profiles** > **Kerberos** and **Add** a server profile.
2. Enter a **Profile Name** to identify the server profile.
3. **Add** each server and specify a **Name** (to identify the server), IPv4 address or FQDN of the **Kerberos Server**, and optional **Port** number for communication with the server (default 88).



If you use an FQDN address object to identify the server and you subsequently change the address, you must commit the change in order for the new server address to take effect.

4. Click **OK** to save your changes to the profile.

STEP 2 | Assign the server profile to an [Configure an Authentication Profile and Sequence](#).

The authentication profile defines authentication settings that are common to a set of users.

STEP 3 | Assign the authentication profile to the firewall application that requires authentication.

- Administrative access to the web interface—[Configure a Firewall Administrator Account](#) and assign the authentication profile you configured.
- End user access to services and applications—Assign the authentication profile you configured to an authentication enforcement object and assign the object to Authentication policy rules. For the full procedure to configure authentication for end users, see [Configure Authentication Policy](#).

STEP 4 | Verify that the firewall can [Test Authentication Server Connectivity](#) to authenticate users.

Configure TACACS+ Authentication

You can configure [TACACS+](#) authentication for end users and firewall or Panorama administrators. You can also use a TACACS+ server to manage administrator authorization (role and access domain assignments) by defining [Vendor-Specific Attributes \(VSAs\)](#). For all users, you must [configure a TACACS+ server profile](#) that defines how the firewall or Panorama connects to the server. You then [assign the server profile to an authentication profile](#) for each set of users who require common authentication settings. What you do with the authentication profile depends on which users the TACACS+ server authenticates:

- **End users**—Assign the authentication profile to an authentication enforcement object and assign the object to Authentication policy rules. For the full procedure, see [Configure Authentication Policy](#).
- **Administrative accounts with authorization managed locally on the firewall or Panorama**—Assign the authentication profile to [firewall administrator](#) or [Panorama administrator](#) accounts.
- **Administrative accounts with authorization managed on the TACACS+ server**—The following procedure describes how to configure TACACS+ authentication and authorization for firewall administrators. For Panorama administrators, refer to [Configure TACACS+ Authentication for Panorama Administrators](#).

STEP 1 | Add a TACACS+ server profile.

The profile defines how the firewall connects to the TACACS+ server.

1. Select **Device > Server Profiles > TACACS+** and **Add** a profile.
2. Enter a **Profile Name** to identify the server profile.
3. (Optional) Select **Administrator Use Only** to restrict access to administrators.
4. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–20).
5. Select the **Authentication Protocol** (default is **CHAP**) that the firewall uses to authenticate to the TACACS+ server.



Select CHAP if the TACACS+ server supports that protocol; it is more secure than PAP.

6. **Add** each TACACS+ server and enter the following:
 - **Name** to identify the server
 - **TACACS+ Server** IP address or FQDN. If you use an FQDN address object to identify the server and you subsequently change the address, you must commit the change for the new server address to take effect.
 - **Secret/Confirm Secret** (a key to encrypt usernames and passwords)
 - **Server Port** for authentication requests (default is 49)
7. Click **OK** to save the server profile.

STEP 2 | Assign the TACACS+ server profile to an authentication profile.

The authentication profile defines authentication settings that are common to a set of users.

1. Select **Device > Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the profile.
3. Set the **Type** to **TACACS+**.
4. Select the **Server Profile** you configured.
5. Select **Retrieve user group from TACACS+** to collect user group information from VSAs defined on the TACACS+ server.

The firewall matches the group information against the groups you specify in the Allow List of the authentication profile.

6. Select **Advanced** and, in the Allow List, **Add** the users and groups that are allowed to authenticate with this authentication profile.
7. Click **OK** to save the authentication profile.

STEP 3 | Configure the firewall to use the authentication profile for all administrators.

1. Select **Device > Setup > Management** and edit the Authentication Settings.
2. Select the **Authentication Profile** you configured and click **OK**.

STEP 4 | Configure the roles and access domains that define authorization settings for administrators.

If you already defined **TACACS+** VSAs on the TACACS+ server, the names you specify for roles and access domains on the firewall must match the VSA values.

1. **Configure an Admin Role Profile** if the administrator will use a custom role instead of a predefined (dynamic) role.
2. Configure an access domain if the firewall has more than one virtual system—Select **Device > Access Domain**, **Add** an access domain, enter a **Name** to identify the access domain, and **Add** each virtual system that the administrator will access, and then click **OK**.

STEP 5 | **Commit** your changes to activate them on the firewall.

STEP 6 | Configure the TACACS+ server to authenticate and authorize administrators.

Refer to your TACACS+ server documentation for the specific instructions to perform these steps:

1. Add the firewall IP address or hostname as the TACACS+ client.
2. Add the administrator accounts.



If you selected CHAP as the Authentication Protocol, you must define accounts with reversibly encrypted passwords. Otherwise, CHAP authentication will fail.

3. Define **TACACS+** VSAs for the role, access domain, and user group of each administrator.



*When you predefine dynamic administrator roles for users, use lower-case to specify the role (for example, enter **superuser**, not **SuperUser**).*

STEP 7 | Verify that the TACACS+ server performs authentication and authorization for administrators.

1. Log in the firewall web interface using an administrator account that you added to the TACACS+ server.
2. Verify that you can access only the web interface pages that are allowed for the role you associated with the administrator.
3. In the **Monitor**, **Policies**, and **Objects** tabs, verify that you can access only the virtual systems that are allowed for the access domain you associated with the administrator.

Configure RADIUS Authentication

You can configure [RADIUS](#) authentication for end users and firewall or Panorama administrators. For administrators, you can use RADIUS to manage authorization (role and access domain assignments) by defining [Vendor-Specific Attributes \(VSAs\)](#). You can also use RADIUS to implement [Multi-Factor Authentication \(MFA\)](#) for administrators and end users. To enable RADIUS authentication, you must configure a RADIUS server profile that defines how the firewall or Panorama connects to the server (see Step 1 below). You then assign the server profile to an authentication profile for each set of users who require common authentication settings (see Step 5 below). What you do with the authentication profile depends on which users the RADIUS server authenticates:

- **End users**—Assign the authentication profile to an authentication enforcement object and assign the object to Authentication policy rules. For the full procedure, see [Configure Authentication Policy](#).



You can also configure client systems to send RADIUS Vendor-Specific Attributes (VSAs) to the RADIUS server by assigning the authentication profile to a GlobalProtect portal or gateway. RADIUS administrators can then perform administrative tasks based on those VSAs.

- **Administrative accounts with authorization managed locally on the firewall or Panorama**—Assign the authentication profile to [firewall administrator](#) or [Panorama administrator](#) accounts.
- **Administrative accounts with authorization managed on the RADIUS server**—The following procedure describes how to configure RADIUS authentication and authorization for firewall administrators. For Panorama administrators, refer to [Configure RADIUS Authentication for Panorama Administrators](#).

STEP 1 | Add a RADIUS server profile.

The profile defines how the firewall connects to the RADIUS server.

1. Select **Device > Server Profiles > RADIUS** and **Add** a profile.
2. Enter a **Profile Name** to identify the server profile.
3. (Optional) Select **Administrator Use Only** to restrict access to administrators.
4. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–120).



If you use the server profile to integrate the firewall with an MFA service, enter an interval that gives users enough time to authenticate. For example, if the MFA service prompts for a one-time password (OTP), users need time to see the OTP on their endpoint device and then enter the OTP in the MFA login page.

5. Enter the number of **Retries**.
6. Select the **Authentication Protocol** (default is **PEAP-MSCHAPv2**) that the firewall uses to authenticate to the RADIUS server.

Depending on which factors you want to use to authenticate users within your multi-factor authentication (MFA) environment, select the appropriate authentication protocol:

- **Username, password, and push (an automatically triggered out-of-band request):** Supported with all authentication protocols
- **Push, password, token, and PIN (when password or token or PIN are provided together):** Supported with PAP, PEAP with GTC, and EAP-TTLS with PAP
- **Username, password, token, and PIN, and challenge-response (when password or token or PIN are provided together):** Supported with PAP and PEAP with GTC

If you select an EAP authentication method (PEAP-MSCHAPv2, PEAP with GTC, or EAP-TTLS with PAP), confirm that your RADIUS server supports Transport Layer Security (TLS) 1.1 or higher and that the root and intermediate certificate authorities (CAs) for your RADIUS server are included in the certificate [profile](#) associated with the RADIUS server profile. If you select an EAP method and you do not associate a correctly configured certificate profile with the RADIUS profile, authentication fails.

7. **Add** each RADIUS server and enter the following:

- **Name** to identify the server
- **RADIUS Server** IP address or FQDN. If you use an FQDN to identify the server and you subsequently change the address, you must commit the change for the new server address to take effect.
- **Secret/Confirm Secret** is a key to encrypt passwords and can be up to 64 characters in length.
- **Server Port** for authentication requests (default is 1812)

8. Click **OK** to save the server profile.

For redundancy, add multiple RADIUS servers in the sequence you want the firewall to use. If you have selected an EAP method, configure an authentication [sequence](#) to ensure that users will be able to successfully respond to the authentication challenge. There is no alternate authentication method with EAP: if the user fails the authentication challenge and you have not configured an authentication sequence that allows another authentication method, authentication fails.

STEP 2 | If you are using PEAP-MSCHAPv2 with GlobalProtect, select **Allow users to change passwords after expiry** to allow GlobalProtect users to changed expired passwords to log in.

STEP 3 | ([PEAP-MSCHAPv2](#), [PEAP with GTC](#), or [EAP-TTLS with PAP only](#)) To anonymize the user's identity in the outer tunnel that is created after authenticating with the server, select **Make Outer Identity Anonymous**.



You must configure the RADIUS server so that the entire chain allows access for anonymous users. Some RADIUS server configurations may not support anonymous outer IDs, and you may need to clear the option. When cleared, the RADIUS server transmits usernames in cleartext.

STEP 4 | If you select an EAP authentication method, select a [Certificate Profile](#).

STEP 5 | Assign the RADIUS server profile to an authentication profile.

The authentication profile defines authentication settings that are common to a set of users.

1. Select **Device > Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the authentication profile.
3. Set the **Type** to **RADIUS**.
4. Select the **Server Profile** you configured.
5. Select **Retrieve user group from RADIUS** to collect user group information from VSAs defined on the RADIUS server.

The firewall matches the group information against the groups you specify in the Allow List of the authentication profile.

6. Select **Advanced** and, in the Allow List, **Add** the users and groups that are allowed to authenticate with this authentication profile.
7. Click **OK** to save the authentication profile.

STEP 6 | Configure the firewall to use the authentication profile for all administrators.

1. Select **Device > Setup > Management** and edit the Authentication Settings.
2. Select the **Authentication Profile** you configured and click **OK**.

STEP 7 | Configure the roles and access domains that define authorization settings for administrators.

If you already defined **RADIUS** VSAs on the RADIUS server, the names you specify for roles and access domains on the firewall must match the VSA values.

1. [Configure an Admin Role Profile](#) if the administrator uses a custom role instead of a predefined (dynamic) role.
2. Configure an access domain if the firewall has more than one virtual system:
 1. Select **Device > Access Domain, Add** an access domain, and enter a **Name** to identify the access domain.
 2. **Add** each virtual system that the administrator will access, and then click **OK**.

STEP 8 | Commit your changes to activate them on the firewall.

STEP 9 | Configure the RADIUS server to authenticate and authorize administrators.

Refer to your RADIUS server documentation for the specific instructions to perform these steps:

1. Add the firewall IP address or hostname as the RADIUS client.
2. Add the administrator accounts.



If the RADIUS server profile specifies CHAP as the Authentication Protocol, you must define accounts with reversibly encrypted passwords. Otherwise, CHAP authentication will fail.

3. Define the vendor code for the firewall (25461) and define the **RADIUS** VSAs for the role, access domain, and user group of each administrator.

When you predefine dynamic administrator roles for users, use lower-case to specify the role (for example, enter **superuser**, not **SuperUser**).



When configuring the advanced vendor options on the ACS, you must set both the Vendor Length Field Size and Vendor Type Field Size to 1. Otherwise, authentication will fail.

4. If you have selected an EAP method, the firewall validates the server but not the client. To ensure client validity, restrict clients by IP address or subdomain.

STEP 10 | Verify that the RADIUS server performs authentication and authorization for administrators.

1. Log in the firewall web interface using an administrator account that you added to the RADIUS server.
2. Verify that you can access only the web interface pages that are allowed for the role you associated with the administrator.
3. In the **Monitor**, **Policies**, and **Objects** tabs, verify that you can access only the virtual systems that are allowed for the access domain you associated with the administrator.
4. In **Monitor > Authentication**, verify the **Authentication Protocol**.
5. Test the connection and the validity of the certificate [profile](#) using the following CLI command:

```
admin@PA-220 > test authentication authentication-profile auth-profile  
username <username> password <password>
```

Configure LDAP Authentication

You can use [LDAP](#) to authenticate end users who access applications or services through Captive Portal and authenticate firewall or Panorama administrators who access the web interface.



You can also connect to an LDAP server to define policy rules based on user groups. For details, see [Map Users to Groups](#).

STEP 1 | Add an LDAP server profile.

The profile defines how the firewall connects to the LDAP server.

1. Select **Device > Server Profiles > LDAP** and **Add** a server profile.
2. Enter a **Profile Name** to identify the server profile.
3. (Multi-vsys only) Select the **Location** in which the profile is available.
4. (Optional) Select **Administrator Use Only** to restrict access to administrators.
5. **Add** the LDAP servers (up to four). For each server, enter a **Name** (to identify the server), **LDAP Server IP** address or FQDN, and server **Port** (default 389).



If you use an FQDN address object to identify the server and you subsequently change the address, you must commit the change for the new server address to take effect.

6. Select the server **Type**.
7. Select the **Base DN**.

To identify the Base DN of your directory, open the **Active Directory Domains and Trusts** Microsoft Management Console snap-in and use the name of the top-level domain.

8. Enter the **Bind DN** and **Password** to enable the authentication service to authenticate the firewall.



The Bind DN account must have permission to read the LDAP directory.

9. Enter the **Bind Timeout** and **Search Timeout** in seconds (default is 30 for both).
10. Enter the **Retry Interval** in seconds (default is 60).
11. (Optional) If you want the endpoint to use SSL or TLS for a more secure connection with the directory server, enable the option to **Require SSL/TLS secured connection** (enabled by default). The protocol that the endpoint uses depends on the server port:
 - 389 (default)—TLS (Specifically, the device uses the [StartTLS operation](#), which upgrades the initial plaintext connection to TLS.)
 - 636—SSL
 - Any other port—The device first attempts to use TLS. If the directory server doesn't support TLS, the device falls back to SSL.
12. (Optional) For additional security, enable the option to **Verify Server Certificate for SSL sessions** so that the endpoint verifies the certificate that the directory server presents for SSL/TLS connections. To enable verification, you must also enable the option to **Require SSL/TLS secured connection**. For verification to succeed, the certificate must meet one of the following conditions:
 - It is in the list of device certificates: **Device > Certificate Management > Certificates > Device Certificates**. If necessary, import the certificate into the device.
 - The certificate signer is in the list of trusted certificate authorities: **Device > Certificate Management > Certificates > Default Trusted Certificate Authorities**.
13. Click **OK** to save the server profile.

STEP 2 | Assign the server profile to [Configure an Authentication Profile and Sequence](#) to define various authentication settings.

STEP 3 | Assign the authentication profile to the firewall application that requires authentication.

- **Administrative access to the web interface**—[Configure a Firewall Administrator Account](#) and assign the authentication profile you configured.
- **End user access to services and applications**—For the full procedure to configure authentication for end users, see [Configure Authentication Policy](#).

STEP 4 | Verify that the firewall can [Test Authentication Server Connectivity](#) to authenticate users.

Connection Timeouts for Authentication Servers

You can configure the firewall to use [External Authentication Services](#) for authenticating administrators who access the firewall or Panorama and end users who access services or applications through Captive Portal. To ensure that the firewall does not waste resources by continuously trying to reach an authentication server that is unreachable, you can set a timeout interval after which the firewall stops trying to connect. You set the timeout in the server profiles that define how the firewall connects to the authentication servers. When choosing timeout values, your goal is to strike a balance between the need to conserve firewall resources and to account for normal network delays that affect how quickly authentication servers respond to the firewall.

- [Guidelines for Setting Authentication Server Timeouts](#)
- [Modify the PAN-OS Web Server Timeout](#)
- [Modify the Captive Portal Session Timeout](#)

Guidelines for Setting Authentication Server Timeouts

The following are some guidelines for setting the timeouts for firewall attempts to connect with [External Authentication Services](#).

- ❑ In addition to the timeouts you set in server profiles for specific servers, the firewall has a global PAN-OS web server timeout. This global timeout applies when the firewall connects to any external server for authenticating administrative access to the firewall web interface or PAN-OS XML API and end user access to applications or services through Captive Portal. The global timeout is 30 seconds by default (range is 3 to 125). It must be the same as or greater than the total time that any server profile allows for connection attempts. The total time in a server profile is the timeout value multiplied by the number of retries and the number of servers. For example, if a RADIUS server profile specifies a 3-second timeout, 3 retries, and 4 servers, the total time that the profile allows for connection attempts is 36 seconds (3 x 3 x 4). [Modify the PAN-OS Web Server Timeout](#) if necessary.



Do not change the PAN-OS web server timeout unless you see authentication failures. Setting the timeout too high could degrade the performance of the firewall or cause it to drop authentication requests. You can review authentication failures in Authentication logs.

- ❑ The firewall applies a Captive Portal session timeout that defines how long end users can take to respond to the authentication challenge in a Captive Portal web form. The web form displays when users request services or applications that match an Authentication policy rule. The session timeout is 30 seconds by default (range is 1 to 1,599,999). It must be the same as or greater than the PAN-OS web server timeout. [Modify the Captive Portal Session Timeout](#) if necessary. Keep in mind that increasing the PAN-OS web server and Captive Portal session timeouts might degrade the performance of the firewall or cause it to drop authentication requests.



The Captive Portal session timeout is not related to the timers that determine how long the firewall retains IP address-to-username mappings.

- ❑ Timeouts are cumulative for authentication sequences. For example, consider the case of an authentication sequence with two authentication profiles. One authentication profile specifies a RADIUS server profile with a 3-second timeout, 3 retries, and 4 servers. The other authentication profile specifies a TACACS+ server profile with a 3-second timeout and 2 servers. The longest possible period in which the firewall can try to authenticate user accounts with that authentication sequence is 42 seconds: 36 seconds for the RADIUS server profile plus 6 seconds for the TACACS+ server profile.

- ❑ The non-configurable timeout for Kerberos servers is 17 seconds for each server specified in the Kerberos server profile.
- ❑ To configure the timeouts and related settings for other server types, see:
 - [Add an MFA server profile.](#)
 - [Add a SAML IdP server profile.](#)
 - [Add a TACACS+ server profile.](#)
 - [Add a RADIUS server profile.](#)
 - [Add an LDAP server profile.](#)

Modify the PAN-OS Web Server Timeout

The PAN-OS web server timeout must be the same as or greater than the timeout in any authentication server profile multiplied by the number of retries and the number of servers in that profile.



Do not change the PAN-OS web server timeout unless you see authentication failures. Setting the timeout too high could degrade the performance of the firewall or cause it to drop authentication requests. You can review authentication failures in [Authentication logs](#).

STEP 1 | Access the firewall [CLI](#).

STEP 2 | Set the PAN-OS web server timeout by entering the following commands, where **<value>** is the number of seconds (default is 30; range is 3 to 125).

```
> configure
# set deviceconfig setting 13-service timeout <value>
# commit
```

Modify the Captive Portal Session Timeout

The Captive Portal session timeout must be the same as or greater than the PAN-OS web server timeout. For details, see [Connection Timeouts for Authentication Servers](#).



The more you raise the PAN-OS web server and Captive Portal session timeouts, the slower Captive Portal will respond to users.

STEP 1 | Select **Device > Setup > Session** and edit the Session Timeouts.

STEP 2 | Enter a new **Captive Portal** value in seconds (default is 30; range is 1 to 1,599,999) and click **OK**.

STEP 3 | **Commit** your changes.

Configure Local Database Authentication

You can configure a user database that is local to the firewall to authenticate administrators who access the firewall web interface and to authenticate end users who access applications through Captive Portal or GlobalProtect. Perform the following steps to configure [Local Authentication](#) with a local database.



[External Authentication Services](#) are usually preferable to local authentication because they provide the benefit of central account management.

You can also configure local authentication without a database, but only for [firewall](#) or [Panorama](#) administrators.

STEP 1 | Add the user account to the local database.

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 | Add the user group to the local database.

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.

The authentication profile defines authentication settings that are common to a set of users. Set the authentication **Type** to **Local Database**.

STEP 4 | Assign the authentication profile to an administrator account or to an Authentication policy rule for end users.

- **Administrators**—[Configure a Firewall Administrator Account](#):
Specify the **Name** of a user you defined earlier in this procedure.
Assign the **Authentication Profile** that you configured for the account.
- **End users**—For the full procedure to configure authentication for end users, see [Configure Authentication Policy](#).

STEP 5 | Verify that the firewall can [Test Authentication Server Connectivity](#) to authenticate users.

Configure an Authentication Profile and Sequence

An authentication profile defines the authentication service that validates the login credentials of administrators who access the firewall web interface and end users who access applications through Captive Portal or GlobalProtect. The service can be [Local Authentication](#) that the firewall provides or [External Authentication Services](#). The authentication profile also defines options such as [Kerberos](#) single sign-on (SSO).

Some networks have multiple databases (such as TACACS+ and LDAP) for different users and user groups. To authenticate users in such cases, configure an *authentication sequence*—a ranked order of authentication profiles that the firewall matches a user against during login. The firewall checks against each profile in sequence until one successfully authenticates the user. A user is denied access only if authentication fails for all the profiles in the sequence. The sequence can specify authentication profiles that are based on any authentication service that the firewall supports excepts [Multi-Factor Authentication](#) (MFA) and [SAML](#).

STEP 1 | ([External service only](#)) Enable the firewall to connect to an external server for authenticating users:

1. Set up the external server. Refer to your server documentation for instructions.
2. Configure a server profile for the type of authentication service you use.
 - [Add a RADIUS server profile.](#)



If the firewall integrates with an MFA service through RADIUS, you must add a RADIUS server profile. In this case, the MFA service provides all the authentication factors. If the firewall integrates with an MFA service through a vendor API, you can still use a RADIUS server profile for the first factor but MFA server profiles are required for additional factors.

- [Add an MFA server profile.](#)
- [Add a SAML IdP server profile.](#)
- [Add a Kerberos server profile.](#)
- [Add a TACACS+ server profile.](#)
- [Add an LDAP server profile.](#)

STEP 2 | ([Local database authentication only](#)) Configure a user database that is local to the firewall.

Perform these steps for each user and user group for which you want to configure [Local Authentication](#) based on a user identity store that is local to the firewall:

1. [Add the user account to the local database.](#)
2. ([Optional](#)) [Add the user group to the local database.](#)

STEP 3 | ([Kerberos SSO only](#)) Create a [Kerberos](#) keytab for the firewall if Kerberos single sign-on (SSO) is the primary authentication service.

[Create a Kerberos keytab.](#) A keytab is a file that contains Kerberos account information for the firewall. To support Kerberos SSO, your network must have a [Kerberos](#) infrastructure.

STEP 4 | Configure an authentication profile.

Define one or both of the following:

- **Kerberos SSO**—The firewall first tries SSO authentication. If that fails, it falls back to the specified authentication **Type**.
- **External authentication or local database authentication**—The firewall prompts the user to enter login credentials, and uses an external service or local database to authenticate the user.

1. Select **Device > Authentication Profile** and **Add** the authentication profile.
2. Enter a **Name** to identify the authentication profile.
3. Select the **Type** of authentication service.

If you use [Multi-Factor Authentication](#), the selected type applies only to the first authentication factor. You select services for additional MFA factors in the **Factors** tab.

If you select **RADIUS**, **TACACS+**, **LDAP**, or **Kerberos**, select the **Server Profile**.

If you select **LDAP**, select the **Server Profile** and define the **Login Attribute**. For Active Directory, enter **sAMAccountName** as the value.

If you select **SAML**, select the **IdP Server Profile**.

4. 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.
5. (**MFA only**) Select **Factors**, **Enable Additional Authentication Factors**, and **Add** the MFA server profiles you configured.

The firewall will invoke each MFA service in the listed order, from top to bottom.

6. 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 the firewall to [Map Users to Groups](#), from an LDAP-based directory service such as Active Directory. By default, the list is empty, meaning no users can authenticate.



You can also select custom groups defined in a [group mapping configuration](#).

7. (**Optional**) To modify the user information before the firewall sends the authentication request to the server, configure a **Username Modifier**.

- **%USERDOMAIN%\%USERINPUT%**—If the source does not include the domain (for example, it uses the sAMAccountName), the firewall adds the **User Domain** you specify before the username. If the source includes the domain, the firewall replaces that domain with the **User Domain**. If the **User Domain** is empty, the firewall removes the domain from the user information that the firewall receives from source before the firewall sends the request to the authentication server.



Because LDAP servers do not support backslashes in the sAMAccountName, do not use this option to authenticate with an LDAP server.

- **%USERINPUT%**—(Default) The firewall sends the user information to the authentication server in the format it receives from the source.
- **%USERINPUT%@%USERDOMAIN%**—If the source does not include the domain, the firewall adds the **User Domain** value after the username. If the source includes domain, the firewall replaces that domain with the **User Domain** value. If the **User Domain** is empty, the firewall removes the domain from the user information that the firewall receives from the source before the firewall sends the request to the authentication server.
- **None**—If you manually enter **None**:
 - For LDAP and Kerberos server profiles, the firewall uses the domain it receives from the source to select the appropriate authentication profile, then removes the domain when it sends the authentication request to the server. This allows you to include the **User Domain** during the authentication sequence but remove the domain before the firewall sends the

authentication request to the server. For example, if you are using an LDAP server profile and the `samAccountName` as the attribute, use this option so that the firewall does not send the domain to the authentication server that expects only a username and not a domain.

- For RADIUS server profiles:
 - If the source sends the user information in **domain\username** format, the firewall sends the user information to the server in the same format.
 - If the source sends the user information in **username@domain** format, the firewall normalizes the user information to the **domain\username** format before sending it to the server.
 - If the source sends only the username, the firewall adds the **User Domain** you specify before sending the information to the server in **domain\username** format.
 - For local databases, TACACS+, and SAML, the firewall sends the user information to the authentication server in the format it receives from the source.
8. Click **OK** to save the authentication profile.

STEP 5 | Configure an authentication sequence.

Required if you want the firewall to try multiple authentication profiles to authenticate users. The firewall 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.



To expedite the authentication process, Use domain to determine authentication profile: the firewall matches 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 uses that profile to authenticate the user. If the firewall does not find a match, or if you disable the option, the firewall tries the profiles in the top-to-bottom sequence.

3. **Add** each authentication profile. To change the evaluation order of the profiles, select a profile and **Move Up** or **Move Down**.
4. Click **OK** to save the authentication sequence.

STEP 6 | Assign the authentication profile or sequence to an administrative account for firewall administrators or to Authentication policy for end users.

- **Administrators**—Assign the authentication profile based on how you manage administrator authorization:

Authorization managed locally on the firewall—[Configure a Firewall Administrator Account](#).

Authorization managed on a SAML, TACACS+, or RADIUS server—Select **Device > Setup > Management**, edit the Authentication Settings, and select the **Authentication Profile**.

- **End users**—For the full procedure to configure authentication for end users, see [Configure Authentication Policy](#).

STEP 7 | Verify that the firewall can [Test Authentication Server Connectivity](#) to authenticate users.

Test Authentication Server Connectivity

The test authentication feature enables you to verify whether the firewall or Panorama can communicate with the authentication server specified in an authentication profile and whether an authentication request succeeds for a specific user. You can test authentication profiles that authenticate administrators who access the web interface or that authenticate end users who access applications through GlobalProtect or Captive Portal. You can perform authentication tests on the candidate configuration to verify the configuration is correct before committing.

STEP 1 | [Configure an authentication profile](#). You do not need to commit the authentication profile or server profile configuration before testing.

STEP 2 | Log into the firewall [CLI](#).

STEP 3 | ([Firewalls with multiple virtual systems](#)) Define the target virtual system that the test command will access.

This is required on firewalls with multiple virtual systems so that the test authentication command can locate the user you will test.

Define the target virtual system by entering:

```
admin@PA-3220> set system setting target-vsys <vsys-name>
```

For example, if the user is defined in vsys2, enter:

```
admin@PA-3220> set system setting target-vsys vsys2
```



The **target-vsys** option is per login session; the firewall clears the option when you log off.

STEP 4 | Test the authentication profile by entering the following command:

```
admin@PA-3220> 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**, enter:

```
admin@PA-3220> test authentication authentication-profile my-profile  
username bsimpson password
```



When running the **test** command, the names of authentication profiles and server profiles are case sensitive. Also, if an 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 ensures that the firewall sends the correct credentials 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 test output.

If the authentication profile is configured correctly, the output displays `Authentication succeeded`. If there is a configuration issue, the output displays information to help you troubleshoot the configuration.



The output results vary based on several factors related to the authentication type that you are testing as well as the type of issue. For example, RADIUS and TACACS+ use different underlying libraries, so the same issue that exists for both of these types will produce different errors. Also, if there is a network problem, such as using an incorrect port or IP address in the authentication server profile, the output error is not specific. This is because the test command cannot perform the initial handshake between the firewall and the authentication server to determine details about the issue.

Authentication Policy

Authentication policy enables you to authenticate end users before they can access services and applications. Whenever a user requests a service or application (such as by visiting a web page), the firewall evaluates Authentication policy. Based on the matching Authentication policy rule, the firewall then prompts the user to authenticate using one or more methods (factors), such as login and password, [Voice, SMS, Push, or One-time Password \(OTP\) authentication](#). For the first factor, users authenticate through a Captive Portal web form. For any additional factors, users authenticate through a [Multi-Factor Authentication \(MFA\)](#) login page.



To implement Authentication policy for GlobalProtect, refer to [Configure GlobalProtect to facilitate multi-factor authentication notifications](#).

After the user authenticates for all factors, the firewall evaluates [Security Policy](#) to determine whether to allow access to the service or application.

To reduce the frequency of authentication challenges that interrupt the user workflow, you can specify a timeout period during which a user authenticates only for initial access to services and applications, not for subsequent access. Authentication policy integrates with Captive Portal to record the timestamps used to evaluate the timeout and to enable user-based policies and reports.

Based on user information that the firewall collects during authentication, User-ID creates a new IP address-to-username mapping or updates the existing mapping for that user (if the mapping information has changed). The firewall generates User-ID logs to record the additions and updates. The firewall also generates an Authentication log for each request that matches an Authentication rule. If you favor centralized monitoring, you can configure reports based on User-ID or Authentication logs and forward the logs to Panorama or external services as you would for any other log types.

- [Authentication Timestamps](#)
- [Configure Authentication Policy](#)

Authentication Timestamps

When configuring an Authentication policy rule, you can specify a timeout period during which a user authenticates only for initial access to services and applications, not for subsequent access. Your goal is to specify a timeout that strikes a balance between the need to secure services and applications and the need to minimize interruptions to the user workflow. When a user authenticates, the firewall records a timestamp for the first authentication challenge (factor) and a timestamp for any additional [Multi-Factor Authentication \(MFA\)](#) factors. When the user subsequently requests services and applications that match an Authentication rule, the firewall evaluates the timeout specified in the rule relative to each timestamp. This means the firewall reissues authentication challenges on a per-factor basis when timeouts expire. If you [Redistribute User Mappings and Authentication Timestamps](#), all your firewalls will enforce Authentication policy timeouts consistently for all users.



The firewall records a separate timestamp for each MFA vendor. For example, if you use [Duo v2](#) and [PingID](#) servers to issue challenges for MFA factors, the firewall records one timestamp for the response to the Duo factor and one timestamp for the response to the PingID factor.

Within the timeout period, a user who successfully authenticates for one Authentication rule can access services or applications that other rules protect. However, this portability applies only to rules that trigger the same authentication factors. For example, a user who successfully authenticates for a rule that triggers TACACS+ authentication must authenticate again for a rule that triggers SAML authentication, even if the access requests are within the timeout period for both rules.

When evaluating the timeout in each Authentication rule and the global timer defined in the Captive Portal settings (see [Configure Captive Portal](#)), the firewall prompts the user to re-authenticate for whichever setting expires first. Upon re-authenticating, the firewall records new authentication timestamps for the rules and resets the time count for the Captive Portal timer. Therefore, to enable different timeout periods for different Authentication rules, set the Captive Portal timer to a value that is the same as or higher than the timeout in any rule.

Configure Authentication Policy

Perform the following steps to configure Authentication policy for end users who access services through Captive Portal. Before starting, ensure that your [Security Policy](#) allows users to access the services and URL categories that require authentication.

STEP 1 | Configure Captive Portal. If you use [Multi-Factor Authentication](#) (MFA) services to authenticate users, you must set the **Mode** to **Redirect**.

STEP 2 | Configure the firewall to use one of the following services to authenticate users.

- [External Authentication Services](#)—Configure a server profile to define how the firewall connects to the service.
- [Local database authentication](#)—Add each user account to the local user database on the firewall.
- [Kerberos single sign-on \(SSO\)](#)—Create a Kerberos keytab for the firewall. Optionally, you can configure the firewall to use Kerberos SSO as the primary authentication service and, if SSO failures occur, fall back to an external service or local database authentication.

STEP 3 | Configure an Authentication Profile and Sequence for each set of users and Authentication policy rules that require the same authentication services and settings.

Select the **Type** of authentication service and related settings:

- **External service**—Select the **Type** of external server and select the **Server Profile** you created for it.
- **Local database authentication**—Set the **Type** to **Local Database**. In the **Advanced** settings, **Add** the Captive Portal users and user groups you created.
- **Kerberos SSO**—Specify the **Kerberos Realm** and **Import** the **Kerberos Keytab**.

STEP 4 | Configure an authentication enforcement object.

The object associates each authentication profile with a Captive Portal method. The method determines whether the first authentication challenge (factor) is transparent or requires a user response.

1. Select **Objects > Authentication** and **Add** an object.
2. Enter a **Name** to identify the object.
3. Select an **Authentication Method** for the authentication service **Type** you specified in the authentication profile:
 - **browser-challenge**—Select this method if you want the client browser to respond to the first authentication factor instead of having the user enter login credentials. For this method, you must have configured Kerberos SSO in the authentication profile or [NT LAN Manager \(NTLM\)](#) authentication in the Captive Portal settings. If the browser challenge fails, the firewall falls back to the **web-form** method.
 - **web-form**—Select this method if you want the firewall to display a Captive Portal web form for users to enter login credentials.
4. Select the **Authentication Profile** you configured.
5. Enter the **Message** that the Captive Portal web form will display to tell users how to authenticate for the first authentication factor.
6. Click **OK** to save the object.

STEP 5 | Configure an Authentication policy rule.

Create a rule for each set of users, services, and URL categories that require the same authentication services and settings.

1. Select **Policies** > **Authentication** and **Add** a rule.
2. Enter a **Name** to identify the rule.
3. Select **Source** and **Add** specific zones and IP addresses or select **Any** zones or IP addresses.

The rule applies only to traffic coming from the specified IP addresses or from [interfaces in the specified zones](#).

4. Select **User** and select or **Add** the source users and user groups to which the rule applies (default is **any**).
5. Select or **Add** the [Host Information Profiles](#) to which the rule applies (default is **any**).
6. Select **Destination** and **Add** specific zones and IP addresses or select **any** zones or IP addresses.

The IP addresses can be resources (such as servers) for which you want to control access.

7. Select **Service/URL Category** and select or **Add** the [services and service groups](#) for which the rule controls access (default is **service-http**).
8. Select or **Add** the [URL Categories](#) for which the rule controls access (default is **any**). For example, you can create a custom URL category that specifies your most sensitive internal sites.
9. Select **Actions** and select the **Authentication Enforcement** object you created.
10. Specify the **Timeout** period in minutes (default 60) during which the firewall prompts the user to authenticate only once for repeated access to services and applications.



Timeout is a tradeoff between tighter security (less time between authentication prompts) and the user experience (more time between authentication prompts). More frequent authentication is often the right choice for access to critical systems and sensitive areas such as a data center. Less frequent authentication is often the right choice at the network perimeter and for businesses for which the user experience is key.

11. Click **OK** to save the rule.

STEP 6 | (MFA only) Customize the MFA login page.

The firewall displays this page so that users can authenticate for any additional MFA factors.

STEP 7 | Verify that the firewall enforces Authentication policy.

1. Log in to your network as one of the source users specified in an Authentication policy rule.
2. Request a service or URL category that matches one specified in the rule.

The firewall displays the Captive Portal web form for the first authentication factor. For example:



If you configured the firewall to use one or more MFA services, authenticate for the additional authentication factors.

3. End the session for the service or URL you just accessed.
4. Start a new session for the same service or application. Be sure to perform this step within the **Timeout** period you configured in the Authentication rule.

-
- The firewall allows access without re-authenticating.
5. Wait until the **Timeout** period expires and request the same service or application.

The firewall prompts you to re-authenticate.

STEP 8 | (Optional) Redistribute User Mappings and Authentication Timestamps to other firewalls that enforce Authentication policy to ensure they all apply the timeouts consistently for all users.

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).  <i>To unlock users, use the following operational command:</i> <pre>> request authentication [unlock-admin unlock-user]</pre>	<pre>PA-220> show authentication locked- users { vsys <value> auth-profile <value> is-seq {yes no} {auth-profile vsys} <value> }</pre>
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, locked user accounts, and Multi-Factor Authentication (MFA) requests. • show-pending-requests displays the number of pending checks for authentication requests, allow lists, locked user accounts, and MFA requests. • 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:	<pre>PA-220> 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 </pre>

Task	Command
- 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.	<pre> 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 } </pre>
Test the connection and validity of the certificate profile .	<pre> PA-220> test authentication authentication-profile auth-profile username <username>password <password> </pre>
Troubleshoot a specific authentication using the Authentication ID displayed in Monitor > Logs > Authentication .	<pre> PA-220> grep <Authentication ID> </pre>

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
- > Default Trusted Certificate Authorities (CAs)
- > 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 ciphertext. 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, multi-factor authentication (MFA), 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).
- External dynamic list (EDL) validation.
- User-ID agent and TS agent access.
- 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 1: 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 Authentication policy identifies users who access HTTPS resources, designate a server certificate for the Captive Portal interface. If you configure Captive Portal to use certificates for identifying users (instead of, or in addition to, interactive authentication), deploy client certificates 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


Key/Certificate Usage	Description
	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.  Beginning in PAN-OS 8.0, firewalls use the Elliptic-Curve Diffie-Hellman Ephemeral (ECDHE) algorithm to perform strict certificate checking. This means that if the firewall uses an intermediate certificate, you must reimport the certificate from your web server to the firewall after you upgrade to a PAN-OS 8.0 or later release and combine the server certificate with the intermediate certificate (install a chained certificate). Otherwise, SSL Inbound Inspection sessions that have an intermediate certificate in the chain will fail. To install a chained certificate: 1. Open each certificate (.cer) file in a plain-text editor such as Notepad. 2. Paste each certificate end-to-end with the Server Certificate at the top with each signer included below. 3. Save the file as a text (.txt) or certificate (.cer) file (the name of the file cannot contain blank spaces). 4. Import the combined (chained) certificate into the firewall.
SSL Exclude Certificate	Certificates for servers to exclude from SSL/TLS decryption. For example, if you enable SSL decryption but your network includes servers for which the firewall should not decrypt traffic (for example, web services for your HR systems), import the corresponding certificates onto the firewall and configure them as SSL Exclude Certificates. See Decryption Exclusions .
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.  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 .

Key/Certificate Usage	Description
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.
Inter-Device Communication	By default, Panorama, firewalls, and Log Collectors use a set of predefined certificates for the SSL/TLS connections used for management and log forwarding. However, you can enhance these connection by deploying custom certificates to the devices in your deployment. These certificates can also be used to secure the SSL/TLS connection between Panorama HA peers.

Default Trusted Certificate Authorities (CAs)

The firewall trusts the most common and trusted authorities (CAs) by default. These trusted certificate providers are responsible for issuing the certificates the firewall requires to secure connections to the internet.

To view and manage the list of CAs that the firewall trusts by default, select **Device > Certificate Management > Certificates > Default Trusted Certificate Authorities**:



The only additional CAs you might want to add are trusted enterprise CAs that your organization requires—see [Obtain Certificates](#).

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, to verify that the certificate hasn't been revoked.

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.



If you configure multiple CRL distribution points (CDPs) and the firewall cannot reach the first CDP, the firewall does not check the remaining CDPs. To redirect invalid CRL requests, [configure a DNS proxy as an alternate server](#).

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](#).

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 (if you are configuring the firewall as an OCSP responder).
- ❑ Enable the HTTP OCSP service on the firewall (if you are configuring the firewall as an OCSP responder).
- ❑ 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, 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.



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).

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). If your enterprise has its own public key infrastructure (PKI), you can use external OCSP responders or you can configure the firewall itself as an OCSP responder. For details on OCSP, see [Certificate Revocation](#).

STEP 1 | Define an external OCSP responder or configure the firewall itself as 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. You can enter an IPv4 or IPv6 address. 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 | If you want the firewall to use the management interface for the OCSP responder interface, enable OCSP communication on the firewall. Otherwise, continue to the next step to configure an alternate interface.

1. Select **Device > Setup > Management**.
2. In the Management Interface Settings section, edit to select the **HTTP OCSP** check box, then click **OK**.

STEP 3 | To use an alternate interface as the OCSP responder interface, [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 OSCP services. The OSCP **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.

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.

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.

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 | 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).

In a high availability (HA) configuration, you must use the same master key on both firewalls or Panorama in the pair. Otherwise, HA 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.

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](#).

STEP 1 | (HA only) Disable HA.

This step is required before you can deploy a new master key to a firewall HA pair. If you do not disable HA before deploying a new master key, Panorama will lose connectivity to the primary firewall.

1. Select **Device > High Availability > General** and edit the Setup.
2. Disable (clear) the **Enable HA** setting and click **OK**.
3. **Commit** your configuration changes.

STEP 2 | Select **Device > Master Key and Diagnostics** and edit the Master Key section.

STEP 3 | Enter the **Current Master Key** if one exists.

STEP 4 | Define a new **New Master Key** and then **Confirm New Master Key**. The key must contain exactly 16 characters.

STEP 5 | To specify the master key **Lifetime**, 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 6 | 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 7 | Enable **Auto Renew Master Key** to configure the firewall to automatically renew the master key. To configure **Auto Renew With Same Master Key**, specify the number of **Days** and/or **Hours** to renew the same master key. The key extension allows the firewall to remain operational and continue securing your network; it is not a replacement for configuring a new key if the existing master key lifetime expires soon.



Consider the number of days until your next available maintenance window when configuring the master key to automatically renew after the lifetime of the key expires.

STEP 8 | (Optional) For added security, select whether to use an **HSM** to encrypt the master key. For details, see [Encrypt a Master Key Using an HSM](#).

STEP 9 | Click **OK** and **Commit**.

STEP 10 | (HA only) Re-enable HA.

1. Select **Device > High Availability > General** and edit the Setup.
2. Select **Enable HA** and click **OK**.
3. **Commit** your configuration changes.

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.

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 63 characters on the firewall or up to 31 characters on Panorama. 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.

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 endpoints](#).

STEP 5 | Enter a **Certificate Name**. The name is case-sensitive and can have up to 63 characters on the firewall or up to 31 characters on Panorama. 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.

You cannot use a [hardware security module \(HSM\)](#) to store ECDSA keys used for SSL/TLS [Decryption](#).

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**.



Client certificates that are used when requesting firewall services that rely on TLSv1.2 (such as administrator access to the web interface) cannot have sha512 as a digest algorithm. The client certificates must use a lower digest algorithm (such as sha384) or you must limit the Max Version to TLSv1.1 when you [Configure an SSL/TLS Service Profile](#) for the firewall services.

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, because the host name populates the Subject Alternate Name (SAN) field of the certificate and some browsers require the SAN to specify the domains the certificate protects; in addition, the Host Name matching the Common Name is mandatory for GlobalProtect.

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.

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.

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 63 characters on the firewall or up to 31 characters on Panorama. 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.

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.

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 63 characters on the firewall or up to 31 characters on Panorama. 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 should 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.

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**.

Deploy Certificates Using SCEP

If you have a Simple Certificate Enrollment Protocol (SCEP) server in your enterprise PKI, you can configure a SCEP profile to automate the generation and distribution of unique client certificates. SCEP operation is dynamic in that the enterprise PKI generates a user-specific certificate when the SCEP client requests it and sends the certificate to the SCEP client. The SCEP client then transparently deploys the certificate to the client device.

You can use a SCEP profile with [GlobalProtect](#) to assign user-specific client certificates to each GlobalProtect user. In this use case, the GlobalProtect portal acts as a SCEP client to the SCEP server in your enterprise PKI. Additionally, you can use a SCEP profile to assign client certificates to [Palo Alto Networks devices for mutual authentication](#) with other Palo Alto Networks devices for management access and inter-device communication.

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.

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 in the PKI infrastructure and then enter the password into the Password field.
- **Dynamic**—Enter a username and password of your choice (possibly the credentials of the PKI administrator) and the SCEP **Server URL** where the portal-client submits these credentials. The uses the credentials to authenticate with the SCEP server which transparently generates an OTP password for the portal upon each certificate request. (You can see this OTP change after a screen

refresh in The enrollment challengepassword is field after each certificate request.) The PKI transparently passes each new password to the portal, which then uses the password for its certificate request.

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.

You can include additional information about the client device or user by specifying tokens in the **Subject** name of the certificate.

The portal includes the token value and host ID in the CSR request to the SCEP server.

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. Enter the **Subject** name to use in the certificates generated by the SCEP server. The subject must be a distinguished name in the `<attribute>=<value>` format and must include a common name (CN) attribute (`CN=<variable>`). The CN supports the following dynamic tokens:
 - **\$USERNAME**—Use this token to enable the portal to request certificates for a specific user. To use this variable with GlobalProtect, you must also [Enable Group Mapping](#). The username entered by the user must match the name in the user-group mapping table.
 - **\$EMAILADDRESS**—Use this token to request certificates associated with a specific email address. To use this variable, you must also [Enable Group Mapping](#) and configure the **Mail Attributes** in the Mail Domains section of the Server Profile. If GlobalProtect cannot identify an email address for the user, it generates a unique ID and populates the CN with that value.
 - **\$HOSTID**—To request certificates for the device only, specify the host ID token. When a user attempts to log in to the portal, the endpoint sends identifying information that includes its host ID value. The host ID value varies by device type, either GUID (Windows) MAC address of the interface (Mac), Android ID (Android devices), UDID (iOS devices), or a unique name that GlobalProtect assigns (Chrome).
 - **\$UDID**—Use the UDID common name attribute to request certificates based on the client's device UDID for GlobalProtect or device serial number for mutual authentication between Palo Alto Networks devices.

When the GlobalProtect portal pushes the SCEP settings to the agent, the CN portion of the subject name is replaced with the actual value (username, host ID, or email address) of the certificate owner (for example, `O=acme, CN=johndoe`).

4. Select the **Subject Alternative Name Type**:



*Use static entries for the Subject Alternative Name Type. The firewall does not support dynamic tokens such as **\$USERNAME**.*

- **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 2,048 bits or larger.

- Select the **Digest for CSR** which indicates the digest algorithm for the certificate signing request (CSR): sha1, sha256, or sha384.

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 firewall. 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 firewall 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**.

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.

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](#)
- [Enable SSL Between GlobalProtect LSVPN Components](#) to configure GlobalProtect agent/app authentication to portals and gateways
- [SSL Forward Proxy](#) decryption
- [Obtain a Certificate from an External CA](#)

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.

Configure a Certificate Profile

Certificate profiles define user and device authentication for Captive Portal, multi-factor authentication (MFA), GlobalProtect, site-to-site IPSec VPN, external dynamic list (EDL) validation, Dynamic DNS (DDNS), User-ID agent and TS agent access, 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 Certificate Revocation List (CRL) status verification for certificate profiles to verify that the certificate hasn't been revoked. Enable both OCSP and CRL so that if the OCSP server isn't available, the firewall uses CRL. For details on these methods, see [Certificate Revocation](#).

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 to 3).

STEP 2 | Identify the certificate profile.

1. Select **Device > Certificate Management > Certificate 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 63 characters on the firewall or up to 31 characters on Panorama 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 "Authority Information Access" (AIA) information from the certificate to extract the OCSP responder information. To override the AIA information, 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.

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 **Block session if certificate status is unknown**. Otherwise, the firewall allows the sessions.
 5. If you want the firewall to block sessions after it registers an OCSP or CRL request timeout, select **Block session if certificate status cannot be retrieved within timeout**. Otherwise, the firewall allows the sessions.
 6. (**GlobalProtect only**) If you want the firewall to block sessions when the serial number attribute in the subject of the client certificate does not match the **host ID** that the GlobalProtect app reports for the endpoint, select **Block sessions if the certificate was not issued to the authenticating device**.

STEP 5 | Click **OK** and **Commit**

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.

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**.



As a best practice, set the Min Version to TLSv1.2 and the Max Version to Max.

On firewalls in FIPS/CC mode running PAN-OS 8.0 or a later release, TLSv1.1 is the earliest supported TLS version; do not select TLSv1.0.

Client certificates that are used when requesting firewall services that rely on TLSv1.2 cannot have SHA512 as a digest algorithm. The client certificates must use a lower digest algorithm (such as SHA384) or you must limit the Max Version to TLSv1.1 for the firewall services.

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.

To secure management traffic, you must also [Configure Administrative Accounts and Authentication](#).

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 TLSv1.2 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**.

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:

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 and the hashing algorithm for the certificates it generates to establish SSL proxy sessions with clients based on the destination server certificate. If the destination server uses a 1,024-bit RSA key, the firewall generates a certificate with a 1,024-bit RSA key. 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. If the destination server uses the SHA-1 hashing algorithm, the firewall generates a certificate with the SHA-1 hashing algorithm. If the destination server uses a hashing algorithm stronger than SHA-1, the firewall generates a certificate with the 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.

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.

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**.

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 and 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-3200 Series, PA-5200 Series, PA-7000 Series, and VM-Series firewalls and with the Panorama management server (both virtual and M-Series appliances) for use with the following HSM vendors:

- **nCipher nShield Connect**—PAN-OS 9.0 and 8.1 support nCipher nShield client version 12.30. PAN-OS 8.0 and earlier releases support client version 11.62, instead.
- **SafeNet Network**—The supported client versions depend on the PAN-OS release:
 - PAN-OS 9.0—SafeNet Network client versions 5.4.2 and 6.3.
 - PAN-OS 8.1—SafeNet Network client versions 5.4.2 and 6.2.2.
 - PAN-OS 8.0.2 and later PAN-OS 8.0 releases (also PAN-OS 7.1.10 and later PAN-OS 7.1 releases)—SafeNet Network client versions 5.2.1, 5.4.2, and 6.2.2.

The HSM server version must be compatible with these client versions. Refer to the HSM vendor documentation for the client-server version compatibility matrix. On the firewall or Panorama, use the following procedure to select the SafeNet Network client version that is compatible with your SafeNet HSM server.



Downgrading HSM servers might not be an option after you upgrade them.

- [Set Up Connectivity with a SafeNet Network HSM](#)
- [Set Up Connectivity with a Thales nShield Connect HSM](#)
- Install the SafeNet Client RPM Packet Manager.
 1. Select **Device > Setup > HSM** and **Select HSM Client Version** (Hardware Security Operations settings).
 2. Select **Version 5.4.2** (default) or **6.2.2** as appropriate for your HSM server version.
 3. Click **OK**.
 4. **(Required only if you change the HSM version on the firewall)** If the version change succeeds, the firewall prompts you to reboot to change to the new HSM version. If prompted, click **Yes**.
 5. If the master key isn't on the firewall, the client version upgrade will fail. **Close** the message and make the master key local to the firewall:

- Edit the Hardware Security Module Provider and disable (clear) the **Master Key Secured by HSM** option.
- Click **OK**.
- Select **Device > Master Key and Diagnostics** to edit the Master Key.
- Enter the **Current Master Key**; you can then enter that same key to be the **New Master Key** and then **Confirm New Master Key**.
- Click **OK**.
- Repeat the first four steps to **Select HSM Client Version** and reboot again.

Set Up Connectivity with a SafeNet Network HSM

To set up connectivity between the Palo Alto Networks firewall (HSM client) and a SafeNet Network HSM server, you must specify the IP address of the server, enter a password for authenticating the firewall to the server, and then register the firewall with the server. Before you begin configuring your HSM client, create a partition for the firewall on the HSM server and then confirm that the SafeNet Network client version on the firewall is compatible with your SafeNet Network HSM server (see [Set Up Connectivity with an HSM](#)).

Before the HSM and firewall connect, the HSM authenticates the firewall based on the firewall IP address. Therefore, you must [configure the firewall](#) to use a static IP address—not a dynamic address assigned through DHCP. Operations on the HSM stop working if the firewall IP address changes during runtime.



HSM configurations are not synchronized between high availability (HA) firewall peers. Consequently, you must configure the HSM separately on each peer. In active/passive HA configurations, you must [manually perform one failover](#) to individually configure and authenticate each HA peer to the HSM. After this initial manual failover, user interaction is not required for failover to function properly.

STEP 1 | Define connection settings for each SafeNet Network HSM.

1. Log in to the firewall web interface and select **Device > Setup > HSM**.
2. Edit the Hardware Security Module Provider settings and set the **Provider Configured to SafeNet Network HSM**.
3. **Add** each HSM server as follows. A high availability (HA) HSM configuration requires at least two servers; you can have a cluster of up to 16 HSM servers. All HSM servers in the cluster must run the same SafeNet version and must authenticate separately. You should use a SafeNet cluster only when you want to replicate the keys across the cluster. Alternatively, you can add up to 16 SafeNet HSM servers to function independently.
 1. Enter a **Module Name** (an ASCII string of up to 31 characters) for the HSM server.
 2. Enter an IPv4 address for the HSM **Server Address**.
4. (**HA only**) Select **High Availability**, specify the **Auto Recovery Retry** value (maximum number of times the HSM client tries to recover its connection to an HSM server before failing over to an HSM HA peer server; range is 0 to 500; default is 0), and enter a **High Availability Group Name** (an ASCII string up to 31 characters long).



If you configure two or more HSM servers, the best practice is to enable High Availability. Otherwise the firewall does not use the additional HSM servers.

5. Click **OK** and **Commit** your changes.

STEP 2 | (Optional) Configure a service route to connect to the HSM if you don't want the firewall to connect through the Management interface (default).



If you configure a service route for the HSM, running the **clear session all** CLI command clears all existing HSM sessions, which brings all HSM states down and then up again. During the several seconds required for HSM to recover, all SSL/TLS operations will fail.

1. Select **Device > Setup > Services** and click **Service Route Configuration**.
2. **Customize** a service route. The **IPv4** tab is active by default.
3. Click **HSM** in the Service column.
4. Select a **Source Interface** for the HSM.
5. Click **OK** and **Commit** your changes.

STEP 3 | Configure the firewall to authenticate to the HSM.

1. Select **Device > Setup > HSM** and **Setup Hardware Security Module**.
2. Select the HSM **Server Name**.
3. Enter the **Administrator Password** to authenticate the firewall to the HSM.
4. Click **OK**.

The firewall tries to authenticate to the HSM and displays a status message.

5. Click **OK** again.

STEP 4 | Register the firewall as an HSM client with the HSM server and assign the firewall to a partition on the HSM server.



If the HSM has a firewall with the same `<cl-name>` already registered, you must first remove the duplicate registration by running the **client delete -client <cl-name>** command, where `<cl-name>` is the name of the registered client (firewall) you want to delete.

1. Log in to the HSM from a remote system.
2. Register the firewall using the **client register -c <cl-name> -ip <fw-ip-addr>** CLI command, 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 for that firewall.
3. Assign a partition to the firewall using the **client assignpartition -c <cl-name> -p <partition-name>** CLI command, where `<cl-name>` is the name you assigned to the firewall using the **client register** command and `<partition-name>` is the name of a previously configured partition that you want to assign to this firewall.

STEP 5 | Configure the firewall to connect to the HSM partition.

1. Select **Device > Setup > HSM** and refresh () the display.
2. **Setup HSM Partition** (Hardware Security Operations settings).
3. Enter the **Partition Password** to authenticate the firewall to the partition on the HSM.
4. Click **OK**.

STEP 6 | (HA only) Repeat the previous authentication, registration, and partition connection steps to add another HSM to the existing HA group.



If you remove an HSM from your configuration, repeat the previous partition connection step to remove the deleted HSM from the HA group.

STEP 7 | Verify firewall connectivity and authentication with the HSM.

1. Select **Device > Setup > HSM** and check the authentication and connection Status:

- **Green**—The firewall is successfully authenticated and connected to the HSM.
 - **Red**—The firewall failed to authenticate to the HSM or network connectivity to the HSM is down.
2. View the following columns in Hardware Security Module Status to determine the authentication status:
 - **Serial Number**—The serial number of the HSM partition if the firewall successfully authenticated to the HSM.
 - **Partition**—The partition name on the HSM that is assigned to the firewall.
 - **Module State**—The current state of the HSM connection. This value is always **Authenticated** if the Hardware Security Module Status displays the HSM.

Set Up Connectivity with an nCipher nShield Connect HSM

You must set up a remote file system (RFS) as a hub to synchronize key data for all firewalls (HSM clients) in your organization that use the nCipher nShield Connect HSM. To ensure the nShield Connect client version on your firewalls is compatible with your nShield Connect server, see [Set Up Connectivity with an HSM](#).

Before the HSM and firewalls connect, the HSM authenticates the firewalls based on their IP addresses. Therefore, you must [configure the firewalls](#) to use static IP addresses—not dynamic addresses assigned through DHCP. (Operations on the HSM stop working if a firewall IP address changes during runtime).



HSM configurations are not synchronized between high availability (HA) firewall peers. Consequently, you must configure the HSM separately on each peer. In active/passive HA configurations, you must [manually perform one failover](#) to individually configure and authenticate each HA peer to the HSM. After this initial manual failover, user interaction is not required for failover to function properly.

STEP 1 | Define connection settings for each nCipher nShield Connect HSM.

1. Log in to the firewall web interface and select **Device > Setup > HSM**.
2. Edit the Hardware Security Module Provider settings and set the **Provider Configured to nShield Connect**.
3. **Add** each HSM server as follows. An HA HSM configuration requires two servers.
 1. Enter a **Module Name** for the HSM server. This can be any ASCII string of up to 31 characters.
 2. Enter an IPv4 address for the HSM **Server Address**.
4. Enter an IPv4 address for the **Remote Filesystem Address**.
5. Click **OK** and **Commit** your changes.

STEP 2 | (Optional) Configure a service route to connect to the HSM if you don't want the firewall to connect through the Management interface (default).



*If you configure a service route for the HSM, running the **clear session all** CLI command clears all existing HSM sessions, which brings all HSM states down and then up again. During the several seconds required for HSM to recover, all SSL/TLS operations will fail.*

1. Select **Device > Setup > Services** and click **Service Route Configuration**.
2. **Customize** a service route. The **IPv4** tab is active by default.
3. Click **HSM** in the Service column.
4. Select a **Source Interface** for the HSM.
5. Click **OK** and **Commit** your changes.

STEP 3 | Register the firewall as an HSM client with the HSM server.

This step briefly describes the procedure for using the front panel interface of the nShield Connect HSM. For more details, refer to nCipher documentation.

1. Log in to the front panel display of the nCipher nShield Connect HSM.
2. Use the right-hand navigation button to select **System > System configuration > Client config > New client**.
3. Enter the firewall IP address.
4. Select **System > System configuration > Client config > Remote file system** and enter the IP address of the client computer where you set up the RFS.

STEP 4 | Configure the RFS to accept connections from the firewall.

1. Log in to the RFS from a Linux client.
2. Obtain the electronic serial number (ESN) and the hash of the K_{NETI} key, which authenticates the HSM to clients, by running the **anonkneti <ip-address>** CLI command, where <ip-address> is the HSM IP address.

For example:

```
anonkneti 192.0.2.1
```

```
B1E2-2D4C-E6A2 5a2e5107e70d525615a903f6391ad72b1c03352c
```

In this example, B1E2-2D4C-E6A2 is the ESN and 5a2e5107e70d525615a903f6391ad72b1c03352c is the hash of the K_{NETI} key.

3. Use the following command from a superuser account to set up the RFS:

```
rfs-setup --force <ip-address> <ESN> <hash-Kneti-key>
```

The <ip-address> is the IP address of the HSM, <ESN> is the electronic serial number, and <hash-Kneti-key> is the hash of the K_{NETI} 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 HSM client submissions on the RFS:

```
rfs-setup --gang-client --write-noauth <FW-IPaddress>
```

where <FW-IPaddress> is the firewall IP address.

STEP 5 | Authenticate the firewall to the HSM.

1. In the firewall web interface, select **Device > Setup > HSM** and **Setup Hardware Security Module**.
2. Click **OK**.

The firewall tries to authenticate to the HSM and displays a status message.

3. Click **OK**.

STEP 6 | Synchronize the firewall with the RFS by selecting **Device > Setup > HSM** and **Synchronize with Remote Filesystem**.

STEP 7 | Verify firewall connectivity and authentication with the HSM.

1. Select **Device > Setup > HSM** and check the authentication and connection Status:
 - **Green**—The firewall is successfully authenticated and connected to the HSM.
 - **Red**—The firewall failed to authenticate to the HSM or network connectivity to the HSM is down.
2. Check the Hardware Security Module Status to determine the authentication status.
 - **Name**—The name of the HSM.
 - **IP address**—The IP address of the HSM.
 - **Module State**—The current state of the HSM connection: Authenticated or NotAuthenticated.

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](#).

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**.

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 nCipher nShield Connect HSMs.

STEP 1 | Log in to the firewall CLI.

STEP 2 | 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.

If you use the DHE or ECDHE key exchange algorithms to enable perfect forward secrecy (PFS) support for SSL decryption, you can use an HSM to store the private keys for SSL Inbound Inspection. You can also use an HSM to store ECDSA keys used for SSL Forward Proxy or SSL Inbound Inspection decryption.

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 | (nCipher nShield Connect only) Synchronize the key data from the nCipher 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. **Synchronize with Remote Filesystem** (Hardware Security Operations settings).

STEP 3 | Import the certificate that corresponds to the HSM-stored key.

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** your changes.

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** your changes.

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

You can perform the following tasks to manage your HSM deployment:

- 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.

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.

- > [HA Overview](#)
- > [HA Concepts](#)
- > [Set Up Active/Passive HA](#)
- > [Set Up Active/Active HA](#)
- > [Refresh HA1 SSH Keys and Configure Key Options](#)
- > [HA Firewall States](#)
- > [Reference: HA Synchronization](#)

HA Overview

You can set up two Palo Alto Networks firewalls as an HA pair; the HA peers should use the same version of PAN-OS and the same content version. 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](#).

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

- The [VM-Series firewall on Azure](#) and [VM-Series firewall on AWS](#) support active/passive HA only.
On AWS, when you deploy the firewall with the Amazon Elastic Load Balancing (ELB) service, it does not support HA (in this case, ELB service provides the failover capabilities).
- The VM-Series firewall on Google Cloud Platform does not support HA.

After you understand the [HA Concepts](#), proceed to [Set Up Active/Passive HA](#) or [Set Up Active/Active HA](#).

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.
- **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.

- For firewalls with dedicated HA ports, use these ports to manage communication and synchronization between the firewalls. For details, see [HA Ports on Palo Alto Networks Firewalls](#).
- For firewalls without dedicated HA ports such as the PA-220 and PA-220R firewalls, as a best practice use the management port for the HA1 port, and use the dataplane port for the HA1 backup.



For firewalls without dedicated HA ports, decide which ports to use for HA1 and HA1 backup based on your environment and understanding which are the least used and least congested. Assign HA1 to the best interface and HA1 backup to the other one.

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. ICMP is used to exchange heartbeats between HA peers. Ports used for HA1—TCP port 28769 and 28260 for clear text communication; port 28 for encrypted communication (SSH over TCP). If you enable encryption on the HA1 link, you can also Refresh HA1 SSH Keys and Configure Key Options.
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.

HA Links and Backup Links	Description
Backup Links	Provide redundancy for the HA1 and the HA2 links. In-band ports can be used for backup links for both HA1 and HA2 connections when dedicated backup links are not available. 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. • PA-3200 Series firewalls don't support an IPv6 address for the HA1-backup link; use an IPv4 address.  <i>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.</i>
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-800 Series, PA-3200 Series, and PA-5200 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-3200 Series, PA-5200 Series, and 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 Palo Alto Networks Firewalls

When connecting two Palo Alto Networks® firewalls in a high availability (HA) configuration, we recommend that you use the dedicated HA ports for [HA Links and Backup Links](#). These dedicated ports include: the HA1 ports labeled HA1, HA1-A, and HA1-B used for HA control and synchronization traffic; and HA2 and the High Speed Chassis Interconnect (HSCI) ports used for HA session setup traffic. The PA-5200 Series firewalls have multipurpose auxiliary ports labeled AUX-1 and AUX-2 that you can configure for HA1 traffic.

You can also configure the HSCI port for HA3, which is used for packet forwarding to the peer firewall during session setup and asymmetric traffic flow (active/active HA only). The HSCI port can be used for HA2 traffic, HA3 traffic, or both.



The HA1 and AUX 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.

If your firewall does not have dedicated HA ports, you can configure data ports as HA interfaces. If your firewall does have dedicated HA ports but does not have a dedicated HA backup port, you can also configure data ports as backups to dedicated HA ports.



Whenever possible, connect HA ports directly between the two firewalls in an HA pair (not through a switch or router) to avoid HA link and communications problems that could occur if there is a network issue.

Use the following table to learn about dedicated HA ports and how to connect the [HA Links and Backup Links](#):

Model	Front-Panel Dedicated Port(s)
PA-800 Series Firewalls	• HA1 and HA2—Ethernet 10Mbps/100Mbps/1000Mbps ports used for HA1 and HA2 in both HA Modes. • For HA1 traffic—Connect the HA1 port on the first firewall directly to the HA1 port on the second firewall in the pair or connect these ports together through a switch or router. • For HA2 traffic—Connect the HA2 port on the first firewall directly to the HA2 port on the second firewall in the pair or connect these ports together through a switch or router.
PA-3200 Series Firewalls	• HA1-A and HA1-B—Ethernet 10Mbps/100Mbps/1000Mbps ports used for HA1 traffic in both HA Modes. • For HA1 traffic—Connect the HA1-A port on the first firewall directly to the HA1-A port on the second firewall in the pair or connect them together through a switch or router. • For a backup to the HA1-A connection—Connect the HA1-B port on the first firewall directly to the HA1-B port on the second firewall in the pair or connect them together through a switch or router.  <i>If the firewall dataplane restarts due to a failure or manual restart, the HA1-B link will also restart. If this occurs and the HA1-A link is not connected and configured, then a split brain condition occurs. Therefore, we recommend that you connect and configure the HA1-A ports and the HA1-B ports to provide redundancy and to avoid split brain issues.</i>  <i>You cannot configure an HA1 or HA1-B connection on the data ports.</i> • HSCI—The HSCI port is a Layer 1 SFP+ interface that connects two PA-3200 Series firewalls in an HA configuration. Use this port for an HA2 connection, HA3 connection, or both. The traffic carried on the HSCI ports is raw Layer 1 traffic, which is not routable or switchable. Therefore, you must connect the HSCI ports directly to each other (from the HSCI port on the first firewall to the HSCI port on the second firewall).
PA-5200 Series Firewalls	• HA1-A and HA1-B—Ethernet 10Mbps/100Mbps/1000Mbps ports used for HA1 traffic in both HA Modes.

Model	Front-Panel Dedicated Port(s)
	• For HA1 traffic—Connect the HA1-A port on the first firewall directly to the HA1-A port on the second firewall in the pair or connect them together through a switch or router. • For a backup to the HA1-A connection—Connect the HA1-B port on the first firewall directly to the HA1-B port on the second firewall in the pair or connect them together through a switch or router. • HSCI—The HSCI port is a Layer 1 interface that connects two PA-5200 Series firewalls in an HA configuration. Use this port for an HA2 connection, HA3 connection, or both.  <i>The HSCI port on the PA-5220 firewall is a QSFP+ port and the HSCI port on the PA-5250, PA-5260, and PA-5280 firewalls is a QSFP28 port.</i> The traffic carried on the HSCI port is raw Layer 1 traffic, which is not routable or switchable. Therefore, you must connect the HSCI ports directly to each other (from the HSCI port on the first firewall to the HSCI port on the second firewall).
PA-5200 Series Firewalls (continued)	• AUX-1 and AUX-2—The auxiliary SFP+ ports are multipurpose ports that you can configure for HA1, management functions, or log forwarding to Panorama. Use these ports when you need a fiber connection for one of these functions. • For HA1 traffic—Connect the AUX-1 port on the first firewall directly to the AUX-1 port on the second firewall in the pair or connect them together through a switch or router. • For a backup to the AUX-1 connection—Connect the AUX-2 port on the first firewall directly to the AUX-2 port on the second firewall in the pair or connect them together through a switch or router.
PA-7000 Series Firewalls	• HA1-A and HA1-B—Ethernet 10Mbps/100Mbps/1000Mbps ports used for HA1 traffic in both HA Modes. • For HA1 traffic—Connect the HA1-A port on the first firewall directly to the HA1-A port on the second firewall in the pair or connect them together through a switch or router. • For a backup to the HA1-A connection—Connect the HA1-B port on the first firewall directly to the HA1-B port on the second firewall in the pair or connect them together through a switch or router.  <i>You cannot configure an HA1 connection on the NPC data ports or the management (MGT) port.</i> • HSCI-A and HSCI-B—The HSCI ports are Layer 1 QSFP+ interfaces that connect two PA-7000 Series firewalls in an HA configuration. Use these ports for an HA2 connection, HA3 connection, or both. The traffic carried on the HSCI ports is raw Layer 1 traffic, which is not routable or switchable. Therefore, you must connect these ports as follows: • For HA2 and HA3 traffic—Connect the HSCI-A port on the first firewall directly to the HSCI-A port on the second firewall.

Model	Front-Panel Dedicated Port(s)
	 For HA2 or HA2/HA3 traffic, the PA-7000 Series firewalls synchronize sessions across the NPCs one-for-one. • For a backup to the HSCI-A connection—Connect the HSCI-B port on the first firewall directly to the HSCI-B port on the second firewall.

Device Priority and Preemption

The firewalls in an Active-Passive HA pair can be assigned a *device priority* value to indicate a preference for which firewall should assume the active 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. The other firewall is the passive firewall.

The same is true for an Active-Active HA pair; however, the *device ID* is used to assign a device priority value. Similarly, the lower numerical value in device ID corresponds to a higher priority. The firewall with the higher priority becomes active-primary and the paired firewall becomes active-secondary.

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. By default, the interval for the heartbeat is 1000 milliseconds. 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 PA-3200 Series, PA-5200 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.

All firewall models except VM-Series firewalls support a pre-negotiation configuration, which depends 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 the step [\(Optional\) Enable LACP and LLDP Pre-Negotiation for Active/Passive HA for faster failover if your network uses LACP or LLDP](#).

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.](#))



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, PA-5200, and PA-3200 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:

7	6	5 4 3 2 1 0	7 6 5 4 3 2 1 0
Device-ID	0	Group-ID	Interface-ID


The format of the virtual MAC address on PA-7000, PA-5200, and PA-3200 Series firewalls is B4-0C-25-xx-xx-xx, where B4-0C-25 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:

7 6 5	4	3 2 1 0 7 6	5 4 3 2	1 0 7 6 5 4 3 2 1 0
111	Device-ID	Group-ID	0000	Interface-ID

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.



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 (Recommended/Aggressive) across the different hardware models; these values are for current reference only and can change in a subsequent release.

Timers	Description	PA-7000 Series PA-5200 Series PA-3200 Series	PA-800 Series PA-220 VM-Series	Panorama Virtual Appliance Panorama M-Series
Monitor Fail Hold Up Time (ms)	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 (min)	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

Timers	Description	PA-7000 Series PA-5200 Series PA-3200 Series	PA-800 Series PA-220 VM-Series	Panorama Virtual Appliance Panorama M-Series
Heartbeat Interval (ms)	Frequency at which the HA peers exchange heartbeat messages in the form of an ICMP (ping).	1000/1000	2000/1000	2000/1000
Promotion Hold Time (ms)	Time that the passive firewall (in active/passive mode) or the active-secondary firewall (in active/active mode) will wait before taking over as the active or active-primary firewall after communications with the HA peer have been lost. This hold time will begin only after the peer failure declaration has been made.	2000/500	2000/500	2000/500
Additional Master Hold Up Time (ms)	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 (ms)	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

Timers	Description	PA-7000 Series PA-5200 Series PA-3200 Series	PA-800 Series PA-220 VM-Series	Panorama Virtual Appliance Panorama M-Series
Maximum No. of Flaps	A flap is counted when one of the following occurs: • A preemption-enabled firewall leaves the active state within 20 minutes after becoming active. • A link or path fails to stay up for 10 minutes after becoming functional. In the case of a failed preemption or non-functional loop, this value indicates the maximum number of flaps that are permitted before the firewall is suspended (range 0-16; default 3).	3/3	3/3	Not Applicable

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 the Session Owner to First Packet reduces traffic across the HA3 link and helps distribute the dataplane load across peers.



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 Layer 2 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.

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



- ❑ The end host sends a packet to FW1.
- ❑ 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.

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.

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). The exceptions are AWS and Azure, where the management interface is configured as DHCP Client and it supports HA1 and HA1 Backup links.

- HA1: Dedicated HA1 port
HA1 Backup: Dedicated HA1 port
Recommendation: Enable Heartbeat Backup
- 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


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	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.

Independent Configuration Settings	PeerA	PeerB
	under load, as this could cause a a path monitoring failure and trigger a failover.	

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.

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.

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 or are on the same VLAN.

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.
 1. Select **Device > Certificate Management > Certificates**.
 2. Select **Export HA key**. Save the HA key to a network location that the peer can access.
 3. 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.
 4. Repeat this process on the second firewall to exchange HA keys on both devices.
2. Select **Device > High Availability > General**, edit the Control Link (HA1) section.
3. Select **Encryption Enabled**.



If you enable encryption, after you finish configuring the HA firewalls, you can [Refresh HA1 SSH Keys and Configure Key Options](#).

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**.



PA-3200 Series firewalls don't support an IPv6 address for the HA1 backup control link; use an IPv4 address.

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.

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) 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.

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:
 1. Select an AE interface in a Layer 2 or Layer 3 deployment.
 2. Select the **LACP** tab.
 3. Select **Enable in HA Passive State**.
 4. 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:
 1. Select an Ethernet interface in a virtual wire deployment.
 2. Select the **Advanced** tab.

3. Select the **LACP** tab.
4. Select **Enable in HA Passive State**.
5. Click **OK**.
5. To enable LLDP active pre-negotiation:
 1. Select an Ethernet interface in a Layer 2, Layer 3, or virtual wire deployment.
 2. Select the **Advanced** tab.
 3. Select the **LLDP** tab.
 4. Select **Enable in HA Passive State**.
 5. Click **OK**.



If you want to allow LLDP passive pre-negotiation for a virtual wire deployment, perform Step 14.e but do not enable LLDP itself.

STEP 15 | Save your configuration changes.

Click **Commit**.

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

Perform the following task to define failover conditions and thus establish what will cause a firewall in an HA pair to fail over, an event where the task of securing traffic passes from the previously active firewall to its HA peer. The [HA Overview](#) describes conditions that cause a failover.



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.

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 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 | **Commit** the configuration.

Verify Failover

To test that your HA configuration works properly, trigger a manual failover and verify that the firewalls transition states successfully.

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.

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. The PA-7000 Series firewalls use the HSCI port for HA3. The PA-5200 Series firewalls can use the HSCI port for HA3 or you can configure aggregate interfaces on the dataplane ports for HA3 for redundancy. On the remaining platforms, you can configure aggregate interfaces on dataplane ports 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.



If you have a switch located between your HA firewalls, the 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.

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.

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 model; choose the HA3 step that corresponds with your model.

- 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 PA-5200 Series firewalls (which have one HSCI port), connect the HSCI port on the first chassis to the HSCI port on the second chassis. You can also use data ports for HA3 on PA-5200 Series firewalls.
 - On PA-3200 Series firewalls (which have one HSCI port), connect the HSCI port on the first chassis to the HSCI port on the second chassis.
 - On any other hardware model, 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**.

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.

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.
 1. Select **Device > Certificate Management > Certificates**.
 2. Select **Export HA key**. Save the HA key to a network location that the peer can access.
 3. 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**.



If you enable encryption, after you finish configuring the HA firewalls, you can [Refresh HA1 SSH Keys and Configure Key Options](#).

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**.



PA-3200 Series firewalls don't support an IPv6 address for the HA1 backup control link; use an IPv4 address.

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. When an HA2 Keep-alive failure occurs, the system either generates a critical system log message or causes a split dataplane depending on your configuration.



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 sends the Keep-alive messages. The other firewall is notified if a failure occurs.



A split dataplane causes the dataplanes of both peers to operate independently while leaving the high-available state as Active-Primary and Active-Secondary. If only one firewall is configured to split dataplane, then split dataplane applies to the other device as well.

7. Edit the **Data Link (HA2 Backup)** section, select the interface, and add the **IPv4/IPv6 Address** and **Netmask**.
8. Click **OK**.

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 recovers post-failure (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**—The firewall performs an XOR operation on the source and destination IP addresses from the packet and based on the result, the firewall chooses which HA peer will set up the session.
 - **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 (recommended setting).



Start with **First Packet** for Session Owner and Session Setup, and then based on load distribution, you can change to one of the other options.

- **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**.

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 | [Define HA Failover Conditions](#).

STEP 20 | **Commit** the configuration.

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.



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 | **Commit** the configuration.

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.

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](#).



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**.

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


STEP 6 | Commit the configuration.

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.



STEP 1 | Perform Step 1 through Step 15 of [Configure Active/Active HA](#).

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](#).

STEP 5 | [Define HA Failover Conditions](#)

STEP 6 | **Commit** the configuration.

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.



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 recommend 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.

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](#).

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](#).

STEP 8 | **Commit** the configuration.

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.

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.



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. 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**.

STEP 7 | Enable jumbo frames on firewalls other than PA-7000 Series firewalls.

STEP 8 | Define HA Failover Conditions.

STEP 9 | Commit the configuration.

STEP 10 | 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 11 | 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**.

STEP 12 | 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 the **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 13 | Commit the configuration.

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.

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.

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 | Commit the configuration.

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.



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](#).

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 | **Commit** the configuration.

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 | **Commit** the configuration.

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.



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. 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 | Commit the configuration.

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**.

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 | Commit the configuration.

Refresh HA1 SSH Keys and Configure Key Options

When you configure [active/passive](#) or [active/active](#) high availability (HA), you can enable encryption for the HA1 (control link) connection between HA firewalls. You may want to secure the HA1 traffic between HA peers with encryption if the firewalls aren't located in the same site.

All Palo Alto Networks firewalls come with Secure Shell (SSH) pre-configured, and the HA firewalls can act as both SSH server and SSH client simultaneously. If you enable encryption on the HA1 control link, you can use the CLI commands that follow to secure the connection between the HA firewalls. You can change the default `host key type`, generate a new pair of public and private SSH host keys for the HA1 control link, and configure other HA1 encryption settings. You can enable the firewalls to use the new host keys and the following settings without restarting the HA peers, thus avoiding your firewalls going offline. The firewall reestablishes HA1 sessions with its peer to synchronize the new SSH key and other settings and also generates system logs (subtype is ha) for re-establishing HA1 and HA1-backup sessions.

The following examples show how to refresh (regenerate) your HA1 SSH keys and change and display various SSH settings for the HA1 control link after you enable encryption and [Access the CLI](#).



You must enable encryption and it must be functioning properly on an HA pair before you can perform the following tasks.



If you are configuring the HA1 control link with encryption in [FIPS-CC mode](#), you must set automatic rekeying parameters for session keys.

- (Optional) Set the default host key type.

If you enable encryption on the HA1 control link, the firewall uses a default host key type of RSA 2048 unless you change it. The HA1 SSH connection uses only the default `host key type` to authenticate the HA peers (before an encrypted session is established between them). You can change the default host key type; the choices are ECDSA 256, 384, or 521, or RSA 2048, 3072, or 4096. Change the default host key type if you prefer a longer RSA key length or if you prefer ECDSA rather than RSA. This example sets the default host key type to an ECDSA key of 256 bits. It also re-establishes the HA1 connection using the new host key without restarting the HA peers.

1. admin@PA-3250PA-3250> **configure**
2. admin@PA-3250# **set deviceconfig system ssh default-hostkey ha key-type ECDSA key-length 256**
3. admin@PA-3250# **commit**
4. admin@PA-3250# **exit**
5. admin@PA-3250> **request high-availability sync-to-remote ssh-key**



An HA connection must already be established between the HA firewalls. If the firewalls have not yet established an HA connection, you must enable encryption on the control link connection, export the HA key to a network location and import the HA key on the peer. See [Configure Active/Passive HA](#) or [Configure Active/Active HA](#).

6. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**
7. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**



You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the two HA peers. (Using the **force** option when an HA1 backup is configured has no effect.)

8. admin@PA-3250> **configure**
9. admin@PA-3250# **show deviceconfig system ssh default-hostkey**

- Establish when automatic rekeying of the session keys occurs for SSH over the HA1 control link by setting parameters.

The session keys are used for encrypting the traffic between the HA peers. After any one rekeying parameter reaches its configured value, SSH uses the new session encryption keys. The parameters are data, interval (seconds), and packet count. If you set more than one parameter, rekeying occurs when the first parameter reaches its configured value and then the firewall resets all rekeying parameters. You can set a second or third parameter if you aren't sure that one parameter you configured will reach its specified value as soon as you want rekeying to occur.

1. admin@PA-3250> **configure**
2. admin@PA-3250# **set deviceconfig system ssh session-rekey ha data 32**

Rekeying occurs after the volume of data (in megabytes) is transmitted following the previous rekeying. The default is based on the type of cipher you use and ranges from 1GB to 4GB; the range is 10MB to 4,000MB. Alternatively, you can use the **set deviceconfig system ssh session-rekey ha data default** command, which sets the data parameter to the default value of the individual cipher you are using.



If you are configuring the HA1 control link with encryption in FIPS-CC mode, you must set a data value (you cannot let it default) and the value must be no greater than 1000MB.

3. admin@PA-3250# **set deviceconfig system ssh session-rekey ha interval 3600**

Rekeying occurs after the specified time interval (in seconds) passes following the previous rekeying. By default, time-based rekeying is disabled (set to **none**); the range is 10 to 3,600.



If you are configuring the HA1 control link with encryption in FIPS-CC mode, you must set a time interval within the range; you cannot leave it disabled.

4. admin@PA-3250# **set deviceconfig system ssh session-rekey ha packets 27**

Rekeying occurs after the defined number of packets (2^n) are transmitted following the previous rekeying. Enter the exponent to which 2 is raised; for example, 14 configures that a maximum of 2^{14} packets are transmitted before a rekeying occurs. The default is 2^{28} ; the range is 12 to 27 (2^{12} to 2^{27}). Alternatively, you can use the **set deviceconfig system ssh session-rekey ha packets default** command, which sets the value to 2^{28} .



Choose rekeying parameters based on your type of traffic and network speeds (in addition to FIPS-CC requirements if they apply to you). Don't set the parameters so low that they affect SSH performance.

5. admin@PA-3250# **commit**
6. admin@PA-3250# **exit**
7. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**
8. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**



You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the two HA peers. (Using the **force** option when an HA1 backup is configured has no effect.)

9. admin@PA-3250> **configure**

10.admin@PA-3250# **show deviceconfig system ssh session-rekey ha**

- (Optional) Set the SSH server to use the specified encryption ciphers for the HA1 sessions.

HA1 SSH allows all supported ciphers by default. When you set one or more ciphers, the SSH server advertises only those ciphers while connecting, and if the SSH client (HA peer) tries to connect using a different cipher, the server terminates the connection.

1. admin@PA-3250> **configure**

2. admin@PA-3250# **set deviceconfig system ssh ciphers ha cipher**

aes128-cbc—AES 128-bit cipher with Cipher Block Chaining

aes128-ctr—AES 128-bit cipher with Counter Mode

aes128-gcm—AES 128-bit cipher with GCM (Galois/Counter Mode)

aes192-cbc—AES 192-bit cipher with Cipher Block Chaining

aes192-ctr—AES 192-bit cipher with Counter Mode

aes256-cbc—AES 256-bit cipher with Cipher Block Chaining

aes256-ctr—AES 256-bit cipher with Counter Mode

aes256-gcm—AES 256-bit cipher with GCM

3. admin@PA-3250# **commit**

4. admin@PA-3250# **exit**

5. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**

6. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**



You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the HA peers. (Using the **force** option when an HA1 backup is configured has no effect.)

7. admin@PA-3250> **configure**

8. admin@PA-3250# **show deviceconfig system ssh ciphers ha**

- (Optional) Delete a cipher from the set of ciphers you selected for SSH over the HA1 control link.

This example deletes the AES CBC cipher with 128-bit key.

1. admin@PA-3250> **configure**

2. admin@PA-3250# **delete deviceconfig system ssh ciphers ha aes128-cbc**

3. admin@PA-3250# **commit**

4. admin@PA-3250# **exit**

5. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**

6. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**



You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the two HA peers. (Using the **force** option when an HA1 backup is configured has no effect.)

7. admin@PA-3250> **configure**
8. admin@PA-3250# **show deviceconfig system ssh ciphers ha**

- (Optional) Set the session key exchange algorithm for HA1 SSH.

By default, the SSH server (HA firewall) advertises all the key exchange algorithms to the SSH client (HA peer firewall).



If you are using an ECDSA default key type, the best practice is to use an ECDH key algorithm.

1. admin@PA-3250> **configure**
2. admin@PA-3250# **set deviceconfig system ssh kex ha value**
diffie-hellman-group14-sha1—Diffie-Hellman group 14 with SHA1 hash
ecdh-sha2-nistp256—Elliptic-Curve Diffie-Hellman over National Institute of Standards and Technology (NIST) P-256 with SHA2-256 hash
ecdh-sha2-nistp384—Elliptic-Curve Diffie-Hellman over NIST P-384 with SHA2-384 hash
ecdh-sha2-nistp521—Elliptic-Curve Diffie-Hellman over NIST P-521 with SHA2-521 hash
3. admin@PA-3250# **commit**
4. admin@PA-3250# **exit**
5. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**
6. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**



You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the two HA peers. (Using the **force** option when an HA1 backup is configured has no effect.)

- (Optional) Set the message authentication code (MAC) for HA1 SSH.

By default, the server advertises all of the MAC algorithms to the client.

1. admin@PA-3250> **configure**
2. admin@PA-3250# **set deviceconfig system ssh mac ha value**
hmac-sha1—MAC with SHA1 cryptographic hash
hmac-sha2-256—MAC with SHA2-256 cryptographic hash
hmac-sha2-512—MAC with SHA2-512 cryptographic hash
3. admin@PA-3250# **commit**
4. admin@PA-3250# **exit**
5. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**
6. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**



You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the two HA peers. (Using the **force** option has no effect when an HA1 backup is configured)

- Regenerate ECDSA or RSA host keys for HA1 SSH to replace the existing keys, and re-establish HA1 sessions between HA peers using the new keys without restarting the HA peers.

The HA peers use the host keys to authenticate each other. This example regenerates the ECDSA 256 default host key.



Regenerating a host key does not change your default host key type. To regenerate the default host key you are using, you must specify your default host key type and length when you regenerate. Regenerating a host key that isn't your default host key type simply regenerates a key that you aren't using and therefore has no effect.

1. admin@PA-3250> **configure**
2. admin@PA-3250# **set deviceconfig system ssh regenerate-hostkeys ha key-type ECDSA key-length 256**
3. admin@PA-3250# **commit**
4. admin@PA-3250# **exit**
5. admin@PA-3250> **request high-availability sync-to-remote ssh-key**



An HA connection must already be established between the HA firewalls. If the firewalls have not yet established an HA connection, you must enable encryption on the control link connection, export the HA key to a network location, and import the HA key on the peer. See [Configure Active/Passive HA](#) or [Configure Active/Active HA](#).

6. (HA1 Backup is configured) admin@PA-3250> **request high-availability session-reestablish**
7. (No HA1 Backup is configured or HA1 Backup link is down) admin@PA-3250> **request high-availability session-reestablish force**



You can force the firewall to reestablish HA1 sessions if there is no HA1 backup, which causes a brief split-brain condition between the two HA peers. (Using the **force** option when an HA1 backup is configured has no effect.)

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.
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

HA Firewall State	Occurs In	Description
		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 silently drop 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	The device is disabled so won't pass data traffic and although HA communications still occur, the device doesn't participate in the HA election process. It can't move to an HA functional state without user intervention.

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 (and associated certificates), Time Zone, Locale, Date, Time, Latitude, Longitude. • 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	You must activate the Virtual Systems license on each firewall in the pair to increase the number of virtual systems beyond the base number provided by default on PA-3200 Series, PA-5200 Series, and PA-7000 Series firewalls. You must also enable Multi Virtual System Capability on each firewall (Device > Setup > Management > General Settings).
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

Configuration Item	What Doesn't Sync in Active/Passive?
Services	Device > Setup > Services
Global Service Routes	Device > Setup > Services > Service Route Configuration
Telemetry and Threat Intelligence Settings	Device > Setup > Telemetry and Threat Intelligence
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 app 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

Configuration Item	What Doesn't Sync in Active/Passive?
Rule Usage Data	Rule usage data, such as hit count, Created, and Modified Dates, are not synced between peers. You need to log in to the each firewall to view the policy rule hit count data for each firewall or use Panorama to view information on the HA firewall peers.
Certificates for Device Management and Syslog Communication over SSL only	Device > Certificate Management > Certificates Certificates used for device management or for syslog communication over SSL don't synchronize with an HA peer.
Certificates in a Certificate Profile	Device > Certificate Management > Certificate Profile
SSL/TLS Service Profile for Device Management only	Device > Certificate Management > SSL/TLS Service Profile SSL/TLS Service Profile for Device Management doesn't synchronize with an HA peer.

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 (and associated certificates), Time Zone, Locale, Date, Time, Latitude, Longitude. 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-vsyz Capability	You must activate the Virtual Systems license on each firewall in the pair to increase the number of virtual systems beyond the base number provided by default on PA-3200 Series, PA-5200 Series, and PA-7000 Series firewalls. You must also enable Multi Virtual System Capability on each firewall (Device > Setup > Management > General Settings).
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

Configuration Item	What Doesn't Sync in Active/Active?
SNMP	Device > Setup > Operations > SNMP Setup
Services	Device > Setup > Services
Global Service Routes	Device > Setup > Services > Service Route Configuration
Telemetry and Threat Intelligence Settings	Device > Setup > Telemetry and Threat Intelligence
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
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 app 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).

Configuration Item	What Doesn't Sync in Active/Active?
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.
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.

Configuration Item	What Doesn't Sync in Active/Active?
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.)
Rule Usage Data	Rule usage data, such as hit count, Created, and Modified Dates, are not synced between peers. You need to log in to the each firewall to view the policy rule hit count data for each firewall or use Panorama to view information on the HA firewall peers.
Certificates for Device Management and Syslog Communication over SSL only	Device > Certificate Management > Certificates Certificates used for device management or for syslog communication over SSL don't synchronize with an HA peer.
Certificates in a Certificate Profile	Device > Certificate Management > Certificate Profile
SSL/TLS Service Profile for Device Management only	Device > Certificate Management > SSL/TLS Service Profile SSL/TLS Service Profile for Device Management doesn't synchronize with an HA peer.

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	
User Mappings across Virtual Systems	Yes	Yes	HA1	
User to IP Address Mappings	Yes	Yes	HA1	
DHCP Lease (as server)	Yes	Yes	HA1	If the PAN-OS versions on the HA peers don't match, the DHCP Lease (as server) config information won't sync.
DNS Cache	No	No	N/A	

Runtime Information	Config Synced?		HA Link	Details
	A/P	A/A		
FQDN Refresh	No	No	N/A	
IKE Keys (phase 2)	Yes	Yes	HA1	
Forward Information Base (FIB)	Yes	Yes	HA1	
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.
BrightCloud URL Database	No	No	N/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.
Content (manual sync)	Yes	Yes	HA1	
PPPoE, PPPoE Lease	Yes	Yes	HA1	
DHCP Client Settings and Lease	Yes	Yes	HA1	If the PAN-OS versions on the HA peers don't match, the DHCP Client Settings and Lease config information won't sync.
SSL VPN Logged in User List	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.  <i>A host session is a session</i>

Runtime Information	Config Synced?		HA Link	Details
	A/P	A/A		
				<i>terminated on one of the firewall interfaces, such as an ICMP session pinging one of the firewall interfaces or a GP tunnel.</i>
ARP Table	Yes	No	HA2	
Neighbor Discovery (ND) Table	Yes	No	HA2	
MAC Table	Yes	No	HA2	
IPSec Sequence Number (anti-replay)	Yes	Yes	HA2	
DoS Block List Entries	No	No	N/A	
Virtual MAC	Yes	Yes	HA2	
SCTP Associations	Yes	No	HA2	

Monitoring

To forestall potential issues and to accelerate incidence response when needed, the firewall provides intelligence about traffic and user patterns using 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. For example, you can 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 with which you can interact to find the information 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
- > Use the App Scope Reports
- > Use the Automated Correlation Engine
- > Take Packet Captures
- > Monitor Applications and Threats
- > View and Manage Logs
- > Monitor Block List
- > View and Manage Reports
- > View Policy Rule Usage
- > Use External Services for Monitoring
- > Configure Log Forwarding
- > Configure Email Alerts
- > Use Syslog for Monitoring
- > SNMP Monitoring and Traps
- > Forward Logs to an HTTP(S) Destination
- > 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.  <i>A Config installed entry indicates configuration changes were committed successfully.</i>
System Resources	Displays the Management CPU usage, Data Plane usage, and the Session Count, which displays the number of sessions established through the firewall.

Dashboard Charts	Descriptions
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.


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](#).

	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.
	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.
	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.
	Application View	The application view allows you filter the ACC view by either the sanctioned and unsanctioned applications in use on your network, or by the risk level of the applications in use on your network. Green indicates sanctioned applications, blue unsanctioned applications, and yellow indicates applications that are partially sanctioned. Partially sanctioned applications are those that have a mixed sanctioned state; it indicates that the application is inconsistently tagged as sanctioned, for example it might be sanctioned on one or more virtual systems on a firewall enabled for multiple virtual systems or across one or more firewalls within a device group on Panorama.
	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

ACC—First Look

		of blocked threats, compromised hosts or traffic to malware hosts/domains.
	Source	The data used for the ACC 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 data from 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 data from 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.
	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.
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

Tab	Description
	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.
Tunnel Activity	Displays the activity of tunnel traffic that the firewall inspected based on your tunnel inspection policies. Information includes tunnel usage based on tunnel ID, monitor tag, user, and tunnel protocols such as Generic Routing Encapsulation (GRE), General Packet Radio Service (GPRS) Tunneling Protocol for User Data (GTP-U), and non-encrypted IPSec.


You can also [Interact with the ACC](#) to create customized tabs with custom layout and widgets that meet your network monitoring needs, export the tab and share with another administrator.

ACC Widgets

The widgets on each tab are interactive; you can set the [ACC Filters](#) and drill down into the details for each table or graph, or customize the widgets included in the tab to focus on the information you need. For details on what each widget displays, see [Widget Descriptions](#).



Widgets

	View	You can sort the data by bytes, sessions, threats, count, content, URLs, malicious, benign, files, applications, data, profiles, objects, users. The available options vary by widget.
	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.
	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. 
	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

Widgets

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.	
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

Widget	Description
	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 endpoint. This information is sourced from entries in the HIP match log that are generated when the data submitted by the GlobalProtect app 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
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	

Widget	Description
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-5200 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
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

Widget	Description
	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 <i>application-default</i> 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
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

Widget	Description
	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, delete, export and import tabs, add and delete widgets, set local and global filters, and interact with the 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.



To save the tab as the default tab, select .

- Export and Import tabs.
 1. Select the tab, and click the pencil icon next to the tab name, to edit the tab.
 2. Select the  icon to export the current tab as a .txt file. You can share this .txt file with another administrator.

3. To import the tab as a new tab on another firewall, select the  icon along the list of tabs, and add a name and click the import icon, browse to select the .txt file.




- 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.
- 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.

- 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.
- 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.



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).



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**).

Received Time	Type	Name	Attacker	Attacker Name	Victim	To Port	Application	Action	Severity
02/02 10:17:02	vulnerability	Microsoft Works File Converter Field Length Remote Code Execution Vulnerability	10.134.18.124	pancademo@pan...	10.134.18.124	42208	imap	Drop	Critical
02/02 10:17:08	vulnerability	Microsoft Works File Converter Field Length Remote Code Execution Vulnerability	10.134.18.124	pancademo@pan...	10.134.18.124	42208	imap	Drop	Critical
02/02 10:17:16	vulnerability	Microsoft Works File Converter Field Length Remote Code Execution Vulnerability	10.134.18.124	pancademo@pan...	10.134.18.124	42208	imap	Drop	Critical
02/02 10:17:24	vulnerability	Microsoft Works File Converter Field Length Remote Code Execution Vulnerability	10.134.18.124	pancademo@pan...	10.134.18.124	42208	imap	Drop	Critical
02/02 10:17:32	vulnerability	Microsoft Works File Converter Field Length Remote Code Execution Vulnerability	10.134.18.124	pancademo@pan...	10.134.18.124	42208	imap	Drop	Critical
02/02 10:17:40	vulnerability	Microsoft Works File Converter Field Length Remote Code Execution Vulnerability	10.134.18.124	pancademo@pan...	10.134.18.124	42208	imap	Drop	Critical
02/02 10:17:48	vulnerability	Microsoft Works File Converter Field Length Remote Code Execution Vulnerability	10.134.18.124	pancademo@pan...	10.134.18.124	42208	imap	Drop	Critical

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.

Use the App Scope Reports

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](#)


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.
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.

Button	Description
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.
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.
Outgoing 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.
Last 6 hours Last 12 hours Last 24 hours Last 7 days Last 30 days	Indicates the period over which the measurements are taken.

Network Monitor Report

The App Scope Network Monitor report (**Monitor > App Scope > Network Monitor**) displays the bandwidth dedicated to different network functions over the specified period of time. Each network function is color-coded as indicated in the legend below the chart. For example, the image below shows application bandwidth for the past 7 days based on session information.



The Network Monitor report contains the following buttons and options.

Button	Description
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.

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 following models support the automated correlation engine:

- Panorama—M-Series appliances and virtual appliances
- PA-7000 Series firewalls
- PA-5200 Series firewalls
- PA-3200 Series firewalls

- [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.

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 or activity seen by a [Traps protected endpoint on Panorama](#). 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

You can view the correlation objects that are currently available on the firewall.

STEP 1 | Select **Monitor > Automated Correlation Engine > Correlation Objects**. All the objects in the list are enabled by default.

A screenshot of the 'Correlation Objects' page in the Palo Alto Networks GUI. The page shows a table with columns for Name, ID, Category, State, and Description. The table lists several correlation objects, including 'Compromise Lifecycle', 'Malware C2', and 'High Activity'. Each row has a link to view the object's details.

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.

Match Time	Update Time	Object Name	Source Address	Source User	Severity	Summary
2018/01/01 15:48:48	2018/01/01 16:02:01	Malware Detection	10.47.26.20	panorama@panorama	Critical	Host visited known malware URL (1/8 times)
2018/01/01 15:53:07	2018/01/01 16:02:01	Malware Detection	10.47.26.20	panorama@panorama	Critical	Host visited known malware URL (2/8 times)
2018/01/01 16:17:20	2018/01/01 16:17:20	Malware Detection	10.47.26.20	panorama@panorama	Critical	Host visited known malware URL (3/8 times)
2018/01/01 16:49:49	2018/01/01 16:49:49	Malware Detection	10.47.26.20	panorama@panorama	Critical	Host visited known malware URL (4/8 times)
2018/01/01 17:10:47	2018/01/01 17:10:47	Malware Detection	10.47.26.20	panorama@panorama	Critical	Host visited known malware URL (5/8 times)

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.
 <i>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.</i>	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:



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 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 it's 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 Authentication Services](#), 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](#).
- **GTP Event Packet Capture**—The firewall captures a single GTP event, such as GTP-in-GTP, end user IP spoofing, and abnormal GTP messages, to make GTP troubleshooting easier for mobile network operators. Enable packet capture in a [GTP Protection profile](#).

Disable Hardware Offload

Packet captures for traffic passing through the network data ports on a Palo Alto Networks firewall are performed by the dataplane CPU. To capture traffic that passes through the management interface, you must [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, the packet capture filter is used differently by the ingress stage, compared to the firewall, drop, and egress capture stages. The ingress stage uses the packet capture filter to copy individual packets that match the filter to the capture file. Packets that fail packet-parsing checks are dropped before being captured. The firewall, drop, and egress capture stages use the same packet capture filter to mark all new sessions that match the filter. Because each session, as recorded in the session tables, identifies both client-to-server and server-to-client connections, any traffic, in either direction, that matches to the flagged session will be copied to the firewall-stage and transmit-stage capture files. Likewise, any dropped traffic (post receive stage) in either direction that matches to a flagged session will be copied to the drop-stage capture file.

On firewall models that include a network processor, traffic that meets certain pre-determined criteria by Palo Alto Networks may be offloaded for handling by the network processor. Such offloaded traffic will not reach the dataplane CPU and will, therefore, not be captured. To capture offloaded traffic, you must use the CLI to turn off the hardware offload feature.

Common types of traffic that may be offloaded include non-decrypted SSL and SSH traffic (which being encrypted cannot be usefully inspected beyond the initial SSL/SSH session setup), network protocols (such as OSPF, BGP, RIP), and traffic that matches an application-override policy. Some types of traffic will never be offloaded, such as ARP, all non-IP traffic, IPSec, and VPN sessions. Individual SYN, FIN, and RST packets, even for session traffic that has been offloaded, will never be offloaded, and will always be passed through to the dataplane CPU, once recognized as such by the network processor.



Hardware offload is supported on the following firewalls: PA-3200 Series, PA-5200 Series, and PA-7000 Series firewall.



Disabling hardware offload may increase the dataplane CPU usage. If dataplane CPU usage is already high, you may want to schedule a maintenance window before disabling 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](#).

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 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.

Detailed Log View		
General	Source	Destination
Session ID: 11840	User	User
Action: allow	Address: 192.168.2.10	Address: 10.43.14.55
Action Source: from-policy	Country: 192.168.0.0-192.168.255.255	Country: 10.0.0.0-10.255.255.255
Application: ping	Port: 0	Port: 0
Rule: rule1	Zone: 03-44n-trust	Zone: 03-trustnet
Session ID Reason: rule	Interface: eth0/2	Interface: ethernet0/1
Category: any	Nat IP: 10.43.14.25	Nat IP: 10.43.14.55
Virtual System: Default VSI	Nat Port: 0	Nat Port: 0

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.



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**.

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**.

The firewall or appliance warns you that system performance can be degraded; acknowledge the warning by clicking **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
```

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.

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	192.168.2.10	10.43.14.55	TCP	66	49525 → Telnet [SYN] Seq=0 win=8192 Len=0 MSS=3440 WS=256 SACK_PERM=3
2	3.002435	192.168.2.10	10.43.14.55	TCP	66	49525 → Telnet [SYN] Seq=0 win=8192 Len=0 MSS=3440 WS=256 SACK_PERM=3
3	8.008678	192.168.2.10	10.43.14.55	TCP	62	49525 → Telnet [SYN] Seq=0 win=8192 Len=0 MSS=3440 SACK_PERM=3

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.



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.

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	192.168.2.10	10.43.14.55	TCP	66	61214 → telnet [SYN] Seq=0 win=65536 Len=0 MSS=1460 WS=256 SACK_Pack=1
2	0.000000	10.43.14.55	10.43.14.25	TCP	66	telnet → 59293 [SYN, ACK] Seq=0 Ack=0 win=14600 Len=0 MSS=1460 SACK_Pack=1 ws=128
3	0.000114	192.168.2.10	10.43.14.55	TCP	66	61214 → telnet [ACK] Seq=1 Ack=1 win=65536 Len=0
4	0.000114	10.43.14.55	10.43.14.25	TELNET	69	Telnet data ...
5	0.000114	192.168.2.10	10.43.14.55	TELNET	60	Telnet data ...
6	0.000114	10.43.14.55	10.43.14.25	TCP	66	telnet → 59293 [ACK] Seq=16 Ack=6 win=14720 Len=0
7	0.000114	192.168.2.10	10.43.14.55	TELNET	66	Telnet data ...
8	0.000114	10.43.14.55	10.43.14.25	TELNET	66	Telnet data ...
9	0.000114	192.168.2.10	10.43.14.55	TELNET	66	Telnet data ...
10	0.000114	10.43.14.55	10.43.14.25	TELNET	66	Telnet data ...
11	0.000114	192.168.2.10	10.43.14.55	TELNET	66	Telnet data ...
12	0.000114	10.43.14.55	10.43.14.25	TELNET	66	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.

STEP 1 | Enable the packet capture option in the security profile.

Some security profiles allow you to define a single-packet capture or an extended-capture. If you choose extended-capture, define the capture length. This allows the firewall to capture more packets and provide additional context related to the threat.



If the action for a given threat is allow, the firewall does not trigger a Threat log and does not capture packets. If the action is alert, you can set the packet capture to single-packet or extended-capture. All blocking actions (drop, block, and reset actions) capture a single packet. The content package on the device determines the default action.

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**.



*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.
 1. Select **Device > Setup > Content-ID** and edit the Content-ID Settings.

2. 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).
3. 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.

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.

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-220>show running application setting | match "Unknown capture"
```

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

```
admin@PA-220>set application dump-unknown yes
```


STEP 2 | Locate unknown application by filtering the traffic logs.

1. Select **Monitor > Logs > Traffic**.
2. Click **Add Filter** and select the filters as shown in the following example.



3. Click **Add** and **Apply Filter**.

STEP 3 | Click the packet capture icon  to view the packet capture or **Export** it to your local system.



Take a Custom Application Packet Capture

You can configure a Palo Alto Networks firewall to take a packet capture based on an application name and filters that you define. You can then use the packet capture to troubleshoot issues with controlling an application. When configuring an application packet capture, you must use the application name defined in the App-ID database. You can view a list of all [App-ID](#) applications using [Applopedia](#) or from the web interface on the firewall in **Objects > Applications**.

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-220>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-220>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.

```
Application setting:
Application cache : yes
Supernode : yes
Heuristics : yes
Cache Threshold : 16
Bypass when exceeds queue limit: no
Traceroute appid : yes
Traceroute TTL threshold : 30
Use cache for appid : no
Unknown capture : on
Max. unknown sessions : 5000
Current unknown sessions : 0
Application capture : on
Max. application sessions : 5000
Current application sessions : 0
Application filter setting:
Rule : rule1
From : any
To : any
Source : any
Destination : any
Protocol : any
Source Port : any
Dest. Port : any
Application : facebook-base
Current APPID Signature
Signature Usage : 21 MB (Max. 32 MB)
TCP 1 C2S : 15503 states
TCP 1 S2C : 5070 states
TCP 2 C2S : 2426 states
TCP 2 S2C : 702 states
UDP 1 C2S : 11379 states
UDP 1 S2C : 2967 states
UDP 2 C2S : 755 states
UDP 2 S2C : 224 states
```


STEP 4 | 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-220>set application dump off
```

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-220 firewalls capture 68 bytes of data from each packet and anything over that is truncated. 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.

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-220>tcpdump filter "<filter-option> <IP-address>" snaplen length
```

For example, to capture the traffic that is generated when an 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-220>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-220>tcpdump filter "net 10.5.104.0/24 and not port 22" snaplen 0
```



Each time **tcpdump** takes a packet capture, it stores the content in a file named **mgmt.pcap**. This file is overwritten each time you run **tcpdump**.

STEP 3 | After the traffic you are interested in has traversed the MGT interface, press Ctrl + C to stop the capture.

STEP 4 | View the packet capture by running the following command:

```
admin@PA-220> 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-220>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-220>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.



[Content Delivery Network Infrastructure](#) 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.

View 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](#)
- [View Logs](#)
- [Filter Logs](#)
- [Export 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](#)
- [Tunnel Inspection Logs](#)
- [Config Logs](#)
- [System Logs](#)
- [HIP Match Logs](#)
- [IP-Tag Logs](#)
- [User-ID Logs](#)
- [Alarms Logs](#)
- [Authentication 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 specified any application, while a deny indicates 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.

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. WildFire Submissions log entries with a malicious verdict and an action set to allow are logged as High.
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. • Threat log entries with a malicious verdict and an action of block or alert, based on the existing WildFire signature severity, 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. • WildFire Submissions log entries with a grayware verdict and any action 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 are logged as Informational. • WildFire Submissions log entries with a benign verdict and any action are logged as Informational. • WildFire Submissions log entries with any verdict and an action set to block and forward are logged as Informational.

Severity	Description
	Log entries with any verdict and an action set to block are logged as Informational.

URL Filtering Logs

URL Filtering logs display entries for traffic that matches the URL Filtering profile attached to a security policy rule. 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.

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 firewall Action for the sample (allow or block), the WildFire verdict for the submitted sample, and the [severity level](#) of the sample.

The following table summarizes the WildFire verdicts:

Verdict	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).
Phishing	Indicates that WildFire assigned a link an analysis verdict of phishing. A phishing verdict indicates that the site to which the link directs users displayed credential phishing activity.
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.
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.

Tunnel Inspection Logs

Tunnel inspection logs are like traffic logs for tunnel sessions; they display entries of non-encrypted tunnel sessions. To prevent double counting, the firewall saves only the inner flows in traffic logs, and sends tunnel sessions to the tunnel inspection logs. The tunnel inspection log entries include Receive Time (date and time the log was received), the tunnel ID, monitor tag, session ID, the Security rule applied to the tunnel session, number of bytes in the session, parent session ID (session ID for the tunnel session), source address, source user and source zone, destination address, destination user, and destination zone.

Click the Detailed Log view to see details for an entry, such as the tunnel protocol used, and the flag indicating whether the tunnel content was inspected or not. Only a session that has a parent session will have the Tunnel Inspected flag set, which means the session is in a tunnel-in-tunnel (two levels of encapsulation). The first outer header of a tunnel will not have the Tunnel Inspected flag set.

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 display 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.

Severity	Description
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.

IP-Tag Logs

IP-tag logs display how and when a source IP address is registered or unregistered on the firewall and what tag the firewall applied to the address. Additionally, each log entry displays the configured timeout (when configured) and the source of the IP address-to-tag mapping information, such as User-ID agent VM information sources and auto-tagging. See how to [Register IP Address and Tags Dynamically](#) for more information.

User-ID Logs

[User-ID](#) logs display information about IP address-to-username mappings and [Authentication Timestamps](#), such as the sources of the mapping information and the times when users authenticated. You can use this information to help troubleshoot User-ID and authentication issues. For example, if the firewall is applying the wrong policy rule for a user, you can view the logs to verify whether that user is mapped to the correct IP address and whether the group associations are correct.

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.

Authentication Logs

Authentication logs display information about authentication events that occur when end users try to access network resources for which access is controlled by [Authentication Policy](#) rules. You can use this information to help troubleshoot access issues and to adjust your Authentication policy as needed. In conjunction with correlation objects, you can also use Authentication logs to identify suspicious activity on your network, such as brute force attacks.

Optionally, you can configure Authentication rules to log timeout events. These timeouts relate to the period when a user need authenticate for a resource only once but can access it repeatedly. Seeing information about the timeouts helps you decide if and how to adjust them (for details, see [Authentication Timestamps](#)).



System logs record authentication events relating to GlobalProtect and to administrator access to the web interface.

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 Role Types](#) 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.

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](#).

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 Role Types](#) 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](#).
- (Traffic, Threat, URL Filtering, WildFire Submissions, Data Filtering, and Unified logs only) View AutoFocus threat data for a log entry.
 1. [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).

1. Hover over an IP address, URL, user agent, threat name (subtype: virus and wildfire-virus only), filename, or SHA-256 hash.
3. Click the drop-down (▾) and select **AutoFocus**.
4. [Content Delivery Network Infrastructure](#).

Next Steps...

- [Filter Logs](#).
- [Export 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.

For example, filtering by the rule UUID makes it easier to pinpoint the specific rule you want to locate, even among many similarly-named rules. If your ruleset is very large and contains many rules, using the rule's UUID as a filter spotlights the particular rule you need to find without having to navigate through pages of results.



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.

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](#).

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.

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 1048576 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](#).

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.*

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).

STEP 1 | Select **Device** > **Scheduled Log Export** and click **Add**.

STEP 2 | Enter a **Name** for the scheduled log export and **Enable** it.

STEP 3 | Select the **Log Type** to export.

STEP 4 | Select the daily **Scheduled Export Start Time**. The options are in 15-minute increments for a 24-hour clock (00:00 - 23:59).

STEP 5 | Select the **Protocol** to export the logs: **SCP** (secure) or **FTP**.

STEP 6 | Enter the **Hostname** or IP address of the server.

STEP 7 | Enter the **Port** number. By default, FTP uses port 21 and SCP uses port 22.

STEP 8 | Enter the **Path** or directory in which to save the exported logs.

STEP 9 | Enter the **Username** and, if necessary, the **Password** (and **Confirm Password**) to access the server.

STEP 10 | (**FTP only**) Select **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**.

Monitor Block List

There are two ways you can cause the firewall to place an IP address on the block list:

- Configure a Vulnerability Protection profile with a rule to Block IP connections and apply the profile to a Security policy, which you apply to a zone.
- Configure a DoS Protection policy rule with the Protect action and a Classified DoS Protection profile, which specifies a maximum rate of connections per second allowed. When incoming packets match the DoS Protection policy and exceed the Max Rate, and if you specified a Block Duration and a Classified policy rule to include source IP address, the firewall puts the offending source IP address on the block list.

In the cases described above, the firewall automatically blocks that traffic in hardware before those packets use CPU or packet buffer resources. If attack traffic exceeds the blocking capacity of the hardware, the firewall uses IP blocking mechanisms in software to block the traffic.

The firewall automatically creates a hardware block list entry based on your Vulnerability Protection profile or DoS Protection policy rule; the source address from the rule is the source IP address in the hardware block list.

Entries on the block list indicate in the Type column whether they were blocked by hardware (hw) or software (sw). The bottom of the screen displays:

- Count of **Total Blocked IPs** out of the number of blocked IP addresses the firewall supports.
- Percentage of the block list that the firewall has used.

To view details about an address on the block list, hover over a Source IP address and click the down arrow link. Click the Who Is link, which displays the [Network Solutions Who Is](#) feature, providing information about the address.

For information on configuring a Vulnerability Protection profile, see [Customize the Action and Trigger Conditions for a Brute Force Signature](#). For more information on block list and DoS Protection profiles, see [DoS Protection Against Flooding of New Sessions](#).

View and Manage Reports

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 Expiration Period and Run Time for Reports](#)
- [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](#)
- [Manage Report Storage Capacity](#)

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. This limit can be reconfigured for PA-7000 series and PA-5200 series firewalls only. For all other firewall models, you can [Configure the](#)

[Expiration Period and Run Time for Reports](#) to allow the firewall to 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.

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 Expiration Period and Run Time for Reports

The expiration period and run time are global settings that apply to all [Report Types](#). After running new reports, the firewall automatically deletes reports that exceed the expiration period.

STEP 1 | Select **Device > Setup > Management**, edit the Logging and Reporting Settings, and select the **Log Export and Reporting** tab.

STEP 2 | Set the **Report Runtime** to an hour in the 24-hour clock schedule (default is 02:00; range is 00:00 [midnight] to 23:00).

STEP 3 | Enter the **Report Expiration Period** in days (default is no expiration; range is 1 to 2,000).



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 4 | Click **OK** and **Commit**.

Disable Predefined Reports

The firewall includes about 40 predefined reports that it automatically generates daily. If you do not use some or all of these, you can disable selected reports to conserve system resources on the firewall.

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



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**.

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, such as threats, as well as the best way to categorize the information, such as grouping by rule UUID, which will allow you to see the rule that applies to each threat type. This consideration guides you in making the following selections in a custom report:

Selection	Description
Database	You can base the report on one of the following database types: • Summary databases—These databases are available for Application Statistics, Traffic, Threat, URL Filtering, and Tunnel Inspection logs. The firewall aggregates the detailed logs at 15-minute intervals. To enable faster response time when generating reports, the firewall condenses the data: duplicate sessions are grouped and incremented with a repeat counter, and some attributes (columns) are excluded from the summary. • Detailed logs—These databases itemize the logs and list all the attributes (columns) for each log entry.  <i>Reports based on detailed logs take much longer to run and are not recommended unless absolutely necessary.</i>
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.
	The following example illustrates how the Selected Columns and Sort By/Group By criteria work together when generating reports:

Selection	Description
	 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.
	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 Frame	The date range for which you want to analyze data. You can define a custom range or select a time period ranging from the last 15 minutes to the last 30

Selection	Description
	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

You can configure custom reports that the firewall generates immediately (on demand) or on schedule (each night). To understand the selections available to create a purposeful custom report, see [Custom Reports](#).



After the firewall has generated a scheduled custom report, you risk invalidating the past results of that report if you modify its configuration to change its future output. If you need to modify a scheduled report configuration, the best practice is to create a new report.

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.

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	878.13 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	gnunella	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

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 blocking, 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.

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 allows you to easily explore SaaS application activity by risk and sanction state. 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 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 presents the key findings for the SaaS applications on your network during the reporting period with a comparison of the sanctioned versus unsanctioned applications and lists the top applications based on sanction state by usage, compliance, and data transfers. To help you identify and explore the extent of high risk application usage, the applications with risky characteristics section of the report lists the SaaS applications with the following unfavorable hosting characteristics: certifications achieved, past data breaches, support for IP-based restrictions, financial viability, and terms of service. You can also view a comparison of sanctioned versus unsanctioned SaaS applications by total number of applications used on your network, bandwidth consumed by these applications, the number of users using these applications, top user groups that use the largest number of SaaS applications, and the top user groups that transfer the largest volume of data through sanctioned and unsanctioned SaaS 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 and risky applications that pose unnecessary risks for malware propagation and data leaks.



The predefined SaaS application usage report is still available as a daily [View Reports](#) that lists the top 100 SaaS applications (which means applications with the SaaS application characteristic, SaaS=yes) running on your network on a given day. This report does not give visibility into applications you have designated as sanctioned, but rather gives visibility into all of the SaaS applications in use on your network.

STEP 1 | Tag applications that you approve for use on your network as Sanctioned.




For generating an accurate and informative report, you need to tag the sanctioned applications consistently across firewalls with multiple virtual systems, and across firewalls that belong to a device group on Panorama. If the same application is tagged as sanctioned in one virtual system and is not sanctioned in another or, on Panorama, if an application is unsanctioned in a parent device group but is tagged as sanctioned in a child device group (or vice versa), the SaaS Application Usage report will report the application as partially sanctioned and will have overlapping results.

Example: If Box is sanctioned on vsys1 and Google Drive is sanctioned on vsys2, Google Drive users in vsys1 will be counted as users of an unsanctioned SaaS application and Box users in vsys2 will be counted as users of an unsanctioned SaaS application. The key finding in the report will highlight that a total of two unique SaaS applications are discovered on the network with two sanctioned applications and two unsanctioned applications.

1. Select **Objects > 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 (). 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 10 pages.

3. Select whether you want the report to **Include logs from**:

- **All User Groups and Zones**—The report includes data on all security zones and user groups available in the logs.

If you want to include specific user groups in the report, select **Include user group information in the report** and click the **manage groups** link to select the groups you want to include. You must add between one and up to a maximum of 25 user groups, so that the firewall or Panorama can filter the logs for the selected user groups. If you do select the groups to include, the report will aggregate all user groups in to one group called Others.

- **Selected Zone**—The report filters data for the specified security zone, and includes data on that zone only.

If you want to include specific user groups in the report, select **Include user group information in the report** and click the **manage groups for selected zone** link to select the user groups within this zone that you want to include in the report. You must add between one and up to a maximum of 25 user groups, so that the firewall or Panorama can filter the logs for the selected user groups within the security zone. If you do select the groups to include, the report will aggregate all user groups in to one group called Others.

- **Selected User Group**—The report filters data for the specified user group only, and includes SaaS application usage information for the selected user group only.



4. Select whether you want to include all the application subcategories in the report (the default) or **Limit the max subcategories in the report** to the top 10, 15, 20 or 25 categories (default is all subcategories).
5. Click **Run Now** to generate the report on-demand for the last 7-day and the last 30-day time period. Make sure that the pop-up blocker is disabled on your browser because the report opens in a new tab.
6. Click **OK** to save your changes.

STEP 3 | Schedule Reports for Email Delivery.

The last 90-days report must be scheduled for email delivery.


On the PA-220R and the PA-800 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.

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.




Selecting Top Threats is displayed as top-attacks in the Predefined Widgets column for the PDF Summary Report.

- 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** to save the report.
 5. **Commit** the changes.

STEP 2 | View the report.

To download and view the PDF Summary Report, see [View Reports](#).



The following summary sections refer to the following PDF Summary Report elements:

- **Top 5 Attacks**—Refers to the Top threats element.
- **Top 5 Threats**—Refers to the High risk user - Top threats element.
- **Top Threats report**—Refers to the full list of threats from the Top threats element.

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.

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 **Log Only the Page a User Visits** (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. (Optional) 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](#).


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.

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.
 1. Select **Monitor > Report Group**.
 2. Click **Add** and then enter a **Name** for the report group.
 3. (Optional) Select **Title Page** and add a **Title** for the PDF output.
 4. 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.

5. Click **OK** to save the settings.
6. 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.

- 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**.

Manage Report Storage Capacity

By default, firewalls contain 200MB of dedicated storage for [reports](#) generated by the firewall. In some instances, especially for PA-7000 series and PA-5200 series firewalls, you may need to increase the capacity of available report storage space in order to successfully generate new reports.

STEP 1 | [Access the firewall CLI](#).

STEP 2 | Confirm the current report storage capacity of the firewall:

The command output displays the report storage size in bytes. For this procedure, the firewall has the default 200MB report storage capacity.

```
admin@ISP-CONDOR-B(active)> request report-storage-size show
209715200
```

STEP 3 | Verify you have sufficient storage across the firewall to allocate toward expanding the report storage capacity:

```
admin> show system disk-space
```

```
admin@ISP-CONDOR-B(active)> show system disk-space

Filesystem      Size  Used Avail  Used Mounted on
/dev/root        12G  8.9G  2.6G   83% /
none             7.9G   52K  7.9G    1% /dev
/dev/sda5        16G  8.5G  5.9G   59% /opt/pancfg
/dev/sda6        12G  5.8G  5.0G   54% /opt/panrepo
tmpfs            7.9G  247M  7.6G    4% /dev/shm
/dev/sda8        22G  8.7G  12G   43% /opt/panlogs
tmpfs            12M    0   12M    0% /opt/pancfg/mgmt/locas/sal/private
```

STEP 4 | Increase the report storage capacity as needed:

For example, we are increasing the report storage size to 1GB.

```
admin> request report-storage-size set size <0-4>
```

```
admin@ISP-CONDOR-B(active)> request report-storage-size set size 1
cfg.report-storage-size-gb: 1
```

STEP 5 | Verify that the report storage capacity is increased to the amount set in the previous step:

```
admin> request report-storage-size show
```

```
admin@ISP-CONDOR-B(active)> request report-storage-size show
1073741024
```

View Policy Rule Usage

View the number of times a Security, NAT, QoS, policy-based forwarding (PBF), Decryption, Tunnel Inspection, Application Override, Authentication, or DoS protection rule matches traffic to help keep your firewall policies up to date as your environment and security needs change. To prevent attackers from exploiting over-provisioned access, such as when a server is decommissioned or when you no longer need temporary access to a service, use the policy rule hit count data to identify and remove unused rules.

Policy rule usage data enables you 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 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 whether traffic is matching the app-based rule instead of the port-based rule. The policy rule hit count helps you determine whether a rule is effective for access enforcement.

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 that data is no longer available after you reset (clear) the hit count.



The rule hit count data is not synchronized across firewalls in a high availability (HA) deployment so you need to log in to each firewall to view the policy rule hit count data for each firewall or use Panorama to view information on the HA firewall peers.



Policy rule usage data is also useful when using [Security Policy Rule Optimization](#) to determine which rules to migrate or clean up first.

STEP 1 | [Launch the Web Interface.](#)

STEP 2 | Verify that **Policy Rule Hit Count** is enabled.

1. Navigate to Policy Rulebase Settings (**Device > Setup > Management**).
2. Verify that **Policy Rule Hit Count** is enabled.



STEP 3 | Select **Policies**.

STEP 4 | View the policy rule usage for each policy rule:

- **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.
- **Modified**—The date and time the policy rule was last modified.
- **Created**—The date and time the policy rule was created.



If the rule was created when Panorama was running PAN-OS 8.1 and the Policy Rule Hit Count setting is enabled, the First Hit date and time is used as the Created date and time on upgrade to PAN-OS 9.0. If the rule was created in PAN-OS 8.1 when the Policy Rule Hit Count setting was disabled or if the rule was created when Panorama was running PAN-OS 8.0 or an earlier release, the Created date for the rule will be the date and time you successfully upgraded Panorama to PAN-OS 9.0

Name	Tags	Type	Zone	Rule Usage				
				Hit Count	Last Hit	First Hit	Modified	Created
7. panorama server	none	universal	any	227173	2018-09-07 15:09:42	2018-05-23 11:14:03	2018-09-25 13:04:30	2018-05-22 19:07:50
8. Allow-Rail-ETD-AS2	none	universal	any	4	2018-05-24 05:28:27	2018-05-15 05:04:45	2018-09-25 13:04:30	2018-05-15 05:05:29
9. Allow-Scapy-viden-A...	none	universal	any	34201	2018-08-14 03:33:48	2018-05-18 13:32:20	2018-09-25 13:04:30	2018-05-18 13:29:13
10. Allow-Scapy-viden	none	universal	any	10421	2018-05-18 13:28:28	2018-05-03 07:34:33	2018-09-25 13:04:30	2018-05-01 00:07:28
11. Allow-Scapy-viden-1	none	universal	any	2223	2018-05-18 13:32:20	2018-05-05 06:17:31	2018-09-25 13:04:30	2018-05-01 00:07:28
12. Allow-Rail-Apps	none	universal	any	13	2018-09-07 09:29:35	2018-08-22 15:01:52	2018-09-25 13:04:30	2018-08-15 15:30:21
13. Allow-Rail	none	universal	any	787	2018-08-22 19:01:57	2018-04-26 05:05:44	2018-09-25 13:04:30	2018-04-11 13:27:30
14. Deny-Driver-Flow-Social-net		universal	any	0	-	-	2018-09-25 13:04:30	2018-08-22 14:42:55
15. Allow-Ping	ping	universal	any	7	2018-04-26 05:03:38	2018-04-12 04:34:37	2018-09-25 13:04:30	2018-04-07 14:23:49

STEP 5 | In the Policy Optimizer dialog, view the **Rule Usage** filter.

STEP 6 | Filter rules in the selected rulebase.



Use the rule usage filter to evaluate the rule usage within a specified period of time. For example, filter the selected rulebase for Unused rules within the last 30 days. You can also evaluate rule usage with other rule attributes, such as the Created and Modified dates, which enables you to filter for the correct set of rules to review. You can use this data to help manage your rule lifecycle and to determine if a rule needs to be removed to reduce your network attack surface.

1. Select the **Timeframe** you want to filter on or specify a **Custom** time frame.
2. Select the rule **Usage** on which to filter.
3. (Optional) If you have reset the rule usage data for any rules, check for **Exclude rules reset during the last <number of days> days** and decide when to exclude a rule based on the number of days you specify since the rule was reset. Only rules that were reset before your specified number of days are included in the filtered results.

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](#).
- ❑ To send an HTTP-based API request directly to any third-party service that exposes an API to automate a workflow or an action. You can, for example, forward logs that match a defined criteria to create an incidence ticket on ServiceNow instead of relying on an external system to convert syslog messages or SNMP traps to an HTTP request. You can modify the URL, HTTP header, parameters, and the payload in the HTTP request to trigger an action based on the attributes in a firewall log. See [Forward Logs to an HTTP\(S\) Destination](#).
- ❑ 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

In an environment where you use multiple firewalls to control and analyze network traffic, any single firewall can display logs and reports only for the traffic it monitors. Because logging in to multiple firewalls can make monitoring a cumbersome task, you can more efficiently achieve global visibility into network activity by forwarding the logs from all firewalls to Panorama or external services. If you [Use External Services for Monitoring](#), the firewall automatically converts the logs to the necessary format: syslog messages, SNMP traps, email notifications, or as an HTTP payload to send the log details to an HTTP(S) server. In cases where some teams in your organization can achieve greater efficiency by monitoring only the logs that are relevant to their operations, you can create forwarding filters based on any log attributes (such as threat type or source user). For example, a security operations analyst who investigates malware attacks might be interested only in Threat logs with the type attribute set to wildfire-virus.



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 [View and Manage Reports](#), but only on a per log type basis, not for the entire log database.

STEP 1 | Configure a server profile for each external service that will receive log information.



You can use separate profiles to send different sets of logs, filtered by log attributes, to a different server. To increase availability, define multiple servers in a single profile.

Configure one or more of the following server profiles:

- [Create an Email server profile](#).
- Step [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.
- If the syslog server requires client authentication, you must also [Create a certificate to secure syslog communication over TLSv1.2](#).
- Configure an HTTP server profile (see [Forward Logs to an HTTP\(S\) Destination](#)).

STEP 2 | Create a Log Forwarding profile.

The profile defines the destinations for Traffic, Threat, WildFire Submission, URL Filtering, Data Filtering, Tunnel and Authentication logs.

1. Select **Objects > Log Forwarding** and **Add** a profile.
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. **Add** one or more *match list* profiles.

The profiles specify log query filters, forwarding destinations, and automatic actions such as tagging. For each match list profile:

1. Enter a **Name** to identify the profile.
 2. Select the **Log Type**.
 3. In the **Filter** drop-down, select **Filter Builder**. Specify the following and then **Add** each query:
 - **Connector** logic (and/or)
 - Log **Attribute**
 - **Operator** to define inclusion or exclusion logic
 - Attribute **Value** for the query to match
 4. Select Panorama if you want to forward logs to Log Collectors or the Panorama management server.
 5. For each type of external service that you use for monitoring (SNMP, Email, Syslog, and HTTP), **Add** one or more server profiles.
4. Click **OK** to save the Log Forwarding profile.

STEP 3 | Assign the Log Forwarding profile to policy rules and network zones.

Security, Authentication, and DoS Protection rules support log forwarding. In this example, you assign the profile to a Security rule.

Perform the following steps for each rule that you want to trigger log forwarding:

1. Select **Policies > Security** and edit the rule.
2. Select **Actions** and select the **Log Forwarding** profile you created.
3. Set the **Profile Type** to **Profiles** or **Group**, and then select the **security profiles** or **Group Profile** required to trigger log generation and forwarding for:
 - Threat logs—Traffic must match any security profile assigned to the rule.
 - WildFire Submission logs—Traffic must match a **WildFire Analysis profile** assigned to the rule.
4. For Traffic logs, select **Log At Session Start** and/or **Log At Session End**.
5. Click **OK** to save the rule.

STEP 4 | Configure the destinations for System, Configuration, User-ID, HIP Match, and Correlation logs.



Panorama generates Correlation logs based on the firewall logs it receives, rather than aggregating Correlation logs from firewalls.

1. Select **Device > Log Settings**.
2. For each log type that the firewall will forward, see Step [Add one or more match list profiles](#).

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. Set the **Interface Type** to **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.

STEP 6 | Commit and verify your changes.

1. **Commit** your changes.
2. Verify the log destinations you configured are receiving firewall logs:
 - Panorama—If the firewall forwards logs to a Panorama virtual appliance in Panorama mode or 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 your syslog server documentation to verify it's receiving logs as syslog messages.
 - SNMP manager—[Use an SNMP Manager to Explore MIBs and Objects](#) to verify it's receiving logs as SNMP traps.
 - HTTP server—[Forward Logs to an HTTP\(S\) Destination](#).

Configure Email Alerts

You can configure email alerts for System, Config, HIP Match, Correlation, Threat, WildFire Submission, and Traffic logs.

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. See Step [Create a Log Forwarding profile](#).
 1. Select **Objects > Log Forwarding**, click **Add**, and enter a **Name** to identify the profile.
 2. For each log type and each severity level or WildFire verdict, select the Email server profile and click **OK**.
2. See Step [Assign the Log Forwarding profile to policy rules and network zones](#).

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 TLS (TLSv1.2 only) 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 TLSv1.2.

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.



If you configure an FQDN and use UDP transport, if the firewall cannot resolve the FQDN, the firewall uses the existing IP address resolution for the FQDN as the Syslog Server address.

- **Transport**—Select **TCP**, **UDP**, or **SSL** (TLS) as the protocol for communicating with the syslog server. For **SSL**, the firewall supports only TLSv1.2.
 - **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/TLS.
 - **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.

STEP 2 | Configure syslog forwarding for Traffic, Threat, and WildFire Submission logs.

1. Configure the firewall to forward logs. For more information, see Step [Create a Log Forwarding profile](#).

1. Select **Objects > Log Forwarding**, click **Add**, and enter a **Name** to identify the profile.
2. 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 a security policy to trigger log generation and forwarding. For more information, See Step [Assign the Log Forwarding profile to policy rules and network zones](#).
 1. Select **Policies > Security** and select a policy rule.
 2. Select the **Actions** tab and select the **Log Forwarding** profile you created.
 3. In the **Profile Type** drop-down, select **Profiles** or **Groups**, and then select the security profiles or **Group Profiles** 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**.

For detailed information about configuring a log forwarding profile and assigning the profile to a policy rule, see [Configure Log Forwarding](#).

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.

STEP 5 | Create a certificate to secure syslog communication over TLSv1.2.

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.

- The connection to a Syslog server over TLS is validated using the Online Certificate Status Protocol (OCSP) or using Certificate Revocation Lists (CRL) so long as each certificate in the trust chain specifies one or both of these extensions. However, you cannot bypass OCSP or CRL failures so you must ensure that the certificate chain is valid and that you can verify each certificate using OCSP or CRL.

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](#)
- [IP-Tag Log Fields](#)
- [User-ID Log Fields](#)
- [Tunnel Inspection Log Fields](#)
- [SCTP Log Fields](#)
- [Config Log Fields](#)
- [Authentication Log Fields](#)
- [System Log Fields](#)
- [Correlated Events Log Fields](#)
- [GTP Log Fields](#)
- [Custom Log/Event Format](#)
- [Escape Sequences](#)

Traffic Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Source Address, Destination Address, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, 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, Source VM UUID, Destination VM UUID, Tunnel ID/IMSI, Monitor Tag/IMEI, Parent Session ID, Parent Start Time, Tunnel Type, SCTP Association ID, SCTP Chunks, SCTP Chunks Sent, SCTP Chunks Received, UUID for rule, HTTP/2 Connection

Field Name	Description
Receive Time (receive_time or cef-formatted-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 the type of log; value is TRAFFIC.
Threat/Content Type (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 or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Source Address (src)	Original session source IP address.
Destination Address (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.
Inbound Interface (inbound_if)	Interface that the session was sourced from.

Field Name	Description
Outbound Interface (outbound_if)	Interface that the session was destined to.
Log Action (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.
Source Port (sport)	Source port utilized by the session.
Destination Port (dport)	Destination port utilized by the session.
NAT Source Port (nat sport)	Post-NAT source port.
NAT Destination Port (nat dport)	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) • 0x40000000—option is enabled to allow a client to use multiple paths to connect to a destination host • 0x20000000—file is submitted to WildFire for a verdict • 0x10000000—enterprise credential submission by end user detected • 0x08000000— source for the flow is an allow list and not subject to recon protection • 0x02000000—IPv6 session • 0x01000000—SSL session is decrypted (SSL Proxy) • 0x00800000—session is denied via URL filtering • 0x00400000—session has a NAT translation performed • 0x00200000—user information for the session was captured through Captive Portal • 0x00100000—application traffic is on a non-standard destination port • 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) • 0x00020000—Client to Server flow is subject to policy based forwarding • 0x00010000—Server to Client flow is subject to policy based forwarding • 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.

Field Name	Description
	0x00000800—symmetric return is used to forward traffic for this session 0x00000400—decrypted traffic is being sent out clear text through a mirror port 0x00000100—payload of the outer tunnel is being inspected
IP 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.
Bytes Received (bytes_received)	Number of bytes in the server-to-client direction of the session.
Packets (packets)	Number of total packets (transmit and receive) for the session.
Start Time (start)	Time of session start.
Elapsed Time (sec)	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.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Source Country (srcloc)	Source country or Internal region for private addresses; maximum length is 32 bytes.
Destination Country (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.
Packets Received (pkts_received)	Number of server-to-client packets for the session.

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—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 <code>clear session all</code> command).

Field Name	Description
	- 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.
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 readonlydg-meta-data API query: <pre>/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show></pre>
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.
Source VM UUID (src_uuid)	Identifies the source universal unique identifier for a guest virtual machine in the VMware NSX environment.
Destination VM UUID (dst_uuid)	Identifies the destination universal unique identifier for a guest virtual machine in the VMware NSX environment.
Tunnel ID/IMSI (tunnelid/imsi)	International Mobile Subscriber Identity (IMSI) is a unique number allocated to each mobile subscriber in the GSM/UMTS/EPS system. IMSI shall consist of decimal digits (0 through 9) only and maximum number of digits allowed are 15.
Monitor Tag/IMEI (monitortag/imei)	International Mobile Equipment Identity (IMEI) is a unique 15 or 16 digit number allocated to each mobile station equipment.

Field Name	Description
Parent Session ID (parent_session_id)	ID of the session in which this session is tunneled. Applies to inner tunnel (if two levels of tunneling) or inside content (if one level of tunneling) only.
Parent Start Time (parent_start_time)	Year/month/day hours:minutes:seconds that the parent tunnel session began.
Tunnel Type (tunnel)	Type of tunnel, such as GRE or IPSec.
SCTP Association ID (assoc_id)	Number that identifies all connections for an association between two SCTP endpoints.
SCTP Chunks (chunks)	Sum of SCTP chunks sent and received for an association.
SCTP Chunks Sent (chunks_sent)	Number of SCTP chunks sent for an association.
SCTP Chunks Received (chunks_received)	Number of SCTP chunks received for an association.
UUID for rule (rule_uuid)	The UUID that permanently identifies the rule.
HTTP/2 Connection	Identifies if traffic used an HTTP/2 Connection by displaying one of the following values: • Parent session ID—HTTP/2 connection • 0—SSL session

Threat Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Source Address, Destination Address, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, URL/Filename, Threat ID, Category, Severity, Direction, Sequence Number, Action Flags, Source Location, Destination Location, FUTURE_USE, Content Type, PCAP_ID, File Digest, 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, Source VM UUID, Destination VM UUID, HTTP Method, Tunnel ID/IMSI, Monitor Tag/IMEI, Parent Session ID, Parent Start Time, Tunnel Type, Threat Category, Content Version, FUTURE_USE, SCTP Association ID, Payload Protocol ID, HTTP Headers, URL Category List, UUID for rule, HTTP/2 Connection

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial #)	Serial number of the firewall that generated the log.

Field Name	Description
Type (type)	Specifies the type of log; value is THREAT.
Threat/Content Type (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, phishing, grayware, or benign, depending on what you are logging) is logged in the WildFire Submissions log. • wildfire-virus—Virus detected via an Antivirus profile.
Generate Time (time_generated or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Source address (src)	Original session source IP address.
Destination address (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.
Inbound Interface (inbound_if)	Interface that the session was sourced from.

Field Name	Description
Outbound Interface (outbound_if)	Interface that the session was destined to.
Log Action (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 Content/Threat Type seen within 5 seconds.
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) • 0x40000000—option is enabled to allow a client to use multiple paths to connect to a destination host • 0x20000000—file is submitted to WildFire for a verdict • 0x10000000—enterprise credential submission by end user detected • 0x08000000— source for the flow is an allow list and not subject to recon protection • 0x02000000—IPv6 session • 0x01000000—SSL session is decrypted (SSL Proxy) • 0x00800000—session is denied via URL filtering • 0x00400000—session has a NAT translation performed • 0x00200000—user information for the session was captured through Captive Portal • 0x00100000—application traffic is on a non-standard destination port • 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) • 0x00020000—Client to Server flow is subject to policy based forwarding • 0x00010000—Server to Client flow is subject to policy based forwarding • 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 is used to forward traffic for this session

Field Name	Description
	0x00000400—decrypted traffic is being sent out clear text through a mirror port 0x00000010—payload of the outer tunnel is being inspected
IP 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 - 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 - block-ip—threat detected and client IP is blocked - random-drop—flood detected and packet was randomly dropped - sinkhole—DNS sinkhole activated - syncookie-sent—syncookie alert - block-continue (URL subtype only)—a HTTP request is blocked and redirected to a Continue page with a button for confirmation to proceed - continue (URL subtype only)—response to a block-continue URL continue page indicating a block-continue request was allowed to proceed - block-override (URL subtype only)—a HTTP request is blocked and redirected to an Admin override page that requires a pass code from the firewall administrator to continue - override-lockout (URL subtype only)—too many failed admin override pass code attempts from the source IP. IP is now blocked from the block-override redirect page - override (URL subtype only)—response to a block-override page where a correct pass code is provided and the request is allowed - block (Wildfire only)—file was blocked by the firewall and uploaded to Wildfire
URL/Filename (misc)	Field with variable length. A Filename has a maximum of 63 characters. A URL has 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-virus File name when the subtype is wildfire

Field Name	Description
	URL or File name when the subtype is vulnerability if applicable
Threat/Content Name (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 —spyware phone home detection • 20000 – 29999 —spyware download detection • 30000 – 44999 —vulnerability exploit detection • 52000 – 52999— filetype detection • 60000 – 69999 —data filtering detection  <i>Threat ID ranges for virus detection, WildFire signature feed, and DNS C2 signatures used in previous releases have been replaced with permanent, globally unique IDs. Refer to the Threat/Content Type (subtype) and Threat Category (thr_category) field names to create updated reports, filter threat logs, and ACC activity.</i>
Category (category)	For URL Subtype, it is the URL Category; For WildFire subtype, it is the verdict on the file and is either 'malicious', 'phishing', '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.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Source Country (srcloc)	Source country or Internal region for private addresses. Maximum length is 32 bytes.
Destination Country (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

Field Name	Description
	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.
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 and virus subtypes; 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.

Field Name	Description
Subject (subject)	Only for WildFire and virus subtypes; 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 and virus subtypes; 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 and virus subtypes; all other types do not use this field. Identifies the analysis request on the WildFire cloud or the WildFire appliance.
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: API query: <pre>/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show></pre>
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.
Source VM UUID (src_uuid)	Identifies the source universal unique identifier for a guest virtual machine in the VMware NSX environment.
Destination VM UUID (dst_uuid)	Identifies the destination universal unique identifier for a guest virtual machine in the VMware NSX environment.
HTTP Method (http_method)	Only in URL filtering logs. Describes the HTTP Method used in the web request. Only the following methods are logged: Connect, Delete, Get, Head, Options, Post, Put.
Tunnel ID/IMSI (tunnel_id/imsi)	International Mobile Subscriber Identity (IMSI) is a unique number allocated to each mobile subscriber in the GSM/UMTS/EPS system. IMSI shall consist of decimal digits (0 through 9) only and maximum number of digits allowed are 15.
Monitor Tag/IMEI (monitortag/imei)	International Mobile Equipment Identity (IMEI) is a unique 15 or 16 digit number allocated to each mobile station equipment.

Field Name	Description
Parent Session ID (parent_session_id)	ID of the session in which this session is tunneled. Applies to inner tunnel (if two levels of tunneling) or inside content (if one level of tunneling) only.
Parent Session Start Time(parent_start_time)	Year/month/day hours:minutes:seconds that the parent tunnel session began.
Tunnel Type (tunnel)	Type of tunnel, such as GRE or IPSec.
Threat Category (thr_category)	Describes threat categories used to classify different types of threat signatures.
Content Version (contentver)	Applications and Threats version on your firewall when the log was generated.
SCTP Association ID (assoc_id)	Number that identifies all connections for an association between two SCTP endpoints.
Payload Protocol ID (ppid)	ID of the protocol for the payload in the data portion of the data chunk.
HTTP Headers (http_headers)	Indicates the inserted HTTP header in the URL log entries on the firewall.
URL Category List	Lists the URL Filtering categories that the firewall used to enforce policy.
UUID for rule (rule_uuid)	The UUID that permanently identifies the rule.
HTTP/2 Connection	Identifies if traffic used an HTTP/2 connection by displaying one of the following values: • TCP connection session ID—session is HTTP/2 • 0—session is not HTTP/2

HIP Match Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, 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, Virtual System ID, IPv6 Source Address, Host ID, User Device Serial Number

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial)	Serial number of the firewall that generated the log.

Field Name	Description
Type (type)	Specifies the type of log; value is HIP-MATCH.
Threat/Content Type (subtype)	Subtype of HIP match log; unused.
Generated Time (time_generated or cef-formatted-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 (os)	The operating system installed on the user's machine or device (or on the client system).
Source Address (src)	IP address of the source user.
HIP (matchname)	Name of the HIP object or profile.
Repeat Count (repeatcnt)	Number of times the HIP profile matched.
HIP Type (matchtype)	Whether the hip field represents a HIP object or a HIP profile.
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space.
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 readonlydg-meta-data API query: <pre>/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show></pre>

Field Name	Description
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.
Virtual System ID (vsys_id)	A unique identifier for a virtual system on a Palo Alto Networks firewall.
IPv6 System Address (srcipv6)	IPv6 address of the user's machine or device.
Host ID (hostid)	Unique ID GlobalProtect assigns to identify the host.
User Device Serial Number (serialnumber)	Serial number of the user's machine or device.

IP-Tag Log Fields

Format: FUTURE_USE , Receive Time, Serial, Type, Threat/Content Type, FUTURE_USE, Generate Time, Virtual System, Source IP, Tag Name , Event ID, Repeat Count , Timeout, Data Source Name, Data Source Type, Data Source Subtype, Sequence Number, Action Flags, DG Hierarchy Level 1 , DG Hierarchy Level 2, DG Hierarchy Level 3, DG Hierarchy Level 4, Virtual System Name, Device Name, Virtual System ID

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	The time the log was received at the management plane.
Serial Number (serial)	The serial number of the firewall that generated the log.
Type (type)	Specifies the type of log; value is IPTAG.
Threat/Content Type (subtype)	The subtype of the HIP match log; unused.
Generated Time (time_generated or cef-formatted-time_generated)	The time the log was generated on the dataplane.
Virtual System (vsys)	The virtual system associated with the HIP match log.
Source IP (src)	The IP address of the source user.
Tag Name (tag_name)	The tag mapped to the source IP address.
Event ID (event_id)	A string showing the name of the event.

Field Name	Description
Repeat Count (repeatcnt)	The number of sessions with the same Source IP, Destination IP, Application, and Subtype seen within 5 seconds.
Timeout (timeout)	The amount of time before the IP address-to-tag mapping expires for the source IP address.
Data Source Name (datasourcename)	The name of the source from which mapping information is collected.
Data Source Type (datasource_type)	The source from which mapping information is collected.
Data Source Subtype (datasource_subtype)	The mechanism used to identify the IP address-to-username mappings within a data source.
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially. Each log type has a unique number space.
Action Flags (actionflags)	A bit field indicating whether the log was forwarded to Panorama.
Device Group Hierarchy (dg_hier_level_1 to dg_hier_level_4)	A sequence of identification numbers that indicates the location of the device group 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 except the shared device group (level 0), which is not included in this structure. If the log values are 12, 34, 45, and 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 readonlydg-meta-data - API query: <pre>/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show></pre>
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.
Virtual System ID (vsys_id)	A unique identifier for a virtual system on a Palo Alto Networks firewall.

User-ID Log Fields

Format: FUTURE_USER, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Virtual System, Source IP, User, Data Source Name, Event ID, Repeat Count, Time Out Threshold, Source Port, Destination Port, Data Source, Data Source Type, 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, Virtual System ID, Factor Type, Factor Completion Time, Factor Number, FUTURE_USE, FUTURE_USE, User Group Flags, Source by User

Field Name	Description
Receive Time (receive_time or cef-formatted-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 the type of log; value is USERID.
Threat/Content Type (subtype)	Subtype of User-ID log; values are login and logout.
Generated Time (time_generated or cef-formatted-time_generated)	The time the log was generated on the dataplane.
Virtual System (vsys)	Virtual System associated with the configuration log.
Source IP (ip)	Original session source IP address.
User (user)	Identifies the end user.
Data Source Name (datasourcename)	User-ID source that sends the IP (Port)-User Mapping.
Event ID (eventid)	String showing the name of the event.
Repeat Count (repeatcnt)	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds.
Time Out Threshold (timeout)	Timeout after which the IP/User Mappings are cleared.
Source Port (beginport)	Source port utilized by the session.
Destination Port (endport)	Destination port utilized by the session.
Data Source (datasource)	Source from which mapping information is collected.
Data Source Type (datasourcetype)	Mechanism used to identify the IP/User mappings within a data source.
Sequence Number (seqno)	Serial number of the firewall that generated the log.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.

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: <code>/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show></code>
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.
Virtual System ID (vsys_id)	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Factor Type (factortype)	Vendor used to authenticate a user when Multi Factor authentication is present.
Factor Completion Time (factorcompletiontime)	Time the authentication was completed.
Factor Number (factorno)	Indicates the use of primary authentication (1) or additional factors (2, 3).
User Group Flags (ugflags)	Displays whether the user group that was found during user group mapping. Supported values are: • User Group Found—Indicates whether the user could be mapped to a group. • Duplicate User—Indicates whether duplicate users were found in a user group. Displays N/A if no user group is found.
User by Source (userbysource)	Indicates the username received from source through IP-User mapping.

Tunnel Inspection Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source Address, Destination Address, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Severity, Sequence Number, Action Flags, Source Location, Destination Location, 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, Tunnel ID/IMSI, Monitor Tag/IMEI, Parent Session ID, Parent Start Time, Tunnel, Bytes, Bytes Sent, Bytes Received,

Packets, Packets Sent, Packets Received, Maximum Encapsulation, Unknown Protocol, Strict Check, Tunnel Fragment, Sessions Created, Sessions Closed, Session End Reason, Action Source, Start Time, Elapsed Time, Tunnel Inspection Rule, UUID for rule

Field Name	Description
Receive Time (receive_time or cef-formatted- receive_time)	Month, day, and 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 as it pertains to the session: START or END.
Threat/Content Type (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 or cef-formatted- time_generated)	Time the log was generated on the dataplane.
Source Address (src)	Source IP address of packets in the session.
Destination Address (dst)	Destination IP address of packets in the session.
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 Security policy rule in effect on the session.
Source User (srcuser)	Source User ID of packets in the session.
Destination User (dstuser)	Destination User ID of packets in the session.
Application (app)	Tunneling protocol used in the session.
Virtual System (vsys)	Virtual System associated with the session.
Source Zone (from)	Source zone of packets in the session.
Destination Zone (to)	Destination zone of packets in the session.

Field Name	Description
Inbound Interface (inbound_if)	Interface that the session was sourced from.
Outbound Interface (outbound_if)	Interface that the session was destined to.
Log Action (logset)	Log Forwarding Profile that was applied to the session.
Session ID (sessionid)	Session ID of the session being logged.
Repeat Count (repeatcnt)	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds.
Source Port (sport)	Source port utilized by the session.
Destination Port (dport)	Destination port utilized by the session.
NAT Source Port (nat sport)	Post-NAT source port.
NAT Destination Port (nat dport)	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
IP 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

Field Name	Description
	- 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
Severity (severity)	Severity associated with the event; values are informational, low, medium, high, critical.
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.
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: <pre>/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show></pre>
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.
Tunnel ID (tunnelid)	ID of the tunnel being inspected or the International Mobile Subscriber Identity (IMSI) ID of the mobile user.
Monitor Tag (monitortag)	Monitor name you configured for the Tunnel Inspection policy rule or the International Mobile Equipment Identity (IMEI) ID of the mobile device.
Parent Session ID (parent_session_id)	ID of the session in which this session is tunneled. Applies to inner tunnel (if two levels of tunneling) or inside content (if one level of tunneling) only.

Field Name	Description
Parent Start Time (parent_start_time)	Year/month/day hours:minutes:seconds that the parent tunnel session began.
Tunnel Type (tunnel)	Type of tunnel, such as GRE or IPSec.
Bytes (bytes)	Number of bytes in the session.
Bytes Sent (bytes_sent)	Number of bytes in the client-to-server direction of the session.
Bytes Received (bytes_received)	Number of bytes in the server-to-client direction of the session.
Packets (packets)	Number of total packets (transmit and receive) for the session.
Packets Sent (pkts_sent)	Number of client-to-server packets for the session.
Packets Received (pkts_received)	Number of server-to-client packets for the session.
Maximum Encapsulation (max_encap)	Number of packets the firewall dropped because the packet exceeded the maximum number of encapsulation levels configured in the Tunnel Inspection policy rule (Drop packet if over maximum tunnel inspection level).
Unknown Protocol (unknown_proto)	Number of packets the firewall dropped because the packet contains an unknown protocol, as enabled in the Tunnel Inspection policy rule (Drop packet if unknown protocol inside tunnel).
Strict Checking (strict_check)	Number of packets the firewall dropped because the tunnel protocol header in the packet failed to comply with the RFC for the tunnel protocol, as enabled in the Tunnel Inspection policy rule (Drop packet if tunnel protocol fails strict header check).
Tunnel Fragment (tunnel_fragment)	Number of packets the firewall dropped because of fragmentation errors.
Sessions Created (sessions_created)	Number of inner sessions created.
Sessions Closed (sessions_closed)	Number of completed/closed sessions created.
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

Field Name	Description
	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-unsupported-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-unsupported-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 <code>clear session all</code> 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.
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.
Start Time (start)	Year/month/day hours:minutes:seconds that the session began.

Field Name	Description
Elapsed Time (elapsed)	Elapsed time of the session.
Tunnel Inspection Rule (tunnel_insp_rule)	Name of the tunnel inspection rule matching the cleartext tunnel traffic.
UUID for rule (rule_uuid)	The UUID that permanently identifies the rule.

SCTP Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, FUTURE_USE, FUTURE_USE, Generated Time, Source Address, Destination Address, FUTURE_USE, FUTURE_USE, Rule Name, FUTURE_USE, FUTURE_USE, FUTURE_USE, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, IP Protocol, Action, 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, Sequence Number, FUTURE_USE, SCTP Association ID, Payload Protocol ID, Severity, SCTP Chunk Type, FUTURE_USE, SCTP Verification Tag 1, SCTP Verification Tag 2, SCTP Cause Code, Diameter App ID, Diameter Command Code, Diameter AVP Code, SCTP Stream ID, SCTP Association End Reason, Op Code, SCCP Calling Party SSN, SCCP Calling Party Global Title, SCTP Filter, SCTP Chunks, SCTP Chunks Sent, SCTP Chunks Received, Packets, Packets Sent, Packets Received, UUID for rule

Field Name	Description
Receive Time (receive_time or cef-formatted-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 the type of log; value is SCTP.
Generated Time (time_generated or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Source Address (src)	Original session source IP address.
Destination Address (dst)	Original session destination IP address.
Rule Name (rule)	Name of the Security policy rule in effect on 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.
Inbound Interface (inbound_if)	Interface that the session was sourced from.

Field Name	Description
Outbound Interface (outbound_if)	Interface that the session was destined to.
Log Action (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.
Source Port (sport)	Source port utilized by the session.
Destination Port (dport)	Destination port utilized by the session.
IP Protocol (proto)	IP protocol associated with the session.
Action (action)	Action taken for the session; possible values are: - allow—session was allowed by the policy - deny—session was denied by the policy
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: <pre>/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show></pre>
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.
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space.
SCTP Association ID (assoc_id)	An internal 56-bit numerical logical identifier applied to each SCTP association.

Field Name	Description
Payload Protocol ID (ppid)	Identifies the Payload Protocol ID (PPID) in the data chunk which triggered this event. PPID is assigned by Internet Assigned Numbers Authority (IANA).
Severity (severity)	Severity associated with the event; values are informational, low, medium, high, critical.
SCTP Chunk Type (sctp_chunk_type)	Describes the type of information contained in a chunk, such as control or data.
SCTP Event Type (sctp_event_type)	Defines the event triggered per SCTP chunk or packet when SCTP protection profile is applied to the SCTP traffic. It is also triggered by start or end of a SCTP association.
SCTP Verification Tag 1 (verif_tag_1)	Used by endpoint1 which initiates the association to verify if the SCTP packet received belongs to current SCTP association and validate the endpoint2.
SCTP Verification Tag 2 (verif_tag_2)	Used by endpoint2 to verify if the SCTP packet received belongs to current SCTP association and validate the endpoint1.
SCTP Cause Code (sctp_cause_code)	Sent by an endpoint to specify reason for an error condition to other endpoint of same SCTP association.
Diameter App ID (diam_app_id)	The diameter application in the data chunk which triggered the event. Diameter Application ID is assigned by Internet Assigned Numbers Authority (IANA).
Diameter Command Code (diam_cmd_code)	The diameter command code in the data chunk which triggered the event. Diameter Command Code is assigned by Internet Assigned Numbers Authority (IANA).
Diameter AVP Code (diam_avp_code)	The diameter AVP code in the data chunk which triggered the event.
SCTP Stream ID (stream_id)	ID of the stream which carries the data chunk which triggered the event.
SCTP Association End Reason (assoc_end_reason)	Reason an association was terminated. If the termination had multiple causes, the highest priority reason is displayed. The possible session end reasons in descending priority are: - shutdown-from-endpoint (highest)—endpoint sends out SHUTDOWN - abort-from-endpoint—endpoint sends out ABORT - unknown (lowest)—the association aged out, or association termination reason is not covered by one of the previous reasons (for example, a clear session all command).
Op Code (op_code)	Identifies the operation code of application layer SS7 protocols, like MAP or CAP, in the data chunk which triggered the event.

Field Name	Description
SCCP Calling Party SSN (sccp_calling_ssn)	The Signaling Connection Control Part (SCCP) calling party subsystem number (SSN) in the data chunk which triggered the event.
SCCP Calling Party Global Title (sccp_calling_gt)	The Signaling Connection Control Part (SCCP) calling party global title (GT) in the data chunk which triggered the event.
SCTP Filter (sctp_filter)	Name of the filter that the SCTP chunk matched.
SCTP Chunks (chunks)	Number of total chunks (transmit and receive) for the association.
SCTP Chunks Sent (chunks_sent)	Number of endpoint1(which initiates association)-to-endpoint2 chunks for the association.
SCTP Chunks Received (chunks_received)	Number of endpoint2-to-endpoint1(which initiates association) chunks for the association.
Packets (packets)	Number of total packets (transmit and receive) for the session.
Packets Sent (pkts_sent)	Number of client-to-server packets for the session.
Packets Received (pkts_received)	Number of server-to-client packets for the session.
UUID for rule (rule_uuid)	The UUID that permanently identifies the rule.

Authentication Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Virtual System, Source IP, User, Normalize User, Object, Authentication Policy, Repeat Count, Authentication ID, Vendor, Log Action, Server Profile, Description, Client Type, Event Type, Factor Number, Sequence Number, Action Flags, Device Group Hierarchy 1, Device Group Hierarchy 2, Device Group Hierarchy 3, Device Group Hierarchy 4, Virtual System Name, Device Name, Virtual System ID, Authentication Protocol, UUID for rule

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial)	Serial number of the device that generated the log.
Type (type)	Specifies the type of log; value is AUTHENTICATION.
Threat/Content Type (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.

Field Name	Description
Generated Time (time_generated or cef-formatted- time_generated)	Time the log was generated on the dataplane.
Virtual System (vsys)	Virtual System associated with the session.
Source IP (ip)	Original session source IP address.
User (user)	End user being authenticated.
Normalize User (normalize_user)	Normalized version of username being authenticated (such as appending a domain name to the username).
Object (object)	Name of the object associated with the system event.
Authentication Policy (authpolicy)	Policy invoked for authentication before allowing access to a protected resource.
Repeat Count (repeatcnt)	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds.
Authentication ID (authid)	Unique ID given across primary authentication and additional (multi factor) authentication.
Vendor (vendor)	Vendor providing additional factor authentication.
Log Action (logset)	Log Action (logset)
Server Profile (serverprofile)	Authentication server used for authentication.
Description (desc)	Additional authentication information.
Client Type (clienttype)	Type of client used to complete authentication (such as authentication portal).
Event Type (event)	Result of the authentication attempt.
Factor Number (factorno)	Indicates the use of primary authentication (1) or additional factors (2, 3).
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially. Each log type has a unique number space.
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.

Field Name	Description
	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 readonlydg-meta-data API query: <pre>/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show></pre>
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.
Virtual System ID (vsys_id)	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Authentication Protocol (authproto)	Indicates the authentication protocol used by the server. For example, PEAP with GTC.
UUID for rule (rule_uuid)	The UUID that permanently identifies the rule.

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, Before Change Detail, After Change Detail, 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 or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial)	Serial number of the device that generated the log.
Type (type)	Specifies the type of log; value is CONFIG.
Threat/Content Type (subtype)	Subtype of the configuration log; unused.
Generated Time (time_generated)	Time the log was generated on the dataplane.

Field Name	Description
or cef-formatted-time_generated)	
Host (host)	Hostname or IP address of the client machine
Virtual System (vsys)	Virtual System associated with the configuration log
Command (cmd)	Command performed by the Admin; values are add, clone, commit, delete, edit, move, rename, set.
Admin (admin)	Username of the Administrator performing the configuration
Client (client)	Client used by the Administrator; values are Web and CLI
Result (result)	Result of the configuration action; values are Submitted, Succeeded, Failed, and Unauthorized
Configuration Path (path)	The path of the configuration command issued; up to 512 bytes in length
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.
Sequence Number (seqno)	A 64bit log entry identifier incremented sequentially; each log type has a unique number space.
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: <pre>/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show></pre>
Virtual System Name (vsys_name)	The name of the virtual system associated with the session; only valid on firewalls enabled for multiple virtual systems.

Field Name	Description
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, Content/Threat Type, 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 or cef-formatted-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 the type of log; value is SYSTEM.
Content/Threat Type (subtype)	Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn.
Generated Time (time_generated or cef-formatted-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.

Field Name	Description
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: <pre>/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show></pre>
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, Content/Threat Type, FUTURE_USE, Generated Time, Source Address, Source User, Virtual System, Category, Severity, 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, Virtual System ID, Object Name, Object ID, Evidence

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Time the log was received at the management plane.
Serial Number (serial)	Serial number of the device that generated the log.
Type (type)	Specifies the type of log; value is CORRELATION.
Content/Threat Type (subtype)	Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn.
Generated Time (time_generated)	Time the log was generated on the dataplane.

Field Name	Description
or cef-formatted-time_generated)	
Source Address (src)	IP address of the user who initiated the event.
Source User (srcuser)	Username of the user who initiated the event.
Virtual System (vsys)	Virtual System associated with the configuration log.
Category (category)	A summary of the kind of threat or harm posed to the network, user, or host.
Severity (severity)	Severity associated with the event; values are informational, low, medium, high, critical.
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 readonlydg-meta-data API query: <pre>/api/?type=op&cmd=<show><dg-hierarchy></dg-hierarchy></show></pre>
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.
Virtual System ID (vsys_id)	A unique identifier for a virtual system on a Palo Alto Networks firewall.
Object Name (objectname)	Name of the correlation object that was matched on.
Object ID (object_id)	Name of the object associated with the system event.
Evidence (evidence)	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 URI (19 times).

GTP Log Fields

Format: FUTURE_USE, Receive Time, Serial Number, Type, Threat/Content Type, FUTURE_USE, Generated Time, Source Address, Destination Address, FUTURE_USE, FUTURE_USE, Rule Name,

FUTURE_USE, FUTURE_USE, Application, Virtual System, Source Zone, Destination Zone, Inbound Interface, Outbound Interface, Log Action, FUTURE_USE, Session ID, FUTURE_USE, Source Port, Destination Port, FUTURE_USE, FUTURE_USE, FUTURE_USE, Protocol, Action, GTP Event Type, MSISDN, Access Point Name, Radio Access Technology, GTP Message Type, End User IP Address, Tunnel Endpoint Identifier1, Tunnel Endpoint Identifier2, GTP Interface, GTP Cause, Severity, Serving Country MCC, Serving Network MNC, Area Code, Cell ID, GTP Event Code, FUTURE_USE, FUTURE_USE, Source Location, Destination Location, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, Tunnel ID/IMSI, Monitor Tag/IMEI, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, FUTURE_USE, Start Time, Elapsed Time, Tunnel Inspection Rule, Remote User IP, Remote User ID, UUID for rule, PCAP ID

Field Name	Description
Receive Time (receive_time or cef-formatted-receive_time)	Month, Day and time the log was received at the management plane.
Serial Number (serial)	Serial number of the firewall that generated the log.
Type (type)	Specifies the type of log; value is GTP.
Threat/Content Type (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 or cef-formatted-time_generated)	Time the log was generated on the dataplane.
Source Address (src)	Source IP address of packets in the session.
Destination Address (dst)	Destination IP address of packets in the session.
Rule Name (rule)	Name of the Security policy rule in effect on the session.
Application (app)	Tunneling protocol used in the session.
Virtual System (vsys)	Virtual System associated with the session.
Source Zone (from)	Source zone of packets in the session.
Destination Zone (to)	Destination zone of packets in the session.
Inbound Interface (inbound_if)	Interface that the session was sourced from.
Outbound Interface (outbound_if)	Interface that the session was destined to.

Field Name	Description
Log Action (logset)	Log Forwarding Profile that was applied to the session.
Session ID (sessionid)	Session ID of the session being logged.
Source Port (sport)	Source port utilized by the session.
Destination Port (dport)	Destination port utilized by the session.
IP 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
GTP Event Type (event_type)	Defines event triggered by a GTP message when checks in GTP protection profile are applied to the GTP traffic. Also triggered by the start or end of a GTP session.
MSISDN (msisdn)	Service identity associated with the mobile subscriber composed of a Country Code, National Destination Code and a Subscriber. Consists of decimal digits (0-9) only with a maximum of 15 digits.
Access Point Name (apn)	Reference to a Packet Data Network Data Gateway (PGW)/ Gateway GPRS Support Node in a mobile network. Composed of a mandatory APN Network Identifier and an optional APN Operator Identifier.
Radio Access Technology (rat)	Type of technology used for radio access. For example, EUTRAN, WLAN, Virtual, HSPA Evolution, GAN and GERAN.
GTP Message Type (msg_type)	Indicates the GTP message type.
End IP Address (end_ip_adr)	IP address of a mobile subscriber allocated by a PGW/GGSN.
Tunnel Endpoint Identifier1 (teid1)	Identifies the GTP tunnel in the network node. TEID1 is the first TEID in the GTP message.
Tunnel Endpoint Identifier2 (teid2)	Identifies the GTP tunnel in the network node. TEID2 is the second TEID in the GTP message.
GTP Interface (gtp_interface)	3GPP interface from which a GTP message is received.
GTP Cause (cause_code)	GTP cause value in logs responses which contain an Information Element that provides information about acceptance or rejection of GTP requests by a network node.
Severity (severity)	Severity associated with the event; values are informational, low, medium, high, critical.
Serving Network MCC (mcc)	Mobile country code of serving core network operator.

Field Name	Description
Serving Network MNC (mnc)	Mobile network code of serving core network operator.
Area Code (area_code)	Area within a Public Land Mobile Network (PLMN).
Cell ID (cell_id)	Base station within an area code.
GTP Event Code (event_code)	Event code describing the GTP event.
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.
Tunnel ID/IMSI (imsi)	International Mobile Subscriber Identity (IMSI) is a unique number allocated to each mobile subscriber in the GSM/UMTS/EPS system. IMSI shall consist of decimal digits (0 through 9) only and maximum number of digits allowed are 15.
Monitor Tag/IMEI (imei)	International Mobile Equipment Identity (IMEI) is a unique 15 or 16 digit number allocated to each mobile station equipment.
Start Time (start)	Time of session start.
Elapsed Time (elapsed)	Elapsed time of the session.
Tunnel Inspection Rule (tunnel_insp_rule)	Name of the tunnel inspection rule matching the cleartext tunnel traffic
Remote User IP (remote_user_ip)	IPv4 or IPv6 address used by a remote user.
Remote User ID (remote_user_id)	IMSI identity of a remote user, and if available, one IMEI identity and/or one MSISDN identity.
UUID for rule (rule_uuid)	Universally Unique ID for rule.
PCAP ID (pcap_id)	Unique packet capture ID that is used to locate the pcap file saved on the firewall.

Syslog Severity

The syslog severity is set based on the log type and contents.

Log Type/Severity	Syslog Severity
Traffic	Info
Config	Info

Log Type/Severity	Syslog Severity
Threat/System—Informational	Info
Threat/System—Low	Notice
Threat/System—Medium	Warning
Threat/System—High	Warning
Threat/System—Critical	Critical

Custom Log/Event Format

To facilitate the integration with external log parsing systems, the firewall allows you to customize the log format; it also allows you to add custom *Key: Value* attribute pairs. Custom message formats can be configured under **Device > Server Profiles > Syslog > Syslog Server Profile > Custom Log Format**.

To achieve ArcSight Common Event Format (CEF) compliant log formatting, refer to the [CEF Configuration Guide](#).

Escape Sequences

Any field that contains a comma or a double-quote is enclosed in double quotes. Furthermore, if a double-quote appears inside a field it is escaped by preceding it with another double-quote. To maintain backward compatibility, the Misc field in threat log is always enclosed in double-quotes.

SNMP Monitoring and Traps

The following topics describe how Palo Alto Networks firewalls, Panorama, and WF-500 appliances implement 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 an 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](#).

SNMPVer	Authentication	Message Privacy	Message	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

[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 Centralized Logging and Reporting](#). 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 2: 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.

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 | (Optional) 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:

SNMP MIBs		Result Table			
MIB Tree		Name/OID	Value	Type	IP-Port
- iso.org.dod.internet - mgmt - private - enterprises - panRoot - panMibs - panCommonMib - panSysVersion		panSysVersion.0		OctetString	10.5.68.19:161
		panSysTimeZoneOffset.0	-28800	Integer	10.5.68.19:161
		panSysDaylightSaving.0	0	Integer	10.5.68.19:161
		panSysThreatVersion.0	0	OctetString	10.5.68.19:161
		panSysUrpfFilteringVersion.0	0	OctetString	10.5.68.19:161
		panSysOperatDatafileVersion.0	0	OctetString	10.5.68.19:161
		1.3.6.1.4.1.25461.2.1.2.1.17.0	0	OctetString	10.5.68.19:161
		1.3.6.1.4.1.25461.2.1.3.1.18.0	0	OctetString	10.5.68.19:161
		panSysVpnClientVersion.0	0.0.0	OctetString	10.5.68.19:161
		panSysGlobalProtectClientVersion.0	0.0.0	OctetString	10.5.68.19:161
		panSysSerialNumber.0	0007PM00001	OctetString	10.5.68.19:161
		panSysAvVersion.0	1751-2367	OctetString	10.5.68.19:161
		panSysAppVersion.0	465-2420	OctetString	10.5.68.19:161
		panSysVersion.0	7.0.0-48	OctetString	10.5.68.19:161
		panSysHkState.0	disabled	OctetString	10.5.68.19:161
		panSysUrpfMode.0	disabled	OctetString	10.5.68.19:161
		panSysUrpfFilteringDatabase.0	paloaltonetworks	OctetString	10.5.68.19:161
		panSysHkPeerState.0	unknown	OctetString	10.5.68.19:161


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.

- 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}
```

- STEP 3** | In a MIB browser, search the MIB tree for the identified object name to display its OID. For example, the panSysHwVersion object has an OID of 1.3.6.1.4.1.25461.2.1.2.1.2.



Enable SNMP Services for Firewall-Secured Network Elements

If you will use Simple Network Management Protocol (SNMP) to monitor or manage network elements (for example, switches and routers) that are within the security zones of Palo Alto Networks firewalls, you must create a security rule that allows SNMP services for those elements.



You don't need a security rule to enable SNMP monitoring of Palo Alto Networks firewalls, Panorama, or WF-500 appliances. For details, see [Monitor Statistics Using SNMP](#).

- 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](#).

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.



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 > Interfaces**, edit the **Management** interface, 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.

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](#).

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:
 1. [Create a Log Forwarding profile](#). For each log type and each severity level or WildFire verdict, select the **SNMP Trap** server profile.
 2. [Assign the Log Forwarding profile to policy rules and network zones](#). The rules and zones will trigger trap generation and forwarding.
2. [Configure the destinations for System, Configuration, User-ID, 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](#).

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.  <i>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.</i>	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-5200 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-5260 firewall, which has multiple dataplanes (DPs), the average is of the cores across 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.
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 <i>child</i> component among its <i>sibling</i> 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.
entPhysicalFirmwareRev	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).

Object	Description
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/ID	Value
entPhysicalDescr 1	Fan: 3020
entPhysicalDescr 2	Fan: 3020
entPhysicalDescr 3	Fan: 3020
entPhysicalDescr 4	Fan: 3020
entPhysicalDescr 5	Temperature @ Serial
entPhysicalDescr 6	Temperature @ Switch
entPhysicalDescr 7	Temperature @ Catalyst
entPhysicalDescr 8	Temperature @ Serial
entPhysicalDescr 9	Temperature @ Switch-Core
entPhysicalDescr 10	Temperature @ Catalyst Core
entPhysicalDescr 11	Temperature @ Catalyst Core



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.



The dot3adTablesLastChanged object indicates the time of the most recent change to dot3adAggTable, dot3adAggPortListTable, and dot3adAggPortTable.

Table	Description
Aggregator Configuration Table (dot3adAggTable)	This table contains information about every aggregate group that is associated with a firewall. Each aggregate group has one entry. 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.  <i>The ifTable MIB (a list of interface entries) does not support logical interfaces and therefore does not have an entry for the aggregate group.</i>
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.

Trap Name	Description
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 `IldpV2StatsRxPortFramesDiscardedTotal` 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 `IldpV2Statistics` objects:
 - `IldpV2StatsRemTablesLastChangeTime`
 - `IldpV2StatsRemTablesInserts`
 - `IldpV2StatsRemTablesDeletes`
 - `IldpV2StatsRemTablesDrops`
 - `IldpV2StatsRemTablesAgeouts`
- The following `IldpV2RemoteSystemsData` objects:
 - The `IldpV2RemOrgDefInfoTable` table
 - In the `IldpV2RemTable` table: `IldpV2RemTimeMark`

[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,

Object Group	Description
	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	Logging statistics for each Log Collector, including logging rate, log quotas, disk usage, retention periods, log redundancy (enabled or disabled), the forwarding status from firewalls to Log Collectors, the forwarding status from Log Collectors to external services, and the status of firewall-to-Log Collector connections.
panDeviceLogging	Logging statistics for each firewall, including logging rate, disk usage, retention periods, the forwarding status from individual firewalls to Panorama and external servers, and the status of firewall-to-Log Collector connections.

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.

Forward Logs to an HTTP/S Destination

The firewall and Panorama™ can forward logs to an HTTP/S server. You can choose to forward all logs or specific logs to trigger an action on an external HTTP-based service when an event occurs. When forwarding logs to an HTTP server, choose from the following options:

- Configure the firewall to send an HTTP-based API request directly to a third-party service to trigger an action that is based on the attributes in a firewall log. You can configure the firewall to work with any HTTP-based service that exposes an API and you can modify the URL, HTTP header, parameters, and the payload in the HTTP request to meet your integration needs.
- Tag the source or destination IP address in a log entry automatically and register the IP address and tag mapping to a User-ID™ agent on the firewall or Panorama, or register to a remote User-ID agent so that you can respond to an event and dynamically enforce security policy. To enforce policy, you must [Use Dynamic Address Groups in Policy](#).

STEP 1 | Create an HTTP server profile to forward logs to an HTTP/S destination.

The HTTP server profile allows you to specify how to access the server and define the format in which to forward logs to the HTTP/S destination. By default, the firewall uses the management port to forward these logs. However, you can assign a different source interface and IP address in **Device > Setup > Services > Service Route Configuration**.

1. Select **Device > Server Profiles > HTTP** and **Add** a new profile.
2. Specify a **Name** for the server profile, and select the **Location**. The profile can be **Shared** across all virtual systems or can belong to a specific virtual system.
3. **Add** the details for each server. Each profile can have a maximum of four servers.
4. Enter a **Name** and **IP Address**.
5. Select the **Protocol (HTTP or HTTPS)**. The default **Port** is 80 or 443 respectively but you can modify the port number to match the port on which your HTTP server listens.
6. Select the **TLS Version** supported on the server—**1.0**, **1.1**, or **1.2** (default).
7. Select the **Certificate Profile** to use for the TLS connection with the server.
8. Select the **HTTP Method** that the third-party service supports—**DELETE**, **GET**, **POST** (default), or **PUT**.
9. (Optional) Enter the **Username** and **Password** for authenticating to the server, if needed.
10. (Optional) Select **Test Server Connection** to verify network connectivity between the firewall and the HTTP/S server.



STEP 2 | Select the Payload Format for the HTTP request.

1. Select the **Log Type** link for each log type for which you want to define the HTTP request format.

2. Select the **Pre-defined Formats** (available through content updates) or create a custom format.


If you create a custom format, the **URI** is the resource endpoint on the HTTP service. The firewall appends the URI to the IP address you defined earlier to construct the URL for the HTTP request. Ensure that the URI and payload format matches the syntax that your third-party vendor requires. You can use any attribute supported on the selected log type within the HTTP Header, the Parameter and Value pairs, and in the request payload.



3. **Send Test Log** to verify that the HTTP server receives the request. When you interactively send a test log, the firewall uses the format as is and does not replace the variable with a value from a firewall log. If your HTTP server sends a 404 response, provide values for the parameters so that the server can process the request successfully.

STEP 3 | Define the match criteria for when the firewall will forward logs to the HTTP server and attach the HTTP server profile you will use.

1. Select the log types for which you want to trigger a workflow:
 - Add a Log Forwarding Profile (**Objects > Log Forwarding**) for logs that pertain to user activity (for example, Traffic, Threat, or Authentication logs).
 - Select **Device > Log Settings** for logs that pertain to system events, such as Configuration or System logs.
2. Select the Log Type and use the new **Filter Builder** to define the match criteria.
3. **Add** the HTTP server profile for forwarding logs to the HTTP destination.



4. Add a tag to the source or destination IP address in the log entry. This capability enables you to use dynamic address groups and Security policy rules to limit network access or isolate the IP address until you can triage the affected user device.

Add Built-in Actions and Name the action. Select the following options to register the tag on the remote User-ID agent:

- **Target**—Select **Source Address** or **Destination Address**.
- **Action**—Select **Add Tag**.



You cannot configure a Timeout with a Remove Tag action.

- **Registration**—Register tag to the local User-ID on a firewall or to the Panorama that is managing the firewall.
- **Timeout (min)**—Enter the amount of time (in minutes) before the firewall removes IP address-to-tag mapping. If you enter a value of zero (0) minutes, the IP address-to-tag mapping does not timeout.



You should set the IP-tag timeout to the same amount of time as the DHCP lease timeout for that IP address. This allows the IP address-to-tag mapping to expire at the same time as the DHCP lease so that you do not unintentionally apply policy when the IP address is reassigned.

- **Tag**—Enter a new tag or select an existing tag.

STEP 4 | Register or unregister a tag on a source or destination IP address in a log entry to a remote User-ID agent.

1. Select **Device > Server Profiles > HTTP**, specify a **Name** for the server profile, and select the **Location**. The profile can be **Shared** across all virtual systems or can belong to a specific virtual system.
2. Select **Tag Registration** to enable the firewall to register the IP address and tag mapping with the User-ID agent on a remote firewall. With tag registration enabled, you cannot specify the payload format.
3. Add the connection details to access the remote User-ID agent.



4. Select the log type (**Objects > Log Forwarding** or **Device > Log Settings**) for which you want to add a tag to the source or destination IP address in the log entry.
5. **Add Built-in Actions** and **Name** the action. Select the following options to register the tag on the remote User-ID agent:
 - **Target**—Select **Source Address** or .
 - **Action**—Select **Add Tag** or **Remove Tag**.



You cannot configure a Timeout with a Remove Tag action.

- **Registration**—Select **Remote User-ID** agent.
- **Timeout (min)**—Enter the amount of time (in minutes) before the firewall removes IP address-to-tag mapping. If you enter a value of zero (0) minutes, the IP address-to-tag mapping does not timeout.



You should set the IP-tag timeout to the same amount of time as the DHCP lease timeout for that IP address. This allows the IP address-to-tag mapping to expire at the same time as the DHCP lease so that you do not unintentionally apply policy when the IP address is reassigned.

- **Tag**—Enter a new tag or select .



For dynamic policy enforcement, [Use Dynamic Address Groups in Policy](#).

NetFlow Monitoring

NetFlow is an industry-standard protocol that the firewall can use to export statistics about the IP traffic on 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. The firewalls support only unidirectional NetFlow, not bidirectional. The firewalls perform NetFlow processing on all IP packets on the interfaces and do not support sampled NetFlow. You can export NetFlow records for Layer 3, Layer 2, virtual wire, tap, VLAN, loopback, and tunnel interfaces. For aggregate Ethernet sub-interfaces, you can export records for the individual sub-interfaces that data flows through within the group. To identify firewall interfaces in a NetFlow collector, see [Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors](#). The firewalls support standard and enterprise (PAN-OS specific) [NetFlow Templates](#), which NetFlow collectors use to decipher the NetFlow fields.

- [Configure NetFlow Exports](#)
- [NetFlow Templates](#)

Configure NetFlow Exports

To use a NetFlow collector for analyzing the network traffic on firewall interfaces, perform the following steps to configure NetFlow record exports.

STEP 1 | Create a NetFlow server profile.

The profile defines which NetFlow collectors will receive the exported records and specifies export parameters.

1. Select **Device > Server Profiles > NetFlow** and **Add** a profile.
2. Enter a **Name** to identify the profile.
3. Specify the rate at which the firewall refreshes [NetFlow Templates](#) in **Minutes** (default is 30) and **Packets** (exported records—default is 20), according to the requirements of your NetFlow collector. The firewall refreshes the templates after either threshold is passed.
4. Specify the **Active Timeout**, which is the frequency in minutes at which the firewall exports records (default is 5).
5. Select **PAN-OS Field Types** if you want the firewall to export App-ID and User-ID fields.
6. **Add** each NetFlow collector (up to two per profile) that will receive records. For each collector, specify the following:
 - **Name** to identify the collector.
 - **NetFlow Server** hostname or IP address.
 - Access **Port** (default 2055).
7. Click **OK** to save the profile.

STEP 2 | Assign the NetFlow server profile to the firewall interfaces that convey 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.



You can export NetFlow records for Layer 3, Layer 2, virtual wire, tap, VLAN, loopback, and tunnel interfaces. For aggregate Ethernet interfaces, you can export records for the aggregate group but not for individual interfaces within the group.

2. Select the NetFlow server profile (**NetFlow Profile**) you configured and click **OK**.

STEP 3 | (Required for PA-7000 Series and PA-5200 Series firewalls) Configure a service route for the interface that the firewall will use to send NetFlow records.

You cannot use the management (MGT) interface to send NetFlow records from the PA-7000 Series and PA-5200 Series firewalls. For other firewall models, a service route is optional. For all firewalls, the interface that sends NetFlow records does not have to be the same as the interface for which the firewall collects the records.

1. Select **Device > Setup > Services**.
2. **(Firewall with multiple virtual systems)** Select one of the following:
 - **Global**—Select this option if the service route applies to all virtual systems on the firewall.
 - **Virtual Systems**—Select this option if the service route applies to a specific virtual system. Set the **Location** to the virtual system.
3. Select **Service Route Configuration** and **Customize**.
4. Select the protocol (**IPv4** or **IPv6**) that the interface uses. You can configure the service route for both protocols if necessary.
5. Click **Netflow** in the Service column.
6. Select the **Source Interface**.

Any, Use default, and MGT are not valid interface options for sending NetFlow records from PA-7000 Series or PA-5200 Series firewalls.
7. Select a **Source Address** (IP address).
8. Click **OK** twice to save your changes.

STEP 4 | Commit your changes.

STEP 5 | Monitor the firewall traffic in a NetFlow collector.

Refer to your NetFlow collector documentation.



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](#).

To troubleshoot NetFlow delivery issues, use the operational CLI command **debug log-receiver netflow statistics**.

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](#), set the refresh rate based on a time interval and a number of exported records according to the requirements of your NetFlow collector. The firewall refreshes the templates after either threshold is passed.

The Palo Alto Networks firewall supports the following NetFlow templates:

Template	ID
IPv4 Standard	256
IPv4 Enterprise	257

Template	ID
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 flow. 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
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	All templates

Value	Field	Description	Templates
		Firewall Interface Identifiers in SNMP Managers and NetFlow Collectors.	
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:	All templates

Value	Field	Description	Templates
		ICMP Type * 256 + ICMP code	
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
233	firewallEvent	Indicates a firewall event: 0 = Ignore (invalid)—Not used. 1 = Flow created—The NetFlow data record is for a new flow. 2 = Flow deleted—The NetFlow data record is for the end of a flow. 3 = Flow denied—The NetFlow data record indicates a flow that firewall policy denied. 4 = Flow alert—Not used. 5 = Flow update—The NetFlow data record is sent for a <i>long-lasting</i> flow, which is a flow that lasts longer than the Active Timeout period configured in the NetFlow server profile.	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	IPv4 with NAT standard

Value	Field	Description	Templates
		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 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
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

Value	Field	Description	Templates
			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

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 [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 3: 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
VM-Series	Number of management ports + physical port offset Number of management ports—This is a constant of 1. Physical port offset—This is the physical port number.	VM-100 firewall, Eth1/4 = 1 (number of management ports) + 4 (physical port) = 5
PA-220, PA-220R, PA-800 Series	Number of management ports + physical port offset Number of management ports—This is a constant of 5. Physical port offset—This is the physical port number.	PA-5200 Series firewall, Eth1/4 = 5 (number of management ports) + 4 (physical port) = 9
PA-3200 Series, PA-5200 Series	Number of management ports + physical port offset Number of management ports—This is a constant of 4. Physical port offset—This is the physical port number.	PA-5200 Series firewall, Eth1/4 = 4 (number of management ports) + 4 (physical port) = 8
PA-7000 Series	(Max. ports * slot) + physical port offset + number of management ports	PA-7000 Series firewall, Eth3/9 =

Firewall Platform	Calculation	Example Interface Index
	• 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. • Number of management ports—This is a constant of 5.	$[64 (\text{max. ports}) * 3 (\text{slot})] + 9 (\text{physical port}) + 5 (\text{number of management ports}) = \mathbf{206}$

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

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 [Authentication Policy](#) 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 4: User-ID

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. [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](#)
- [Authentication Policy and Captive Portal](#)
- [GlobalProtect](#)
- [XML API](#)
- [Client Probing](#)

Server Monitoring

With server monitoring a User-ID agent—either a Windows-based agent running on a domain server in your network, or the PAN-OS integrated 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. 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](#)

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

If you have a proxy server deployed between the users on your network and the firewall, the firewall might see the proxy server IP address as the source IP address in HTTP/HTTPS 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 extract the end user IP address from the XFF so that User-ID can map the IP address to a username. This enables you to [Use XFF Values for Policies and Logging Source Users](#) so that you can enforce user-based policy to safely enable access to web-based for your users behind a proxy server.

Authentication Policy and Captive Portal

In some cases, the User-ID agent can't map an IP address to a username using server monitoring or other methods—for example, if the user isn't logged in or uses an operating system such as Linux that your domain servers don't support. In other cases, you might want users to authenticate when accessing sensitive applications regardless of which methods the User-ID agent uses to perform user mapping. For all these cases, you can configure [Configure Authentication Policy](#) and [Map IP Addresses to Usernames Using Captive Portal](#). Any web traffic (HTTP or HTTPS) that matches an Authentication policy rule prompts the user to authenticate through Captive Portal. You can use the following [Captive Portal Authentication Methods](#):

- Browser challenge—Use [Kerberos](#) single sign-on (recommended) or [NT LAN Manager \(NTLM\)](#) authentication if you want to reduce the number of login prompts that users must respond to.
- Web form—Use [Multi-Factor Authentication](#), [SAML](#) single sign-on, [Kerberos](#), [TACACS+](#), [RADIUS](#), [LDAP](#), or [Local Authentication](#).
- [Client Certificate Authentication](#).

Syslog

Your environment might have existing network services that authenticate users. These services include wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, and other Network Access Control (NAC) mechanisms. You can configure these services to send syslog messages that contain information about login and logout events and configure the User-ID agent to parse those messages. The User-ID agent parses for login events to map IP addresses to usernames and parses for logout events to delete outdated mappings. Deleting outdated mappings is particularly useful in environments where IP address assignments change often.

Both the PAN-OS integrated User-ID agent and Windows-based User-ID agent use Syslog Parse profiles to parse syslog messages. In environments where services send the messages in different formats, you can create a custom profile for each format and associate multiple profiles with each syslog sender. If you use

the PAN-OS integrated User-ID agent, you can also use predefined Syslog Parse profiles that Palo Alto Networks provides through Applications content updates.

Syslog messages must meet the following criteria for a User-ID agent to parse them:

- Each message must be a single-line text string. The allowed delimiters for line breaks are a new line (\n) or a carriage return plus a new line (\r\n).
- 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.

See [Configure User-ID to Monitor Syslog Senders for User Mapping](#) for configuration details.



Figure 5: User-ID Integration with Syslog

GlobalProtect

For mobile or roaming users, the GlobalProtect endpoint provides the user mapping information to the firewall directly. In this case, every GlobalProtect user has an app running on the endpoint 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](#).

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 [Authentication Policy and 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:

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 services and applications.

1. Select **Network > Zones** and click the **Name** of the zone.
2. **Enable User Identification** and click **OK**.

STEP 2 | Create a Dedicated Service Account for the User-ID Agent.



As a best practice, 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, Microsoft Exchange servers, or 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.

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 Monitor Syslog Senders for User Mapping.](#)
- [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.

The way you specify which networks to include and exclude depends on whether you are using the [Windows-based](#) User-ID agent or the [PAN-OSintegrated](#) User-ID agent.

STEP 6 | Configure [Authentication Policy](#) and [Captive Portal](#).

The firewall uses Captive Portal to authenticate end users when they request services, applications, or URL categories that match [Authentication Policy](#) rules. Based on user information collected during authentication, the firewall creates new user mappings or updates existing mappings. The mapping information collected during authentication overrides information collected through other User-ID methods.

1. [Configure Captive Portal](#).
2. [Configure Authentication Policy](#).

STEP 7 | 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 **User** and specify which users and groups to match in the rule in one of the following ways:
 - If you want to select specific users or groups as matching criteria, click **Add** in the Source User section to display a list of users and groups discovered by the firewall group mapping function. Select the users 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](#).

STEP 8 | 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 9 | Configure the firewall to obtain user IP addresses from X-Forwarded-For (XFF) headers.

When the firewall is between the Internet and a proxy server, the IP addresses in the packets that the firewall sees are for the proxy server rather than users. To enable visibility of user IP addresses instead, configure the firewall to use the XFF headers 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 **X-Forwarded-For Header in User-ID**.



Selecting Strip-X-Forwarded-For Header 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 | If you use a high availability (HA) configuration, enable synchronization.



As a best practice, always enable the Enable Config Sync option for an HA configuration to ensure that the group mappings and user mappings are synchronized between the active and passive firewall.

1. Select **Device > High Availability > General** and edit the Setup section.
2. Select **Enable HA**.
3. Select **Enable Config Sync**.
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 11 | Commit your changes.

Commit your changes to activate them.

STEP 12 | [Verify the User-ID Configuration](#).

After you configure user mapping and group mapping, verify that the configuration works properly and that you can safely enable and monitor user and group access to your applications 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. The number of distinct user groups that each firewall or Panorama can reference across all policies varies by model. For more information, [refer](#) to the Compatibility Matrix.

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 group mapping configuration with an LDAP server profile that connects the firewall to the domain controller with the best connectivity. You can add up to four domain controllers to the LDAP server profile for redundancy. Note that you cannot increase redundancy beyond four domain controllers for a single domain by adding multiple group mapping configurations for that domain.*
- *If you have multiple domains and/or multiple forests, you must create a group mapping configuration with an LDAP server profile that connects the firewall to a domain server in each domain/forest. Take steps to ensure unique usernames in separate forests.*
- *If you have Universal Groups, create an LDAP server profile to connect to the Global Catalog server.*
- *Before using group mapping, configure a Primary Username for user-based security policies, since this attribute will identify users in the policy configuration, logs, and reports.*

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.

1. Select **Device > Server Profiles > LDAP** and **Add** a server profile.
2. Enter a **Profile Name** to identify the server profile.
3. **Add** the LDAP servers. You can add up to four servers to the profile but they must be the same **Type**. For each server, enter a **Name** (to identify the server), **LDAP Server** IP address or FQDN, and server **Port** (default 389).
4. Select the server **Type**.

Based on your selection (such as **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.

5. For the **Base DN**, enter the Distinguished Name (DN) of the LDAP tree location where you want the firewall to start searching for user and group information.
6. For the **Bind DN**, **Password** and **Confirm Password**, enter the authentication credentials for binding to the LDAP tree.

The **Bind DN** can be a fully qualified LDAP name (such as `cn=administrator,cn=users,dc=acme,dc=local`) or a user principal name (such as `administrator@acme.local`).

7. Enter the **Bind Timeout** and **Search Timeout** in seconds (default is 30 for both).
8. Click **OK** to save the server profile.

STEP 2 | Configure the server settings in a group mapping configuration.

1. Select **Device > User Identification > Group Mapping Settings**.

2. **Add** the group mapping configuration.
3. 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) and **Object Class** (group definition).
7. (Optional) To filter the users that the firewall tracks for group mapping, in the User Objects section, enter a **Search Filter** (LDAP query), and **Object Class** (user definition).
8. Make sure the group mapping configuration is **Enabled** (default is enabled).

STEP 3 | (Optional) Define User and Group Attributes to collect for user and group mapping. This step is required if you want to map users based on directory attributes other than the domain.

1. If your User-ID sources only send the username and the username is unique across the organization, select **Device > User Identification > User Mapping > Setup** and **Edit** the Setup section to **Allow matching usernames without domains** 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.



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.

2. Select **Device > User Identification > Group Mapping Settings > Add > User and Group Attributes > User Attributes** and enter the **Directory Attribute** you want to collect for user identification. Specify a **Primary Username** to identify the user on the firewall and to represent the user in reports and logs that will override any other format the firewall receives from the User-ID source.

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:

- **User Principal Name (UPN):** `userPrincipalName`
- **NetBios Name:** `sAMAccountName`
- **Email ID:** Directory attribute for that email
- **Multiple formats:** Retrieve the user mapping attributes from the user directory before enabling your User-ID sources.

If you do not specify a primary username, the firewall uses the following default values for each server profile type:

Attribute	Active Directory	Novell eDirectory or Sun ONE Directory Server
Primary Username	<code>sAMAccountName</code>	<code>uid</code>
E-Mail	<code>mail</code>	<code>mail</code>

Attribute	Active Directory	Novell eDirectory or Sun ONE Directory Server
Alternate Username 1	userPrincipalName	None.
Group Name	name	cn
Group Member	member	member

3. (Optional) Specify an **E-Mail** address format and up to three **Alternate Username** formats.
4. Select **Device > User Identification > Group Mapping Settings > Add > User and Group Attributes > Group Attributes** and specify the **Group Name**, **Group Member**, and **E-Mail** address formats.

You must commit before the firewall collects the directory attributes from the LDAP server.

STEP 4 | Limit which groups will be available in policy rules.

Required only if you want to limit policy rules to specific groups. The combined maximum for the **Group Include List** and **Custom Group** list is 640 entries per group mapping configuration. Each entry can be a single group or a list of 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 (Configure an Authentication Profile and Sequence).

1. Add existing groups from the directory service:
 1. Select **Group Include List**.
 2. Select the Available Groups you want to appear in policy rules and add (⊕) them to the Included Groups.
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:
 1. Select **Custom Group** and **Add** the group.
 2. 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 (such as in policies and logs).

3. 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** to save your changes.

You must commit before custom groups will be available in policies and objects.

STEP 5 | Commit your changes.

You must commit before you can use custom groups in policies and objects and before the firewall can collect the attributes from the LDAP server.



After configuring the firewall to retrieve group mapping information from an LDAP server, but before configuring policies based on the groups it retrieves, the best practice is to either wait for the firewall to refresh its group mappings cache or refresh the cache

*manually. To verify which groups you can currently use in policies, access the firewall CLI and run the **show user group** command. To determine when the firewall will next refresh the group mappings cache, run the **show user group-mapping statistics** command and check the *Next Action*. To manually refresh the cache, run the **debug user-id refresh group-mapping all** command.*

STEP 6 | Verify that the user and group mapping has correctly identified users.

1. Select **Device > User Identification > Group Mapping > Group Include List** to confirm the firewall has fetched all of the groups.
2. To verify that all of the user attributes have been correctly captured, use the following CLI command:

```
show user user-attributes user all
```

The normalized format for the User Principal Name (UPN), primary username, email attributes, and any configured alternate usernames display for all users:

```
admin@PA-VM-8.1> show user user-attributes user all
```

```
Primary: nam\sam-user    Email: sam-user@nam.com
```

```
Alt User Names:1) nam.com\sam-user
```

```
2) nam\sam-user-upn
```

```
3) sam-user-upn@nam.local
```

```
4) sam-user@nam.com
```

3. Verify that the usernames are correctly displayed in the **Source User** column under **Monitor > Logs > Traffic**.

Relative Time	Type	From Zone	To Zone	Source	Source User	Destination	To Port	Application	Action	Rule	Session Rule Reason	Action
02/11/2018 08:00:02	end	Trust	Trust		user@pan.com		120	ssh/ftp	allow	allow decrypted traffic	agent not	240
02/11/2018 08:00:07	end	Trust	Untrust		user@pan.com		443	ssl	allow	allow decrypted traffic	exp-fs	08:00
02/11/2018 08:01:21	start	Trust	Untrust		user@pan.com		443	ssl	allow	allow decrypted traffic	n/a	402
02/11/2018 08:01:23	end	Trust	Untrust		user@pan.com		80	web-browsing	allow	allow decrypted traffic	exp-fs	1:04
02/11/2018 08:01:26	end	Trust	Untrust		user@pan.com		80	web-browsing	allow	allow decrypted traffic	exp-fs	1:04
02/11/2018 08:07:06	end	Trust	Trust		user@pan.com		120	ssh/ftp	allow	allow decrypted traffic	agent not	240
02/11/2018 08:08:07	end	Trust	Untrust		user@pan.com		443	ssl	allow	allow decrypted traffic	exp-fs	6:36
02/11/2018 08:08:52	start	Trust	Untrust		user@pan.com		443	ssl	allow	allow decrypted traffic	n/a	402
02/11/2018 08:09:29	end	Trust	Untrust		user@pan.com		443	non-updates	allow	allow decrypted traffic	exp-fs	6:36
02/11/2018 08:09:29	end	Trust	Untrust		user@pan.com		443	non-updates	allow	allow decrypted traffic	exp-fs	6:36
02/11/2018 08:09:34	start	Trust	Untrust		user@pan.com		443	non-updates	allow	allow decrypted traffic	n/a	402
02/11/2018 08:09:34	start	Trust	Untrust		user@pan.com		443	ssl	allow	allow decrypted traffic	n/a	402
02/11/2018 08:09:43	start	Trust	Untrust		user@pan.com		443	non-updates	allow	allow decrypted traffic	n/a	402
02/11/2018 08:09:43	start	Trust	Untrust		user@pan.com		443	ssl	allow	allow decrypted traffic	n/a	402
02/11/2018 08:09:58	end	Trust	Untrust		user@pan.com		443	ssl	allow	allow decrypted traffic	exp-fs from check	08:30
02/11/2018 08:09:58	start	Trust	Untrust		user@pan.com		443	ssl	allow	allow decrypted traffic	n/a	402
02/11/2018 08:10:00	end	Trust	Trust		user@pan.com		120	ssh/ftp	allow	allow decrypted traffic	agent not	240

4. Verify that the users are mapped to the correct usernames in the **User Provided by Source** column under **Monitor > Logs > User-ID**.

Relative Time	IP	User	Username	Data Source	Source Name	Source Type	Factor Type	Factor Completion Time	Factor Number	User Provided by Source
02/11/2018 08:04:12		user@pan.com	user@pan.com	logfire-portal		authenticating		2018/10/10 10:04:12	1	user@pan.com
02/11/2018 08:04:13		user@pan.com	user@pan.com	logfire-portal		authenticating		2018/10/10 10:04:13	1	user@pan.com
02/11/2018 08:04:13		user@pan.com	user@pan.com	logfire-portal		authenticating		2018/10/10 10:04:13	1	user@pan.com
02/11/2018 08:04:13		user@pan.com	user@pan.com	logfire-portal		authenticating		2018/10/10 10:04:13	1	user@pan.com
02/11/2018 08:04:13		user@pan.com	user@pan.com	logfire-portal		authenticating		2018/10/10 10:04:13	1	user@pan.com
02/11/2018 08:04:13		user@pan.com	user@pan.com	logfire-portal		authenticating		2018/10/10 10:04:13	1	user@pan.com

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.

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:

- ❑ 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](#). Using Captive Portal in conjunction with [Authentication Policy](#) also ensures that all users authenticate to access your most sensitive applications and data.
- ❑ 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 Monitor Syslog Senders for User Mapping](#).



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.

- ❑ To [Share User-ID Mappings Across Virtual Systems](#), you can configure a virtual system as a User-ID hub.
- ❑ 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

To use 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, 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 the user mapping methods and settings you plan to use. For example, if you are using the PAN-OS integrated User-ID agent, the service account requires Server Operator privileges to monitor user sessions. If you are using the Windows-based User-ID agent, the service account does not require Server Operator privileges to monitor user sessions. To reduce

the risk of compromising the User-ID service account, always configure the account with the minimum set of permissions necessary for the agent.

- If you are installing the Windows-based User-ID agent on a supported Windows server, [Configure a Service Account for the Windows User-ID Agent](#).
- If you are using the PAN-OS integrated User-ID agent on the firewall, [Configure a Service Account for the PAN-OS Integrated User-ID Agent](#).



User-ID provides many methods for safely collecting user mapping information. Some legacy features designed for environments that only required user mapping on Windows desktops attached to the local network require privileged service accounts. If the privileged service account is compromised, this would open your network to attack. As a best practice, avoid using legacy features that require privileges that would pose a threat if compromised, such as client probing, NTLM authentication, and session monitoring.

Configure a Service Account for the Windows User-ID Agent

Create a dedicated Active Directory service account for the Windows User-ID agent to access the services and hosts it will monitor to collect user mappings. You must create a service account in each domain the agent will monitor. After you enable the required permissions for the service account, [Configure User Mapping Using the Windows User-ID Agent](#).



The following workflow details all required privileges and provides guidance for the User-ID features which require privileges that could pose a threat so that you can decide how to best identify users without compromising your overall security posture.

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**, then click **Next** and **Finish**.

STEP 2 | Configure either local or group policy to allow the service account to log on as a service.


The permission to log on as a service is only needed locally on the Windows server that is the agent host.

- To assign permissions locally:
 1. select **Control Panel > Administrative Tools > Local Security Policy**.


2.



3. Select **Local Policies > User Rights Assignment > Log on as a service**.




4. Add User or Group to add the service account.




5. Enter the object names to select (the service account name) in **domain\username** format and click **OK**.




- To configure group policy if you are installing Windows User-ID agents on multiple servers, use the Group Policy Management Editor.
1. Select **Start > Group Policy Management > <your domain> > Default Domain Policy > Action > Edit** for the Windows server that is the agent host.



2. Select **Computer Configuration > Policies > Windows Settings > Security Settings > Local Policies > User Rights Assignment**.



2. Navigate to the Builtin folder for the domain, right-click the **Event Log Readers** group, and select **Add to a Group**.



3. Enter the name of the service account then click **Check Names** to validate that you have the proper object name.




4. Click **OK** twice to save the settings.
5. Confirm that the builtin Event Log Reader group lists the service account as a member.



STEP 4 | Assign account permissions to the installation folder to allow the service account to access the agent's installation folder to read the configuration and write logs.

You only need to perform this step if the service account you configured for the User-ID agent is not either a domain administrator or a local administrator on the User-ID agent server host.

1. From the Windows Explorer, navigate to **C:\Program Files(x86)\Palo Alto Networks**, right-click the folder, and select **Properties**.
2. On the **Security** tab, click **Edit**.




3. **Add** the User-ID agent service account and **Allow** permissions to **Modify**, **Read & execute**, **List folder contents**, **Read**, and **Write**, and then click **OK** to save the account settings.



If you do not want to configure individual permissions, you can Allow the Full Control permission instead.

STEP 5 | To allow the agent to make configuration changes (for example, if you select a different logging level), give the service account permissions to the User-ID agent registry sub-tree.

1. Select **Start > Run** and enter **regedt32** 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
2. Right-click the **Palo Alto Networks** node and select **Permissions**.



3. Assign the User-ID service account **Full Control** and then click **OK** to save the setting.




STEP 6 | Disable service account privileges that are not required.

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).
 1. Select **Group Policy Management Editor > Default Domain Policy > Computer Configuration > Policies > Windows Settings > Security Settings > User Rights Assignment**.
 2. For **Deny log on as a batch job**, **Deny log on locally**, and **Deny log on through Remote Desktop Services**, right-click **Properties**.
 3. Select **Define these policy settings > Add User or Group** and add the service account name, then click **OK**.



- **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.
 1. Select **Start > Run**, enter **MMC**, and select **File > Add/Remove Snap-in > Active Directory Users and Computers > Users**.
 2. Right-click the service account name, then select **Properties**.
 3. Select **Dial-in**, then **Deny** the **Network Access Permission**.



STEP 7 | As a next step, [Configure User Mapping Using the Windows User-ID Agent](#).

Configure a Service Account for the PAN-OS Integrated User-ID Agent

Create a dedicated Active Directory service account for the PAN-OS Integrated User-ID agent to access the services and hosts it will monitor to collect user mappings. You must create a service account in each domain the agent will monitor. After you enable the required permissions for the service account, [Configure User Mapping Using the PAN-OS Integrated User-ID Agent](#).



The following workflow details all required privileges and provides guidance for the User-ID features which require privileges that could pose a threat so that you can decide how to best identify users without compromising your overall security posture.


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**, then click **Next** and **Finish**.

STEP 2 | If you want to use [Server Monitoring](#) to identify users, add the service account to the Event Log Reader builtin group to allow the service account to read the security log events.

1. On the domain controller or Exchange server that contains the logs you want the User-ID agent to read, or on the member server that receives events from Windows log forwarding, select **Start > Run**, enter **MMC**, and select **File > Add/Remove Snap-in > Active Directory Users and Computers > Add**, then click **OK** to run the MMC and launch the Active Directory Users and Computers snap-in.



2. Navigate to the Builtin folder for the domain, right-click the **Event Log Readers** group, and select **Add to a group**.



3. Enter the name of the service account then click **Check Names** to validate that you have the proper object name.



4. Click **OK** twice to save the settings.
5. Confirm that the builtin Event Log Reader group lists the service account as a member.



STEP 3 | If you want to use **WMI** to collect user data, assign DCOM privileges to the service account so that it can use WMI queries on monitored servers.

1. Select **Active Directory Users and Computers** > <your domain> > **Builtin** > **Distributed COM Users**.
2. Right-click **Properties** > **Members** > **Add** and enter the service account name.

STEP 4 | If you plan to use [WMI](#), enable the account to read the CIMV2 namespace and assign the required permissions on the client systems to be probed.



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 **wmicmgmt.msc**, and launch the WMI Management Console.
2. In the console tree, right-click **WMI Control** and select **Properties**.



3. Select the **Security** tab, then select **Root > CIMV2**, and click the **Security** button.



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 **Remote Enable** permissions.



6. Click **OK** twice.
7. Use the Local Users and Groups MMC snap-in (lusrmgr.msc) to add the service account to the local Distributed Component Object Model (DCOM) Users and Remote Desktop Users groups on the system that will be probed.

STEP 5 | (Not Recommended) To allow the agent to monitor user sessions to poll Windows servers for user mapping information, assign Server Operator privileges to the service account.



Because this group also has privileges for shutting down and restarting servers, only assign the account to this group if monitoring user sessions is very important.

1. Select **Active Directory Users and Computers** > <your domain> > **Builtin** > **Server Operators Group**.
2. Right-click **Properties** > **Members** > **Add** add service account name

STEP 6 | If you want to configure **NTLM** authentication for **Captive Portal**, configure the firewall to join the domain.


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).

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.




*The PAN-OS integrated agent requires privileged operations to join the domain, which poses a security threat if the account is compromised. As a best practice, configure **Kerberos** single sign-on (SSO) or **SAML** SSO authentication for Captive Portal instead of NTLM. Kerberos and SAML are stronger, more secure authentication methods and do not require the firewall to join the domain.*

1. Select **Start** > **Run**, enter **MMC**, and select **File** > **Add/Remove Snap-in** > **Active Directory Users and Computers** > **Users**.
2. Right-click the domain and select **Delegate Control**.



3. Click **Next**, then **Add** the service account name and click **OK**.
4. Click **Next**, then **Join a computer to the domain**.



5. Click **Next**, verify the service account information, then **Finish**.


STEP 7 | Disable service account privileges that are not required.

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).
 1. Select **Group Policy Management Editor > Default Domain Policy > Computer Configuration > Policies > Windows Settings > Security Settings > User Rights Assignment**.

2. For **Deny log on as a batch job**, **Deny log on locally**, and **Deny log on through Remote Desktop Services**, right-click **Properties**, then select **Define these policy settings > Add User or Group** and add the service account name, then click **OK**.



- **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.
 1. **Start > Run**, enter **MMC**, and select **File > Add/Remove Snap-in > Active Directory Users and Computers > Users**.
 2. Right-click the service account name, then select **Properties**.
 3. Select **Dial-in**, then **Deny** the **Network Access Permission**.



STEP 8 | As a next step, [Configure User Mapping Using the PAN-OS Integrated User-ID Agent](#).

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 Windows-Based User-ID Agent](#)
- [Configure the Windows User-ID Agent for User Mapping](#)

Install the Windows-Based 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 the [Palo Alto Networks Compatibility Matrix](#).

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](#) and grant the necessary permissions for the Windows User-ID agent.

1. Enable the service account to log on as a service by configuring either local or group policy.
 1. To configure the group policy if you are installing Windows-based User-ID agents on multiple servers, select **Group Policy Management > Default Domain Policy > Computer Configuration > Policies > Windows Settings > Security Settings > Local Policies > User Rights Assignment** for the Windows server that is the agent host.
 2. Right-click **Log on as a service**, then select **Properties**.
 3. Add the service account username or builtin group (Administrators have this privilege by default).




The permission to log on as a service is only needed locally on the Windows server that is the agent host. If you are using only one User-ID agent, you can grant the permissions locally on the agent host using the following instructions.


1. To assign permissions locally, select **Control Panel > Administrative Tools > Local Security Policy**.



2. Select Local Policies > User Rights Assignment > Log on as a service.



3. Add User or Group to add the service account.




4. Enter the service account name in **domain\username** format in the **Enter the object names to select** entry field and click **OK**.



To confirm the service account name is valid, **Check Names**.

2. If you want to use [server monitoring](#) to identify users, add the service account to the Event Log Reader builtin group to enable privileges for reading the security log events.
 1. On the domain controller or Exchange server that contains the logs you want the User-ID agent to read, or on the member server that receives events from Windows log forwarding, run the MMC and launch the Active Directory Users and Computers snap-in.
 2. Navigate to the Builtin folder for the domain, right-click the **Event Log Reader** group and select **Add to Group** to open the properties dialog.
 3. 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.
 4. Click **OK** twice to save the settings.
 5. Confirm that the builtin Event Log Reader group lists the service account as a member.



3. Assign account permissions to the installation folder to allow the service account to access the agent's installation folder to read the configuration and write logs.

You only need to perform this step if the service account you configured for the User-ID agent is not either a domain administrator or a local administrator on the User-ID agent server host.

1. From the Windows Explorer, navigate to **C:\Program Files(x86)\Palo Alto Networks** for 32-bit systems, right-click the folder, and select **Properties**.
2. On the **Security** tab, click **Edit**.



3. **Add** the User-ID agent service account and assign it permissions to **Modify**, **Read & execute**, **List folder contents**, **Read**, and **Write**, and then click **OK** to save the account settings.



If you want to allow the service account to access the User-ID agent's registry keys, Allow the Full Control permission.

4. Give the service account permissions to the User-ID Agent registry sub-tree:
 1. Run **regedt32** 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
 2. Right-click the Palo Alto Networks node and select **Permissions**.
 3. Assign the User-ID service account **Full Control** and then click **OK** to save the setting.

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). During the initial connection, the agent transfers the most recent 50,000 events from the log to map users. On each subsequent connection, the agent transfers events with a timestamp later than the last communication with the domain controller. Therefore, always install one or more User-ID agents at each site that has servers to be monitored.

- 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 [Compatibility Matrix](#).
- 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 domain controllers that process authentication for users you want to map. You might need to install multiple User-ID agents to efficiently monitor all of your resources.
- If you are using the User-ID agent for credential detection, you must install it on the read-only domain controller (RODC). As a best practice deploy a separate agent for this purpose. Do not use the User-ID agent installed on the RODC to map IP addresses to users. The User-ID agent installer for credential detection is named UaCredInstall64-x.x.x.msi.

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.

1. Log in to the [Palo Alto Networks Customer Support Portal](#).
2. Select **Updates > Software Updates**.
3. Set **Filter By** to **User Identification Agent** and select the version of the User-ID agent you want to install from the corresponding Download column. For example, to download the 9.0 version of the User-ID agent, select **UaInstall-9.0.0-0.msi**.

If you are using the User-ID agent for [credential detection](#), make sure you download the **UaCredInstall164-x.x.x.msi** file instead of the regular User-ID installation file, which is named **UaInstall-x.x.x.msi**.

4. Save the **UaCredInstall164-x.x.x-xx.msi** or **UaInstall-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 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, 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 **C:\Program Files(x86)\Palo Alto Networks** for 32-bit systems or **C:\Program Files\Palo Alto Networks** for 64-bit systems, 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 as an administrator.

Open the Windows **Start** menu, right-click the **User-ID Agent** program, and select **Run as administrator**.



You must run the User-ID Agent application as an administrator to install the application, commit configuration changes, or uninstall the application.

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. To change the account to a restricted account:

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.
4. **Commit** the changes to the User-ID agent configuration to restart the service using the service account credentials.

STEP 7 | (Optional) Assign your own certificates for mutual authentication between the Windows User-ID agent and the firewall.

1. Obtain your certificate for the Windows User-ID agent using one of the following methods. Upload the server certificate in Privacy Enhanced Mail (PEM) format and the server certificate's encrypted key.
 - [Generate a Certificate](#) and export it for upload to the Windows User-ID agent.
 - Export a certificate from your enterprise certificate authority (CA) and the upload it to the Windows User-ID agent.
2. Add a server certificate to Windows User-ID agent.
 1. On the Windows User-ID agent, select **Server Certificate** and click **Add**.
 2. Enter the path and name of the certificate file received from the CA or browse to the certificate file.
 3. Enter the private key passphrase.
 4. Click **OK** and then **Commit**.
3. Upload a certificate to the firewall to validate the Windows User-ID agent's identity.
4. Configure the certificate profile for the client device (firewall or Panorama).
 1. Select **Device > Certificate Management > Certificate Profile**.
 2. [Configure a Certificate Profile](#).



You can only assign one certificate profile for Windows User-ID agents and Terminal Services (TS) agents. Therefore, your certificate profile must include all certificate authorities that issued certificates uploaded to connected User-ID and TS agents.

5. Assign the certificate profile on the firewall.
 1. Select **Device > User Identification > Connection Security** and click the edit button.
 2. Select the **User-ID Certificate Profile** you configured in the previous step.
 3. Click **OK**.
6. **Commit** your changes.

STEP 8 | [Configure Credential Detection with the Windows-based User-ID Agent](#).

To use the Windows-based User-ID agent to detect credential submissions and [Prevent Credential Phishing](#), you must install the User-ID credential service on the Windows-based User-ID agent. You can only install this add-on on a read-only domain controller (RODC).

Configure the Windows 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](#).

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 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.

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.

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.

Save the ignore-user list as a text document on the agent host using the title `ignore_user_list` and use the `.txt` file extension to 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`. You can also use the ignore-user list to identify users whom you want to force to authenticate using Captive Portal.

STEP 7 | Configure the firewall to connect to the User-ID agent.



The firewall can connect to only one Windows-based User-ID agent that is using the User-ID credential service add-on to detect corporate credential submissions. See [Configure Credential Detection with the Windows-based User-ID Agent](#) for more details on how to use this service for credential phishing prevention.

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).

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 describes 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).

STEP 1 | Create an Active Directory service account for the User-ID agent to access the services and hosts that the firewall 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 that the firewall can monitor the Security log files on all servers that contain login events.

1. Select **Device > User Identification > User Mapping**.
2. **Add** a server (Server Monitoring section).
3. Enter a **Name** to identify the server.
4. Select the **Type** of server.
 - **Microsoft Active Directory**
 - **Microsoft Exchange**
 - **Novell eDirectory**
 - **Syslog Sender**
5. (**Microsoft Active Directory and Microsoft Exchange only**) Select the **Transport Protocol** you want to use to monitor security logs and session information on the server.
 - **WMI**—The firewall and the monitored servers use Windows Management Instrumentation ([WMI](#)) to communicate.
 - **WinRM-HTTP**—The firewall and the monitored servers use Kerberos for mutual authentication and the monitored server encrypts the communication with the firewall using a negotiated Kerberos session key.
 - **WinRM-HTTPS**—The firewall and the monitored servers use HTTPS to communicate and use basic authentication or Kerberos for mutual authentication.

If you select a Windows Remote Management (WinRM) option, you must [Configure Server Monitoring Using WinRM](#).
6. (**Microsoft Active Directory, Microsoft Exchange, and Novell eDirectory only**) Enter the **Network Address** of the server.



If you are using [WinRM with Kerberos](#), you must enter a fully qualified domain name (FQDN). If you want to use [WinRM with basic authentication](#) or use WMI to monitor the server, you can enter an IP address or FQDN.

*To monitor servers using WMI, specify an IP address, the service account name (if all server monitoring is in the same domain), or a fully qualified domain name (FQDN). If you specify an FQDN, use the down-level logon name in the (DLN)\sAMAccountName format instead of the FQDN\sAMAccountName format. For example, use **example***

`\user.services not example.com\user.services`. If you specify an FQDN, the firewall will attempt to authenticate using Kerberos, which does not support WMI.

7. (Syslog Sender only) If you select **Syslog Sender** as the server **Type**, [Configure the PAN-OS Integrated User-ID Agent as a Syslog Listener](#).
8. (Novell eDirectory only) Make sure the **Server Profile** you select is **Enabled** and click **OK**.
9. (Optional) Configure 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.

STEP 3 | (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 domain controller is processing many requests, delays between queries may exceed the specified value.

1. **Edit** the Palo Alto Networks User ID Agent Setup.
2. Select the **Server Monitor** tab and specify the **Server Log Monitor Frequency** in seconds (range is 1 to 3,600; default is 2). In environments with older domain controllers or high-latency links, set this frequency to a minimum of five seconds.



*Ensure that the **Enable Session** setting is not enabled. 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 network access control (NAC) devices.*

3. Click **OK** to save your changes.

STEP 4 | Specify the subnetworks that 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, which networks to exclude from User-ID to ensure that the agent is communicating only with internal resources and to prevent unauthorized users from being mapped. You should enable user mapping only 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. Ensure that the entry is **Enabled**.
3. Enter the **Network Address** and then select whether to include or exclude it:
 - **Include**—Select this option to limit user mapping to only users logged in to the specified subnetwork. 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 to configure 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 5 | Set the domain credentials for the account that 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 **Server Monitor Account** tab and enter the **User Name** and **Password** for the [service account](#) that the User-ID agent will use to probe the clients and monitor servers. Enter the username using the **domain\username** syntax.
3. If you are using WinRM to monitor servers, configure the firewall to authenticate with the server you are monitoring.
 - If you want to use [WinRM with basic authentication](#), enable WinRM on the server, configure basic authentication, and specify the service account **Domain's DNS Name**.
 - If you want to use [WinRM with Kerberos](#), [Configure a Kerberos server profile](#) if you have not already done so and then select the **Kerberos Server Profile**.

STEP 6 | (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. On the **Client Probing** tab, **Enable Probing**.
2. (Optional) Specify the **Probe Interval** to define the interval (in minutes) between the end of the last probe request and the start of the next request.

If necessary, increase the value to ensure the User-ID agent has sufficient time to probe all the learned IP addresses (range is 1 to 1440; default is 20).



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.

STEP 7 | (Optional) Define the set of user accounts that don't require IP address-to-username mappings, such as kiosk accounts.



Define the ignore user list on the firewall that is the User-ID agent, not the client. If you define the ignore user list on the client firewall, the users in the list are still mapped during redistribution.

On the **Ignore User List** tab, **Add** each username you want to exclude from user mapping. You can also use the ignore user list to identify the users you want to force to use Captive Portal to authenticate. 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 8 | Activate your configuration changes.

Click **OK** and **Commit**.

STEP 9 | 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 Server Monitoring Using WinRM

You can [configure the PAN-OS integrated User-ID agent](#) to monitor servers using Windows Remote Management (WinRM). Using the WinRM protocol improves speed, efficiency, and security when monitoring server events to map user events to IP addresses. The PAN-OS integrated User-ID agent supports the WinRM protocol on Windows Server 2008 Active Directory and Microsoft Exchange Server 2008 or later versions of both.

There are three ways to configure server monitoring using WinRM:

- [Configure WinRM over HTTPS with Basic Authentication](#)—The firewall authenticates to the monitored server using the username and password of the service account for the User-ID agent and the firewall authenticates the monitored server using the User-ID certificate profile.
- [Configure WinRM over HTTP with Kerberos](#)—The firewall and the monitored servers use Kerberos for mutual authentication and the monitored server encrypts the communication with the firewall using a negotiated Kerberos session key.
- [Configure WinRM over HTTPS with Kerberos](#)—The firewall and the monitored server use HTTPS to communicate and use Kerberos for mutual authentication.

Configure WinRM over HTTPS with Basic Authentication

When you configure WinRM to use HTTPS with basic authentication, the firewall transfers the credentials for the service account in a secure tunnel using SSL.

STEP 1 | Configure the [service account](#) with Remote Management User and CIMV2 privileges for the server you want to monitor.

STEP 2 | On the Windows server you are monitoring, obtain the thumbprint from the certificate for the Windows server to use with WinRM and enable WinRM.



The account you use to configure WinRM on the server you want to monitor must have administrator privileges.

1. Verify the certificate is installed in the Local Computer certificate store (**Certificates (Local Computer) > Personal > Certificates**).
If you do not see the Local Computer certificate store, launch the Microsoft Management Console (**Start > Run > MMC**) and add the Certificates snap-in (**File > Add/Remove Snap-in > Certificates > Add > Computer account > Next > Finish**).
2. Open the certificate and select **General > Details > Show: <All>**.
3. Select the **Thumbprint** and copy it.

4. To enable the firewall to connect to the Windows server using WinRM, enter the following command: **winrm quickconfig**.
5. Enter **y** to confirm the changes and then confirm the output displays **winrm service started**.

If WinRM is enabled, the output displays **winrm service is already running on this machine**. You will be prompted to confirm any additional required configuration changes.

6. To verify that WinRM is communicating using HTTPS, enter the following command: **winrm enumerate winrm/config/listener** and confirm that the output displays **Transport = HTTPS**.

By default, WinRM/HTTPS uses port 5986.

7. From the Windows server command prompt, enter the following command:
winrm create winrm/config/Listener?Address=*&Transport=HTTPS @{Hostname="<hostname>";CertificateThumbprint="Certificate Thumbprint"} , where *hostname* is the hostname of the Windows server and *Certificate Thumbprint* is the value you copied from the certificate.



Use the command prompt (not Powershell) and remove any spaces in the Certificate Thumbprint to ensure that WinRM can validate the certificate.

8. From the Windows server command prompt, enter the following command:

```
c:\> winrm set winrm/config/client/auth @{Basic="true"}
```

9. Enter the following command: **winrm get winrm/config/service/Auth** and confirm that **Basic = true**.

STEP 3 | Enable Basic Authentication between the PAN-OS integrated User-ID agent and the monitored servers.

1. Select **Device > User Identification > User Mapping > Palo Alto Networks User-ID Agent Setup > Server Monitor Account**.
2. In **domain\username** format, enter the **User Name** for the service account that the User-ID agent will use to monitor servers.
3. Enter the **Domain's DNS Name** of the server monitor account.
4. Enter the **Password** and **Confirm Password** for the service account.

Palo Alto Networks User-ID Agent Setup

Server Monitor Account | Server Monitor | Client Probing | Cache | NTLM | Redistribution | Syslog Filters | Ignore User List

User Name: [text field]

Domain's DNS Name: [text field]

Password: [password field]

Confirm Password: [password field]

Kerberos Server Profile: None [dropdown]

OK Cancel

5. Click **OK**.

STEP 4 | Configure [server monitoring](#) for the PAN-OS integrated User-ID agent.

1. Select the Microsoft server **Type** (**Microsoft Active Directory** or **Microsoft Exchange**).
2. Select **Win-RM-HTTPS** as the **Transport Protocol** to use Windows Remote Management (WinRM) over HTTPS to monitor the server security logs and session information.
3. Enter the IP address or FQDN **Network Address** of the server.



STEP 5 | To enable the PAN-OS integrated User-ID agent to communicate with the monitored servers using WinRM-HTTPS, verify that you successfully imported the root certificate for the service certificates that the Windows server uses for WinRM on to the firewall and associate the certificate with the User-ID Certificate Profile.

1. Select **Device > User Identification > Connection Security**.
2. Click **Edit**.
3. Select the Windows server certificate to use for the **User-ID Certificate Profile**.



4. Click **OK**

STEP 6 | **Commit** your changes.

STEP 7 | Verify that the status of each monitored server is Connected (**Device > User Identification > User Mapping**).

Configure WinRM over HTTP with Kerberos

When you configure WinRM over HTTP with Kerberos, the firewall and the monitored servers use Kerberos for mutual authentication and the monitored server encrypts the communication with the firewall using a negotiated Kerberos session key.



WinRM with Kerberos supports the aes128-cts-hmac-sha1-96 and aes256-cts-hmac-sha1-96 ciphers. If the server you want to monitor uses RC4, you must download the Windows update and disable RC4 for Kerberos in the registry settings of the server you want to monitor.

STEP 1 | Configure the service account with Remote Management User and CIMV2 privileges for the server you want to monitor.

STEP 2 | Confirm that WinRM is enabled on the Windows server you are monitoring.



The account you use to configure WinRM on the server you want to monitor must have administrator privileges.

1. To enable the firewall to connect to the Windows server using WinRM, enter the following command: **winrm quickconfig**.
2. Enter **y** to confirm the changes and then confirm the output displays **winrm service started**.

If WinRM is enabled, the output displays `winRM service is already running on this machine`. You will be prompted to confirm any additional required configuration changes.

3. To verify that WinRM is communicating using HTTP, enter the following command: **winrm enumerate winrm/config/listener** and confirm that the output displays `Transport = HTTP`.

By default, WinRM/HTTP uses port 5985.


4. Enter the following command: **winrm get winrm/config/service/Auth** and confirm that `Kerberos = true`.

STEP 3 | Enable the PAN-OS integrated User-ID agent and the monitored servers to authenticate using Kerberos.

1. If you did not do so during the [initial configuration](#), configure date and time (NTP) settings to ensure successful Kerberos negotiation.
2. [Configure a Kerberos server profile](#) on the firewall to authenticate with the server to monitor the security logs and session information.
3. Select **Device > User Identification > User Mapping > Palo Alto Networks User-ID Agent Setup > Server Monitor Account**.
4. In **domain\username** format, enter the **User Name** for the service account that the User-ID agent will use to monitor servers.
5. Enter the **Domain's DNS Name** of the server monitor account.

Kerberos uses the domain name to locate the service account.

6. Enter the **Password** and **Confirm Password** for the service account.
7. Select the **Kerberos Server Profile** you configured in Step 3.2.

The screenshot shows the 'Palo Alto Networks User-ID Agent Setup' window with the 'Server Monitor Account' tab selected. The window has several tabs: 'Server Monitor Account', 'Server Monitor', 'Client Probing', 'Cache', 'NTLM', 'Redistribution', 'Syslog Filters', and 'Ignore User List'. The 'Server Monitor Account' tab contains the following fields: 'User Name' (text input), 'Domain's DNS Name' (text input), 'Password' (password input), 'Confirm Password' (password input), and 'Kerberos Server Profile' (dropdown menu). The 'Kerberos Server Profile' dropdown is currently set to 'None'. At the bottom right, there are 'OK' and 'Cancel' buttons.

8. Click **OK**.

STEP 4 | Configure [server monitoring](#) for the PAN-OS integrated User-ID agent.

1. Configure the Microsoft server type (**Microsoft Active Directory** or **Microsoft Exchange**).
2. Select **WinRM-HTTP** as the **Transport Protocol** to use Windows Remote Management (WinRM) over HTTP to monitor the server security logs and session information.
3. Enter the FQDN **Network Address** of the server.

If you are using Kerberos, the network address must be a fully qualified domain name (FQDN).



STEP 5 | Commit your changes.

STEP 6 | Verify that the status of each monitored server is Connected (**Device > User Identification > User Mapping**).

Configure WinRM over HTTPS with Kerberos

When you configure WinRM over HTTPS with Kerberos, the firewall and the monitored server use HTTPS to communicate and use Kerberos for mutual authentication.



WinRM with Kerberos supports the `aes128-cts-hmac-sha1-96` and `aes256-cts-hmac-sha1-96` ciphers. If the server you want to monitor uses RC4, you must download the Windows update and disable RC4 for Kerberos in the registry settings of the server you want to monitor.

STEP 1 | Configure the [service account](#) with Remote Management User and CIMV2 privileges for the server you want to monitor.

STEP 2 | On the Windows server you are monitoring, obtain the thumbprint from the certificate for the Windows server to use with WinRM and enable WinRM.



The account you use to configure WinRM on the server you want to monitor must have administrator privileges.

1. Verify the certificate is installed in the Local Computer certificate store (**Certificates (Local Computer) > Personal > Certificates**).
If you do not see the Local Computer certificate store, launch the Microsoft Management Console (**Start > Run > MMC**) and add the Certificates snap-in (**File > Add/Remove Snap-in > Certificates > Add > Computer account > Next > Finish**).
2. Open the certificate and select **General > Details > Show: <All>**.
3. Select the **Thumbprint** and copy it.
4. To enable the firewall to connect to the Windows server using WinRM, enter the following command: **winrm quickconfig**.
5. Enter **y** to confirm the changes and then confirm the output displays `winrm service started`.
If WinRM is enabled, the output displays `winrm service is already running on this machine`. You will be prompted to confirm any additional required configuration changes.
6. To verify that WinRM is communicating using HTTPS, enter the following command: **winrm enumerate winrm/config/listener**. Then confirm that the output displays `Transport = HTTPS`.

By default, WinRM/HTTPS uses 5986.

7. From the Windows server command prompt, enter the following command:
**winrm create winrm/config/Listener?Address=*&Transport=HTTPS
@{Hostname="<hostname>";CertificateThumbprint="Certificate Thumbprint"};**
where *hostname* is the hostname of the Windows server and *Certificate Thumbprint* is the value you copied from the certificate.



Use the command prompt (not Powershell) and remove any spaces in the Certificate Thumbprint to ensure that WinRM can validate the certificate.

8. Enter the following command: **winrm get winrm/config/service/Auth** and confirm that Basic = false and Kerberos= true.

STEP 3 | Enable the PAN-OS integrated User-ID agent and the monitored servers to authenticate using Kerberos.

1. If you did not do so during the [initial configuration](#), configure date and time (NTP) settings to ensure successful Kerberos negotiation.
2. [Configure a Kerberos server profile](#) on the firewall to authenticate with the server to monitor the security logs and session information.
3. Select **Device > User Identification > User Mapping > Palo Alto Networks User-ID Agent Setup > Server Monitor Account**.
4. In **domain\username** format, enter the **User Name** for the service account that the User-ID agent will use to monitor servers.
5. Enter the **Domain's DNS Name** of the server monitor account.

Kerberos uses the domain name to locate the service account.

6. Enter the **Password** and **Confirm Password** for the service account.
7. Select the **Kerberos Server Profile** you created in Step 3.2.

Palo Alto Networks User-ID Agent Setup

Server Monitor Account | Server Monitor | Client Probing | Cache | NTLM | Redistribution | Syslog Filters | Ignore User List

User Name: _____

Domain's DNS Name: _____

Password: _____

Confirm Password: _____

Kerberos Server Profile: None

OK Cancel

8. Click **OK**.

STEP 4 | Configure [server monitoring](#) for the PAN-OS integrated User-ID agent.

1. Configure the Microsoft server type (**Microsoft Active Directory** or **Microsoft Exchange**).
2. Select **Win-RM-HTTPS** as the **Transport Protocol** to use Windows Remote Management (WinRM) over HTTPS to monitor the server security logs and session information.
3. Enter the FQDN **Network Address** of the server.

If you are using Kerberos, the network address must be a fully qualified domain name (FQDN).



STEP 5 | To enable the PAN-OS integrated User-ID agent to communicate with the monitored servers using WinRM-HTTPS, verify that you successfully imported the root certificate for the service certificates that the Windows server uses for WinRM on to the firewall and associate the certificate with the User-ID Certificate Profile.

The firewall uses the same certificate to authenticate with all monitored servers.

1. Select **Device > User Identification > Connection Security**.
2. Click **Edit**.
3. Select the Windows server certificate to use for the **User-ID Certificate Profile**.



4. Click **OK**.
5. **Commit** your changes.

STEP 6 | Verify that the status of each monitored server is Connected (**Device > User Identification > User Mapping**).

Configure User-ID to Monitor Syslog Senders for User Mapping

To obtain IP address-to-username mappings from existing network services that authenticate users, you can configure the PAN-OS integrated User-ID agent or Windows-based User-ID agent to parse [Syslog](#) messages from those services. To keep user mappings up to date, you can also configure the User-ID agent to parse syslog messages for logout events so that the firewall automatically deletes outdated mappings.

- [Configure the PAN-OS Integrated User-ID Agent as a Syslog Listener](#)
- [Configure the Windows User-ID Agent as a Syslog Listener](#)

Configure the PAN-OS Integrated User-ID Agent as a Syslog Listener

To configure the PAN-OS Integrated User-ID agent to create new user mappings and remove outdated mappings through syslog monitoring, start by defining Syslog Parse profiles. The User-ID agent uses the profiles to find login and logout events in syslog messages. In environments where *syslog senders* (the network services that authenticate users) deliver syslog messages in different formats, configure a profile for each syslog format. Syslog messages must meet certain criteria for a User-ID agent to parse them (see [Syslog](#)). This procedure uses examples with the following formats:

- **Login events**—[Tue Jul 5 13:15:04 2016 CDT] Administratorauthentication success
User:johndoe1 Source:192.168.3.212

- **Logout events**—[Tue Jul 5 13:18:05 2016CDT] User logout successful
User:johndoe1 Source:192.168.3.212

After configuring the Syslog Parse profiles, you specify syslog senders for the User-ID agent to monitor.

STEP 1 | Determine whether there is a predefined Syslog Parse profile for your particular syslog senders.

Palo Alto Networks provides several predefined profiles through Application content updates. The predefined profiles are global to the firewall, whereas custom profiles apply to a single virtual system only.



Any new Syslog Parse profiles in a given content release is documented in the corresponding release note along with the specific regex used to define the filter.

1. Install the latest Applications or Applications and Threats update:
 1. Select **Device > Dynamic Updates** and **Check Now**.
 2. **Download** and **Install** any new update.
2. Determine which predefined Syslog Parse profiles are available:
 1. Select **Device > User Identification > User Mapping** and click **Add** in the Server Monitoring section.
 2. Set the **Type** to **Syslog Sender** and click **Add** in the Filter section. If the Syslog Parse profile you need is available, skip the steps for defining custom profiles.

STEP 2 | Define custom Syslog Parse profiles to create and delete user mappings.

Each profile filters syslog messages to identify either login events (to create user mappings) or logout events (to delete mappings), but no single profile can do both.

1. Review the syslog messages that the syslog sender generates to identify the syntax for login and logout events. This enables you to define the matching patterns when creating Syslog Parse profiles.



While reviewing syslog messages, also determine whether they include the domain name. If they don't, and your user mappings require domain names, enter the Default Domain Name when defining the syslog senders that the User-ID agent monitors (later in this procedure).

2. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User-ID Agent Setup.
3. Select **Syslog Filters** and **Add** a Syslog Parse profile.
4. Enter a name to identify the **Syslog Parse Profile**.
5. Select the **Type** of parsing to find login or logout events in syslog messages:
 - **Regex Identifier**—Regular expressions.
 - **Field Identifier**—Text strings.

The following steps describe how to configure these parsing types.

STEP 3 | (Regex Identifier parsing only) Define the regex matching patterns.



*If the syslog message contains a standalone space or tab as a delimiter, use **\s** for a space and **\t** for a tab.*

1. Enter the **Event Regex** for the type of events you want to find:
 - **Login events**—For the example message, the regex **(authentication\ success){1}** extracts the first **{1}** instance of the string authenticationsuccess.

- **Logout events**—For the example message, the regex (**logout\ successful**){1} extracts the first {1} instance of the string logoutsuccessful.

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.

2. Enter the **Username Regex** to identify the start of the username.

In the example message, the regex **User: ([a-zA-Z0-9\\._]+)** matches the string **User: johndoe1** and identifies johndoe1 as the username.

3. Enter the **Address Regex** to identify the IP address portion of syslog messages.

In the example 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**.

The following is an example of a completed Syslog Parse profile that uses regex to identify login events:

4. Click **OK** twice to save the profile.

STEP 4 | (Field Identifier parsing only) Define string matching patterns.

1. Enter an **Event String** to identify the type of events you want to find.

- **Login events**—For the example message, the string **authentication success** identifies login events.
- **Logout events**—For the example message, the string **logoutsuccessful** identifies logout events.

2. Enter a **Username Prefix** 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 example messages, **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. Enter an **Address Prefix** 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 example messages, **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.

The following is an example of a completed Syslog Parse profile that uses string matching to identify login events:

6. Click **OK** twice to save the profile.

STEP 5 | Specify the syslog senders that the firewall monitors.

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 discards any syslog messages received from senders that are not on this list.

1. Select **Device > User Identification > User Mapping** and **Add** an entry to the Server Monitoring list.
2. Enter a **Name** to identify the sender.
3. Make sure the sender profile is **Enabled** (default is enabled).
4. Set the **Type** to **Syslog Sender**.
5. Enter the **Network Address** (IP address) of the syslog sender.
6. Select **SSL** (default) or **UDP** as the **Connection Type**.



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 sender. 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.



Always use SSL to listen for syslog messages because the traffic is encrypted (UDP sends the traffic in cleartext). If you must use UDP, make sure that the syslog sender and client are both on a dedicated, secure network to prevent untrusted hosts from sending UDP traffic to the firewall.

A syslog sender using SSL to connect will show a Status of Connected only when there is an active SSL connection. Syslog senders using UDP will not show a Status value.

7. For each syslog format that the sender supports, **Add** a Syslog Parse profile to the Filter list. Select the **Event Type** that each profile is configured to identify: **login** (default) or **logout**.
8. (Optional) If the syslog messages don't contain domain information and your user mappings require domain names, enter a **Default Domain Name** to append to the mappings.
9. Click **OK** to save the settings.

STEP 6 | Enable syslog listener services on the interface that the firewall uses to collect user mappings.

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** or **User-ID Syslog Listener-UDP** or both, based on the protocols you defined for the syslog senders in the Server Monitoring list.



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 only accepts syslog connections from senders that have a corresponding entry in the User-ID monitored servers configuration. The firewall discards connections or messages from senders that are not on the list.

4. Assign the Interface Management profile to the interface that the firewall uses to collect user mappings:
 1. Select **Network > Interfaces** and edit the interface.
 2. Select **Advanced > Other info**, select the Interface **Management Profile** you just added, and click **OK**.
5. **Commit** your changes.

STEP 7 | Verify that the firewall adds and deletes user mappings when users log in and out.



You can use [CLI commands](#) to see additional information about syslog senders, syslog messages, and user mappings.

1. Log in to a client system for which a monitored syslog sender generates login and logout event messages.
2. [Log in to the firewall CLI](#).
3. Verify that the firewall mapped the login username to the client IP address:

```
> show user ip-user-mapping ip <ip-address>
IP address:    192.0.2.1 (vsys1)
User:         localdomain\username
From:         SYSLOG
```

4. Log out of the client system.
5. Verify that the firewall deleted the user mapping:

```
> show user ip-user-mapping ip <ip-address>
No matched record
```

Configure the Windows User-ID Agent as a Syslog Listener

To configure the Windows-based User-ID agent to create new user mappings and remove outdated mappings through syslog monitoring, start by defining Syslog Parse profiles. The User-ID agent uses the profiles to find login and logout events in syslog messages. In environments where *syslog senders* (the network services that authenticate users) deliver syslog messages in different formats, configure a profile for each syslog format. Syslog messages must meet certain criteria for a User-ID agent to parse them (see [Syslog](#)). This procedure uses examples with the following formats:

- Login events—[Tue Jul 5 13:15:04 2016 CDT] Administrator authentication success User:johndoe1 Source:192.168.3.212
- Logout events—[Tue Jul 5 13:18:05 2016 CDT] User logout successful User:johndoe1 Source:192.168.3.212

After configuring the Syslog Parse profiles, you specify the syslog senders that the User-ID agent monitors.



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 sender. 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 sender and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending syslogs to the User-ID agent.

STEP 1 | Deploy the Windows-based User-ID agents if you haven't already.

1. [Install the Windows-Based User-ID Agent.](#)
2. [Configure the firewall to connect to the User-ID agent.](#)

STEP 2 | Define custom Syslog Parse profiles to create and delete user mappings.

Each profile filters syslog messages to identify either login events (to create user mappings) or logout events (to delete mappings), but no single profile can do both.

1. Review the syslog messages that the syslog sender generates to identify the syntax for login and logout events. This enables you to define the matching patterns when creating Syslog Parse profiles.



While reviewing syslog messages, also determine whether they include the domain name. If they don't, and your user mappings require domain names, enter the Default Domain Name when defining the syslog senders that the User-ID agent monitors (later in this procedure).

2. Open the Windows **Start** menu and select **User-ID Agent**.
3. Select **User Identification > Setup** and **Edit** the Setup.
4. Select **Syslog**, **Enable Syslog Service**, and **Add** a Syslog Parse profile.
5. Enter a **Profile Name** and **Description**.
6. Select the **Type** of parsing to find login and logout events in syslog messages:
 - **Regex**—Regular expressions.
 - **Field**—Text strings.

The following steps describe how to configure these parsing types.

STEP 3 | (Regex parsing only) Define the regex matching patterns.

If the syslog message contains a standalone space or tab as a delimiter, use `\s` for a space and `\t` for a tab.

1. Enter the **Event Regex** for the type of events you want to find:
 - **Login events**—For the example message, the regex `(authentication\ success){1}` extracts the first `{1}` instance of the string `authentication success`.
 - **Logout events**—For the example message, the regex `(logout\ successful){1}` extracts the first `{1}` instance of the string `logout successful`.

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.


2. Enter the **Username Regex** to identify the start of the username.

In the example message, the regex `User: ([a-zA-Z0-9\\\. _]+)` matches the string `User: johndoe1` and identifies `johndoe1` as the username.

3. Enter the **Address Regex** to identify the IP address portion of syslog messages.

In the example 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`.

The following is an example of a completed Syslog Parse profile that uses regex to identify login events:



4. Click **OK** twice to save the profile.

STEP 4 | (Field Identifier parsing only) Define string matching patterns.

1. Enter an **Event String** to identify the type of events you want to find.
 - **Login events**—For the example message, the string `authentication success` identifies login events.
 - **Logout events**—For the example message, the string `logout successful` identifies logout events.
2. Enter a **Username Prefix** 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 example messages, `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. Enter an **Address Prefix** 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 example messages, `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.

The following is an example of a completed Syslog Parse profile that uses string matching to identify login events:



6. Click **OK** twice to save the profile.

STEP 5 | Specify the syslog senders that the User-ID agent monitors.

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 discards any syslog messages received from senders that are not on this list.

1. Select **User Identification > Discovery** and **Add** an entry to the Servers list.
2. Enter a **Name** to identify the sender.
3. Enter the **Server Address** of the syslog sender (IP address or FQDN).
4. Set the **Server Type** to **Syslog Sender**.
5. (Optional) If you want to override the current domain in the username of your syslog message or prepend the domain to the username if your syslog message doesn't contain a domain, enter a **Default Domain Name**.
6. For each syslog format that the sender supports, **Add** a Syslog Parse profile to the Filter list. Select the **Event Type** that you configured each profile to identify—**login** (default) or **logout**—and then click **OK**.
7. Click **OK** to save the settings.
8. **Commit** your changes to the User-ID agent configuration.

STEP 6 | Verify that the User-ID agent adds and deletes user mappings when users log in and out.



You can use [CLI commands](#) to see additional information about syslog senders, syslog messages, and user mappings.

1. Log in to a client system for which a monitored syslog sender generates login and logout event messages.
2. Verify that the User-ID agent mapped the login username to the client IP address:
 1. In the User-ID agent, select **Monitoring**.
 2. Enter the username or IP address in the filter field, **Search**, and verify that the list displays the mapping.
3. Verify that the firewall received the user mapping from the User-ID agent:
 1. [Log in to the firewall CLI](#).
 2. Run the following command:

```
> show user ip-user-mapping ip <ip-address>
```

If the firewall received the user mapping, the output resembles the following:

```
IP address:    192.0.2.1 (vsys1)
User:          localdomain\username
From:          SYSLOG
```

4. Log out of the client system.
5. Verify that the User-ID agent removed the user mapping:
 1. In the User-ID agent, select **Monitoring**.
 2. Enter the username or IP address in the filter field, **Search**, and verify that the list does not display the mapping.
6. Verify that the firewall deleted the user mapping:
 1. Access the firewall CLI.
 2. Run the following command:

```
> show user ip-user-mapping ip <ip-address>
```

If the firewall deleted the user mapping, the output displays:

```
No matched record
```

Map IP Addresses to Usernames Using Captive Portal

When a user initiates web traffic (HTTP or HTTPS) that matches an [Authentication Policy](#) rule, the firewall prompts the user to authenticate through Captive Portal. This ensures that you know exactly who is accessing your most sensitive applications and data. Based on user information collected during authentication, the firewall creates a new IP address-to-username mapping or updates the existing mapping for that user. This method of user mapping is useful in environments where the firewall cannot learn mappings through other methods such as monitoring servers. For example, you might have users who are not logged in to your monitored domain servers, such as users on Linux clients.

- [Captive Portal Authentication Methods](#)
- [Captive Portal Modes](#)
- [Configure Captive Portal](#)

Captive Portal Authentication Methods

Captive Portal uses the following methods to authenticate users whose web requests match [Authentication Policy](#) rules:

Authentication Method	Description
Kerberos SSO	The firewall uses Kerberos single sign-on (SSO) to transparently obtain user credentials from the browser. 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. 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 Authentication policy and Captive Portal configuration.  <i>Kerberos SSO is preferable to 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.</i>
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.

Authentication Method	Description
	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 Authentication policy and 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. For this method, you can configure Authentication policy to use Multi-Factor Authentication (MFA), SAML, Kerberos, TACACS+, RADIUS, or LDAP authentication. Although users have to manually enter their login 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 Authentication policy 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, use this mode only 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 timeouts 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

Mode	Description
	trusted sites. Redirect mode is also required if you use Multi-Factor Authentication to authenticate Captive Portal users.

Configure Captive Portal

The following procedure shows how to set up Captive Portal authentication by configuring the PAN-OS integrated User-ID agent to redirect web requests that match an [Authentication Policy](#) rule to a firewall interface (redirect host). Based on their sensitivity, the applications that users access through Captive Portal require different authentication methods and settings. To accommodate all authentication requirements, you can use default and custom authentication enforcement objects. Each object associates an Authentication rule with an authentication profile and a Captive Portal authentication method.

- **Default authentication enforcement objects**—Use the default objects if you want to associate multiple Authentication rules with the same global authentication profile. You must [configure this authentication profile](#) before configuring Captive Portal, and then assign it in the Captive Portal Settings. For Authentication rules that require [Multi-Factor Authentication](#) (MFA), you cannot use default authentication enforcement objects.
- **Custom authentication enforcement objects**—Use a custom object for each Authentication rule that requires an authentication profile that differs from the global profile. Custom objects are mandatory for Authentication rules that require MFA. To use custom objects, create authentication profiles and assign them to the objects after configuring Captive Portal—when you [Configure Authentication Policy](#).

Keep in mind that authentication profiles are necessary only if users authenticate through a Captive Portal [Web Form](#), [Kerberos SSO](#), or [NT LAN Manager \(NTLM\)](#). Alternatively, or in addition to these methods, the following procedure also describes how to implement [Client Certificate Authentication](#).



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.

STEP 1 | Configure the interfaces that the firewall will use for incoming web requests, authenticating users, and communicating with directory servers to map usernames to IP addresses.

When the firewall connects to authentication servers or User-ID agents, it uses the management interface by default. As a best practice, isolate your management network by configuring service [routes](#) to connect to the authentication servers or User-ID agents.

1. (**MGT interface only**) Select **Device > Setup > Interfaces**, edit the **Management** interface, select **User-ID**, and click **OK**.
2. (**Non-MGT interface only**) [Assign an Interface Management Profile](#) to the Layer 3 interface that the firewall will use for incoming web requests and communication 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, TACACS+, or Multi-Factor Authentication**—Configure a service route for any authentication services that you use.
 - **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-220> ping host dc1.acme.com
```

STEP 3 | 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.



If you don't assign an SSL/TLS Service Profile, the firewall uses TLS 1.2 by default. To use a different TLS version, configure an SSL/TLS Service Profile for the TLS version you want to use.

5. Configure clients to trust the certificate:
 1. [Export the CA certificate](#) you created or imported.
 2. 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 4 | (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 through 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 5 | (Optional) Enable NT LAN Manager (NTLM) authentication.



As a best practice, choose [Kerberos](#) single sign-on (SSO) or [SAML SSO](#) authentication over NTLM authentication. Kerberos and SAML are stronger, more robust authentication methods than NTLM and do 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](#).



As a best practice, you use a User-ID agent account that is separate from your firewall administrator account.

2. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User ID Agent Setup section.
3. Select **NTLM** and **Enable NTLM authentication processing**.
4. Enter the **NTLM Domain** against which the User-ID agent on the firewall will check NTLM credentials.
5. Enter the **Admin User Name** 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.

6. Click **OK** to save your settings.

STEP 6 | (Optional) Configure Captive Portal for the Apple Captive Network Assistant.

This step is only required if you are using Captive Portal with the Apple Captive Network Assistant (CNA). To use Captive Portal with CNA, perform the following steps.

1. Verify you have specified an FQDN for the redirect host (not just an IP address).
2. Select an [SSL/TLS service profile](#) that uses a publicly-signed certificate for the specified FQDN.
3. Enter the following command to adjust the number of requests supported for Captive Portal: **set deviceconfig setting ctd cap-portal-ask-requests <threshold-value>**

By default, the firewall has a rate limit threshold for Captive Portal that limits the number of requests to one request every two seconds. The CNA sends multiple requests that can exceed this limit, which can result in a TCP reset and an error from the CNA. The recommended threshold value is 5 (default is one). This value will allow up to 5 requests every two seconds. Based on your environment, you may need to configure a different value. If the current value is not sufficient to handle the number of requests, increase the value.

STEP 7 | Configure the Captive Portal settings.

1. Select **Device > User Identification > Captive Portal Settings** and edit the settings.
2. **Enable Captive Portal** (default is enabled).
3. Specify the **Timer**, which is the maximum time in minutes that the firewall retains an IP address-to-username mapping for a user after that user authenticates through Captive Portal (default is 60; range is 1 to 1,440). After the **Timer** expires, the firewall removes the mapping and any associated [Authentication Timestamps](#) used to evaluate the **Timeout** in Authentication policy rules.



When evaluating the Captive Portal Timer and the Timeout value in each Authentication policy rule, the firewall prompts the user to re-authenticate for whichever setting expires first. Upon re-authenticating, the firewall resets the time

count for the Captive Portal Timer and records new authentication timestamps for the user. Therefore, to enable different Timeout periods for different Authentication rules, set the Captive Portal Timer to a value the same as or higher than any rule Timeout.

4. Select the **SSL/TLS Service Profile** you created for redirect requests over TLS. See [Configure an SSL/TLS Service Profile](#).
5. Select the **Mode** (in this example, **Redirect**).
6. (**Redirect mode only**) Specify the **Redirect Host**, which 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 web requests are redirected.

If users authenticate through [Kerberos](#) single sign-on (SSO), the **Redirect Host** must be the same as the hostname specified in the Kerberos keytab.

7. Select the authentication method to use if NTLM fails (or if you don't use NTLM):
 - To use client certificate authentication, select the **Certificate Profile** you created.
 - To use global settings for interactive or SSO authentication, select the **Authentication Profile** you configured.
 - To use Authentication policy rule-specific settings for interactive or SSO authentication, assign authentication profiles to authentication enforcement objects when you [Configure Authentication Policy](#).
8. Click **OK** and **Commit** the Captive Portal configuration.

STEP 8 | Next steps...

The firewall does not display the Captive Portal web form to users until you [Configure Authentication Policy](#) rules that trigger authentication when users request services or applications.

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 identify specific users on Windows-based terminal servers, the Palo Alto Networks Terminal Server agent (TS agent) allocates a port range to each user. The TS agent 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, configure the PAN-OS XML API to extract user mapping information. The following values apply for both methods:

- Default port range: 1025 to 65534
- Per user block size: 200
- Maximum number of multi-user systems: 2,500

For information about the terminal servers supported by the TS agent and the number of TS agents supported on each firewall model, refer to the [Palo Alto Networks Compatibility Matrix](#) and the [Product Comparison Tool](#).

The following sections describe how to configure user mapping for terminal server users:

- [Configure the Palo Alto Networks Terminal Server \(TS\) Agent for User Mapping](#)
- [Retrieve User Mappings from a Terminal Server Using the PAN-OS XML API](#)

Configure the Palo Alto Networks Terminal Server (TS) 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 to which your users log in.



- If you are using TS agent 7.0 or a later version, disable any Sophos antivirus software on the TS agent host. Otherwise, the antivirus software overwrites the source ports that the TS agent allocates.
- If you are installing a TS agent running TS agent 8.0 or a later version on Windows Server 2008 R2, download and install the [security advisory patch](#) on the server before you install the TS agent. Otherwise, the TS agent service fails to launch and writes the following error to the logs: `servicefails` with error 577.

For information about default values, ranges, and other specifications, refer to [Configure User Mapping for Terminal Server Users](#). For information about the terminal servers supported by the TS agent and the number of TS agents supported on each firewall model, refer to the [Palo Alto Networks Compatibility Matrix](#).

STEP 1 | Download the TS agent installer.

1. Log in to the [Palo Alto Networks Customer Support Portal](#).
2. Select **Updates** > **Software Updates**.
3. Set **Filter By** to **Terminal Services Agent** and select the version of the agent you want to install from the corresponding Download column. For example, to download TS agent 9.0, select **TaInstall-9.0.msi**.
4. Save the `TaInstall1.x64-x.x.x-xx.msi` or `TaInstall-x.x.x-xx.msi` file on the systems where you plan to install the agent; be sure to select the appropriate version based on whether the Windows system is running a 32-bit or a 64-bit OS.

Version	Release Date	Release Notes	Download	Size	Checksum
Terminal Services Agent					
8.0.9	05/02/2018	TS_Agent_8.0.9_RN.pdf	TaInstall-8.0.9.msi	13 MB	Checksum
8.0.9-64	05/02/2018	TS_Agent_8.0.9_RN.pdf	TaInstall64x64-8.0.9.msi	15 MB	Checksum
8.1.1	05/02/2018	TS_Agent_8.1.1_RN.pdf	TaInstall-8.1.1.msi	13 MB	Checksum
8.1.1-64	05/02/2018	TS_Agent_8.1.1_RN.pdf	TaInstall64x64-8.1.1.msi	15 MB	Checksum
8.1.1-64	03/21/2018	TS_Agent-8.1.1-RN.pdf	TaInstall64x64-8.1.1.msi	15 MB	Checksum
8.1.1	03/21/2018	TS_Agent-8.1.1-RN.pdf	TaInstall-8.1.1.msi	13 MB	Checksum
8.0.8-64	03/08/2018	TS_Agent_8.0_RN.pdf	TaInstall64x64-8.0.8.msi	15 MB	Checksum
8.0.8	03/08/2018	TS_Agent_8.0_RN.pdf	TaInstall-8.0.8.msi	13 MB	Checksum
8.1.0-64	03/06/2018	TS_Agent_8.1_RN.pdf	TaInstall64x64-8.1.0.msi	15 MB	Checksum

STEP 2 | Run the installer as an administrator.

1. Open the Windows **Start** menu, right-click the **Command Prompt** program, and **Run as administrator**.
2. From the command line, run the .msi file you downloaded. For example, if you saved the `TaInstall-9.0.msi` file to the Desktop, then enter the following:

```
C:\Users\administrator.acme>cd Desktop
C:\Users\administrator.acme\Desktop>TaInstall-9.0.0-1.msi
```

3. Follow the setup prompts to install the agent using the default settings. The setup installs the agent in C:\ProgramFiles (x86)\Palo Alto Networks\Terminal Server Agent.



To ensure correct port allocation, you must use the default Terminal Server agent installation folder location.

4. When the installation completes, **Close** the setup dialog.



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 you upgrade.

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 specify the range of ports that are allocated to non-user sessions. Make sure the values in these fields do not overlap with the ports you designate for user traffic. These values can be changed only by editing the corresponding Windows registry settings. The TS agent does not allocate ports for network traffic emitted by session 0.

1. Open the Windows **Start** menu and select **Terminal Server Agent** to launch the Terminal Server agent application.
2. **Configure** (side menu) the agent.
3. Enter the **Source Port Allocation Range** (default is 20,000-39,999). This is the full range of port numbers that the TS agent will allocate for user mapping. The port range you specify cannot overlap the **System Source Port Allocation Range**.
4. (Optional) If there are ports or 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 (**Port Allocation Start Size Per User**); default is 200.
6. Specify the **Port Allocation Maximum Size Per User**, which is the maximum number of ports the Terminal Server agent can allocate to an individual user.
7. Specify whether to continue processing traffic from the user if the user runs out of allocated ports. The **Fail port binding when available ports are used up** option is enabled by default, 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, disable (clear) this option, but if you do, this traffic may not be identified with User-ID.

STEP 4 | (Optional) Assign your own certificates for mutual authentication between the TS agent and the firewall.

1. Obtain your certificate for the TS agent from your enterprise PKI or generate one on your firewall. The private key of the server certificate must be encrypted and the certificate must be uploaded in PEM file format. Perform one of the following tasks to assign a certificate:
 - [Generate a Certificate](#) and export it.
 - Export a certificate from your enterprise certificate authority (CA).
2. Add a server certificate to the TS agent.
 1. On the TS agent, select **Server Certificate** and **Add** a new certificate.

2. Enter the path and name of the certificate file received from the CA or browse to the certificate file.
 3. Enter the private key password.
 4. Click **OK**.
 5. **Commit** your changes.
3. Configure and assign the certificate profile for the firewall.

1. Select **Device > Certificate Management > Certificate Profile** to [Configure a Certificate Profile](#).



You can assign only one certificate profile for Windows User-ID agents and TS agents. Therefore, your certificate profile must include all certificate authorities that issued certificates uploaded to connected Windows User-ID and TS agents.

2. Select **Device > User Identification > Connection Security**.
3. Edit  and select the certificate profile you configured in the previous step as the **User-ID Certificate Profile**.
4. Click **OK**.
5. **Commit** your changes.

STEP 5 | Configure the firewall to connect to the Terminal Server agent.

Complete the following steps on each firewall you want to connect to the Terminal Server agent to receive user mappings:

1. Select **Device > User Identification > Terminal Server Agents** and **Add** a new TS agent.
2. Enter a **Name** for the Terminal Server agent.
3. Enter the hostname or IP address of the Windows **Host** on which the Terminal Server agent is installed.

The hostname or IP address must resolve to a static IP address. If you change the existing hostname, the TS agent resets when you commit the changes to resolve the new hostname. If the hostname resolves to multiple IP addresses, the TS agent uses the first address in the list.

4. (Optional) Enter the hostname or IP address for any **Alternative IP Addresses** that can appear as the source IP address for the outgoing traffic.

The hostname or IP address must resolve to a static IP address. You can enter up to 8 IP addresses or hostnames.

5. Enter the **Port** number on which the agent will listen for user mapping requests. This value must match the value configured on the Terminal Server agent. By default, the port is set to 5009 on the firewall and on the agent. If you change it on the firewall, you must also change the **Listening Port** on the Terminal Server agent **Configure** dialog to the same port.
6. Make sure that the configuration is **Enabled** and then click **OK**.
7. **Commit** your changes.
8. Verify that the **Connected** status displays as connected (a green light).

STEP 6 | Verify that the Terminal Server 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 Server agent is successfully mapping port ranges to usernames (**Monitor** in the side menu) and confirm that the mapping table is populated.

STEP 7 | (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 **Settings > Internet options > Advanced** and scroll to the Security section.
3. Disable (clear) the **Enable Enhanced Protected Mode** option.
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.

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>
```

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.

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=<input_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=k7J335J6hI7nBxIqyfa62sZugWx7ot%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"
```

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-5250> 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-5250> 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.

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.

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.

STEP 2 | Configure the rules that control user access based on the groups you just configured.

For more information, refer to [Enable user- and group-based policy enforcement](#).

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.

Verify the User-ID Configuration

After you configure user and group mapping, enable User-ID in your Security policy, and configure Authentication policy, you should verify that User-ID works properly.

STEP 1 | [Access the firewall CLI.](#)

STEP 2 | Verify that group mapping is working.

From the CLI, enter the following operational command:

```
> show user group-mapping statistics
```

STEP 3 | 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 4 | Test your Security policy 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 troubleshoot the running configuration 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:
 1. Select **Device > Troubleshooting**, and select **Security Policy Match** from the Select Test drop-down.
 2. Enter **0.0.0.0** as the Source and Destination IP addresses. This executes the policy match test against any source and destination IP addresses.
 3. Enter the Destination Port.
 4. Enter the Protocol.
 5. **Execute** the security policy match test.

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 (such as regional directory services) but need access to remote services and applications (such as global data center applications).

If you [Configure Authentication Policy](#), your firewalls must also redistribute the [Authentication Timestamps](#) associated with user responses to authentication challenges. Firewalls use the timestamps to evaluate the timeouts for Authentication policy rules. The timeouts allow a user who successfully authenticates to later request services and applications without authenticating again within the timeout periods. Redistributing timestamps enables you to enforce consistent timeouts for each user even if the firewall that initially grants a user access is not the same firewall that later controls access for that user.

If you have configured multiple virtual systems, you can share IP address-to-username mapping information across virtual systems by selecting a virtual system as a User-ID hub.

- [Deploy User-ID for Numerous Mapping Information Sources](#)
- [Redistribute User Mappings and Authentication Timestamps](#)
- [Share User-ID Mappings Across Virtual Systems](#)

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.

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.

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 all the *Windows Event Collectors*—the member servers that collect login events from domain controllers. The following is an overview of the tasks; consult your [Windows Server documentation](#) for the specific steps.

STEP 1 | On each Windows Event Collector, 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) and 2016, 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).



To forward events as quickly as possible, Minimize Latency when configuring the subscription.

User-ID agents monitor the Security log, not the default forwarded events location, on Windows Event Collectors. Therefore, perform the following steps on each Windows Event Collector to change the event logging path to the Security log.

1. Open the Event Viewer.
2. Right-click the **Security** log and select **Properties**.
3. Copy the **Log path** (default `%SystemRoot%\System32\Winevt\Logs\security.evtx`) and click **OK**.
4. Right-click the **Forwarded Events** folder and select **Properties**.
5. Replace the default **Log path** (`%SystemRoot%\System32\Winevt\Logs\ForwardedEvents.evtx`) by pasting the value from the **Security** log, and then click **OK**.

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

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 Windows-Based User-ID Agent](#) on a Windows server that can access the member servers. Make sure the system that will host the User-ID agent is a member of the same domain as the servers it will monitor.

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:
 1. In the Servers section, click **Add** and enter a **Name** to identify the member server.
 2. In the **Server Address** field, enter the FQDN or IP address of the member server.
 3. For the **Server Type**, select **Microsoft Active Directory**.
 4. Click **OK** to save the server entry.
3. Configure the remaining User-ID agent settings: see [Configure the Windows-Based User-ID Agent for User Mapping](#).
4. If the User-ID sources provide usernames in multiple formats, specify the format for the **Primary Username** when you [Map Users to Groups](#).

The primary username is the username that identifies the user on the firewall and represents the user in reports and logs, regardless of the format that the User-ID source provides.

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**.

Redistribute User Mappings and Authentication Timestamps

Every firewall that enforces user-based policy requires user mapping information. In a large-scale network, instead of configuring all your firewalls to directly query the mapping information sources, you can streamline resource usage by configuring some firewalls to collect mapping information through redistribution. Redistribution also enables the firewalls to enforce user-based policies when users rely on local sources for authentication (such as regional directory services) but need access to remote services and applications (such as global data center applications).



You can redistribute user mapping information collected through any method except Terminal Services (TS) agents. You cannot redistribute [Group Mapping](#) or [HIP match](#) information.

If you use Panorama and Dedicated Log Collectors to manage firewalls and aggregate firewall logs, you can use Panorama to [manage User-ID redistribution](#). Leveraging Panorama and your distributed log collection infrastructure is a simpler solution than creating extra connections between firewalls to redistribute User-ID information.

If you [Configure Authentication Policy](#), your firewalls must also redistribute the [Authentication Timestamps](#) that are generated when users authenticate to access applications and services. Firewalls use the timestamps to evaluate the timeouts for Authentication policy rules. The timeouts allow a user who successfully authenticates to later request services and applications without authenticating again within the timeout periods. Redistributing timestamps enables you to enforce consistent timeouts across all the firewalls in your network.

Firewalls share user mappings and authentication timestamps as part of the same redistribution flow; you don't have to configure redistribution for each information type separately.

- [Firewall Deployment for User-ID Redistribution](#)
- [Configure User-ID Redistribution](#)

Firewall Deployment for User-ID Redistribution

To aggregate User-ID information, 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 map IP addresses to usernames. Each higher layer has firewalls that receive the mapping information and authentication timestamps from up to 100 redistribution points in the layer beneath it. The top-layer firewalls aggregate the mappings and timestamps from all layers. This deployment provides the option to configure policies for all users in top-layer firewalls

and region- or function-specific policies for a subset of users in the corresponding domains served by lower-layer firewalls.

Figure 6: User-ID and Timestamp Redistribution shows a deployment with three layers of firewalls that redistribute mappings and timestamps from local offices to regional offices and then to a global data center. The data center firewall that aggregates all the information shares it with other data center firewalls so that they can all enforce policy and generate reports for users across your entire network. Only the bottom layer firewalls use User-ID agents to query the directory servers.

The information sources that the User-ID agents query 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 information in one data center firewall and the second to share the information with other data center firewalls.




Figure 6: User-ID and Timestamp Redistribution

Configure User-ID Redistribution

Before you configure User-ID redistribution:

- ❑ Plan the redistribution architecture. Some factors to consider are:
 - Which firewalls will enforce 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 all User-ID information? The maximum allowed number of hops is ten.
 - How can you minimize the number of firewalls that query the user mapping information sources? The fewer the number of querying firewalls, the lower the processing load is on both the firewalls and sources.
- ❑ Configure user mapping using [PAN-OS Integrated User-ID agents](#) or [Windows-based User-ID agents](#).
- ❑ [Configure Authentication Policy](#).

Perform the following steps on the firewalls in the User-ID redistribution sequence.

STEP 1 | Configure the firewall to redistribute User-ID information.

Skip this step if the firewall receives but does not redistribute User-ID information.

1. Select **Device > User Identification > User Mapping**.
2. (**Firewalls with multiple virtual systems only**) Select the **Location**. You must configure the User-ID settings for each virtual system.



You can redistribute 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.

3. Edit the Palo Alto Networks User-ID Agent Setup and select **Redistribution**.
4. Enter a **Collector Name** and **Pre-Shared Key** to identify this firewall or virtual system as a User-ID agent.
5. Click **OK** to save your changes.

STEP 2 | Configure the service route that the firewall uses to query other firewalls for User-ID information.

Skip this step if the firewall receives user mapping information from Windows-based User-ID agents or directly from the information sources (such as directory servers) instead of from other firewalls.

1. Select **Device > Setup > Services**.
2. (**Firewalls with multiple virtual systems only**) Select **Global** (for a firewall-wide service route) or **Virtual Systems** (for a virtual system-specific service route), and then [configure the service route](#).
3. Click **Service Route Configuration**, select **Customize**, and select **IPv4** or **IPv6** based on your network protocols. Configure the service route for both protocols if your network uses both.
4. Select **UID Agent** and then select the **Source Interface** and **Source Address**.
5. Click **OK** twice to save the service route.

STEP 3 | Enable the firewall to respond when other firewalls query it for User-ID information.

Skip this step if the firewall receives but does not redistribute User-ID information.

[Configure an Interface Management Profile](#) with the **User-ID** service enabled and assign the profile to a firewall interface.

STEP 4 | Commit and verify your changes.

1. **Commit** your changes to activate them.
2. [Access the CLI](#) of a firewall that redistributes User-ID information.
3. Display all the user mappings by running the following command:

```
> show user ip-user-mapping all
```

4. Record the IP address associated with any username.
5. Access the CLI of a firewall that receives redistributed User-ID information.
6. Display the mapping information and authentication timestamp for the *<ip_address>* you recorded:

```
> show user ip-user-mapping <ip_address>
IP address:    192.0.2.0 (vsys1)
User:         corpdomain\username1
From:         UIA
Idle Timeout:  10229s
Max. TTL:     10229s
MFA Timestamp: first(1) - 2016/12/09 08:35:04
Group(s):     corpdomain\groupname(621)
```



This example output shows the authentication timestamp for one response to an authentication challenge (factor). For Authentication policy rules that use [Multi-Factor Authentication \(MFA\)](#), the output shows multiple [Authentication Timestamps](#).

Share User-ID Mappings Across Virtual Systems

To simplify User-ID™ source configuration when you have multiple virtual systems, configure the User-ID sources on a single [virtual system](#) to share IP address-to-username mappings with all other virtual systems on the firewall.


Configuring a single virtual system as a *User-ID hub* simplifies user mapping by eliminating the need to configure the sources on multiple virtual systems, especially if a user's traffic will pass through multiple virtual systems based on the resources the user is trying to access (for example, in an academic networking environment where a student will be accessing different departments whose traffic is managed by different virtual systems).

To map the user, the firewall uses the mapping table on the local virtual system and applies the policy for that user. If the firewall does not find the mapping for a user on the virtual system where that user's traffic originated, the firewall queries the hub to fetch the IP address-to-username information for that user. If the firewall locates the mapping on both the User-ID hub and the local virtual system, the firewall uses the mapping it learns locally.

After you configure the User-ID hub, the virtual system can use the mapping table on the User-ID hub when it needs to identify a user for user-based policy enforcement or to display the username in a log or report but the source is not available locally. When you select a hub, the firewall retains the mappings on other virtual systems so we recommend consolidating the User-ID sources on the hub. However, if you don't want to share mappings from a specific source, you can configure an individual virtual system to perform user mapping.

STEP 1 | Assign the [virtual system](#) as a User-ID hub.

1. Select **Device > Virtual Systems** and then select the virtual system where you consolidated your User-ID sources.
2. On the **Resource** tab, **Make this vsys a User-ID data hub** and click **Yes** to confirm. Then click **OK**.



STEP 2 | Consolidate your User-ID sources and migrate them to the virtual system that you want to use as a User-ID hub.

This consolidates the User-ID configuration for operational simplicity. By configuring the hub to monitor servers and connect to agents that were previously monitored by other virtual systems, the hub collects the user mapping information instead of having each virtual system collect it independently. If you don't want to share mappings from specific virtual systems, configure those mappings on a virtual system that will not be used as the hub.

1. Remove any sources that are unnecessary or outdated.
2. Identify all configurations for your [Windows-based](#) or [integrated](#) agents and any sources that send user mappings using the [XML API](#) and copy them to the virtual system you want to use as a User-ID hub.



On the hub, you can configure any User-ID source that is currently configured on a virtual system. However, IP address-and-port-to-username mapping information from Terminal Server agents and group mappings are not shared between the User-ID hub and the connected virtual systems.

3. Specify the subnetworks that User-ID should [include in or exclude from](#) mapping.
4. [Define](#) the **Ignore User List**.
5. On all other virtual systems, remove any sources that are on the User-ID hub.

STEP 3 | **Commit** the changes to enable the User-ID hub and begin collecting mappings for the consolidated sources.

STEP 4 | Confirm the User-ID hub is mapping the users.

1. Use the **show user ip-user-mapping all** command to show the IP address-to-username mappings and which virtual system provides the mappings.
2. Use the **show user user-id-agent statistics** command to show which virtual system is serving as the User-ID hub.

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.

New and modified App-IDs are released as part of Applications and Threat Content Updates—follow the Best Practices for Applications and Threats Content Updates to seamlessly keep your application and threat signatures up-to-date.

- > App-ID Overview
- > App-ID and HTTP/2 Inspection
- > Manage Custom or Unknown Applications
- > Manage New and Modified App-IDs
- > Use Application Objects in Policy
- > Safely Enable Applications on Default Ports
- > Applications with Implicit Support
- > Security Policy Rule Optimization
- > Application Level Gateways
- > Disable the SIP Application-level Gateway (ALG)
- > Use HTTP Headers to Manage SaaS Application Access
- > Maintain Custom Timeouts for Legacy Applications

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.

App-ID and HTTP/2 Inspection

You can now safely enable applications running over HTTP/2, without any additional configuration on the firewall. As more websites continue to adopt HTTP/2, the firewall can enforce security policy and all threat detection and prevention capabilities on a stream-by-stream basis. This visibility into HTTP/2 traffic enables you to secure web servers that provide services over HTTP/2, and allow your users to benefit from the speed and resource efficiency gains that HTTP/2 provides.



The firewall processes and inspects HTTP/2 traffic by default when [SSL decryption](#) is enabled. For HTTP/2 inspection to work correctly, the firewall must be enabled to use ECDHE (elliptic curve Diffie-Hellman) as a key exchange algorithm for SSL sessions. ECDHE is enabled by default, but you can check to confirm that it's enabled by selecting **Objects > Decryption > Decryption Profile > SSL Decryption > SSL Protocol Settings**.



Follow these steps to disable HTTP/2 inspection for targeted traffic, or globally.



Make sure that you're running Applications and Threats content version 8202 or later if you're performing Safe Search Enforcement with HTTP/2 Inspection. This content release makes it possible to enforce Safe Search while HTTP/2 Inspection is enabled.

- Disable HTTP/2 inspection for targeted traffic.

You'll need to specify for the firewall to remove any value contained in the Application-Layer Protocol Negotiation (ALPN) TLS extension. ALPN is used to secure HTTP/2 connections—when there is no value specified for this TLS extension, the firewall either downgrades HTTP/2 traffic to HTTP/1.1 or classifies it as unknown TCP traffic.



1. Select **Objects > Decryption > Decryption Profile > SSL Decryption > SSL Forward Proxy** and then select **Strip ALPN**.
 2. Attach the decryption profile to a decryption policy (**Policies > Decryption**) to turn off HTTP/2 inspection for traffic that matches the policy.
 3. **Commit** your changes.
- Disable HTTP/2 inspection globally.

Use the CLI command: `set deviceconfig setting http2 enable no` and **Commit** your changes. The firewall will classify HTTP/2 traffic as unknown TCP traffic.

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 and Modified App-IDs

New and modified App-IDs are delivered to the firewall as part of [Applications and Threat Content Updates](#). While new and modified App-IDs enable the firewall to enforce your security policy with ever-increasing precision, changes in security policy enforcement that can occur when a content update release is installed can impact application availability. For this reason, you will need to think about how to best deploy content updates so that you can get the latest threat prevention as it's made available, and adjust your security policy to best leverage new and modified App-IDs.

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:

- [Workflow to Best Incorporate New and Modified App-IDs](#)
- [See the New and Modified App-IDs in a Content Release](#)
- [See How New and Modified App-IDs Impact Your Security Policy](#)
- [Ensure Critical New App-IDs are Allowed](#)
- [Monitor New App-IDs](#)
- [Disable and Enable App-IDs](#)

Workflow to Best Incorporate New and Modified App-IDs

Refer to this master workflow to first set up Application and Threat content updates, and then to best incorporate new and modified App-IDs into your security policy. Everything you need to deploy content updates is referenced here.

STEP 1 | Align your business needs with an approach to deploying Application and Threat content updates.

Learn how [Applications and Threat Content Updates](#) work, and identify your organization as either [mission-critical or security-first](#). Understanding which of these is most important to your business will help you to decide how to best deploy content updates and apply best practices to meet your business needs. You might find that you want to apply a mix of both approaches, perhaps depending on firewall deployment (data center or perimeter) or office location (remote or headquarters).

STEP 2 | Review and apply the [Best Practices for Applications and Threats Content Updates](#) based on your organization's network security and application availability requirements.

STEP 3 | Configure a security policy rule to always allow new App-IDs that might have network-wide impact, like authentication or software development applications.

The New App-ID characteristic matches to only the App-IDs introduced in the latest content release. When used in a security policy, this gives you a month's time to fine tune your security policy based on new App-IDs while ensuring constant availability for App-IDs that fall into critical categories ([Ensure Critical New App-IDs are Allowed](#)).

STEP 4 | Set the schedule to [Deploy Application and Threat Content Updates](#); this includes the option to delay new App-ID installation until you've had time to make necessary security policy updates (using the **New App-ID Threshold**).

STEP 5 | After you've setup a content updates installation schedule, you'll want to regularly check in and [See the New and Modified App-IDs in a Content Release](#).

STEP 6 | You can then [See How New and Modified App-IDs Impact Your Security Policy](#), and make adjustments to your security policy as needed.

STEP 7 | [Monitor New App-IDs](#) to get a view into new App-ID activity on your network, so that you're best equipped to make the most effective security policy updates.

See the New and Modified App-IDs in a Content Release

For both downloaded and installed content updates, you can see a list of the new and modified App-IDs the update includes. Full application details are provided, and importantly, updates to applications with network-wide impact (for example, LDAP or IKE) are prominently flagged as a recommended for policy review. For modified App-IDs, application details also describe how coverage is either now expanded or more precise.

STEP 1 | Select **Device > Dynamic Updates** and select **Check Now** to refresh the list of available content updates.

STEP 2 | For either a downloaded or currently installed content release, click **Review Apps** link in the **Actions** column to view details on newly-identified and modified applications in that release:

Applications and Threats		Last checked: 2018/02/24 16:54:59 PST	Schedule: Every Sunday at 03:00 (Download and Install)			
771-4450	panup2-all-content-771-4450	Apps, Threats	36 MB			Download
772-4461	panup2-all-content-772-4461	Apps, Threats	36 MB			Download
778-4489	panup2-all-content-778-4489	Apps, Threats	44 MB	✓ previously		Revert
779-4500	panup2-all-content-779-4500	Apps, Threats	44 MB			Download
780-4504	panup2-all-content-780-4504	Apps, Threats	44 MB	✓		Install
781-4515	panup2-all-content-781-4515	Apps, Threats	44 MB	✓	✓	Review Policies Review Apps
782-4517	panup2-all-content-782-4517	Apps, Threats	44 MB			Download
783-4530	panup2-all-content-783-4530	Apps, Threats	44 MB			Download

STEP 3 | Review the App-IDs this content release introduces or modifies since the last content version.

New and modified App-IDs are listed separately. Full application details are provided for each, and App-IDs that Palo Alto Networks foresees as having network-wide impact are flagged as recommended for policy review.

The screenshot shows the Palo Alto Networks management console interface. On the left, a sidebar lists navigation options: 'New and Modified Applications since last installed content', 'New', 'Modified', and 'Policy Review Recommendations'. The main panel displays details for a new application, 'OpenText Enterprise Connect'. The details include: 'Name: OpenText Enterprise Connect', 'Description: Enterprise Connect simplifies the user experience and pulls together business content from multiple heterogeneous data sources into a logical business-centric view. It accelerates delivery of content applications by allowing users to access and interact with business content directly from familiar desktop applications, such as Microsoft Office, Microsoft Outlook, and Windows Explorer. Enterprise Connect is one of a set of Open Text Content Services that are the foundation of the Open Text ECM Suite.' 'Additional Information: OpenText Inc. (NYSE:OTEX)', 'Standard Ports: tcp/818,208,209,119', 'Depends on: web browsing', 'Explicitly Uses: none', 'Data Action: download', 'Expanded Coverage: web browsing, unknown top-level OpenText Enterprise Connect', 'Characteristics: Excludes: none, Used by Malware: none, Capable of File Transfer: yes, Has Known Vulnerabilities: yes, Turns On Other Applications: none, Ports to Monitor: none, Stability Index: none', 'Categories: business systems, Subcategories: management, Technology: client server, Risk: 3', 'Options: TCP Timeout (seconds): 3000, TCP half closed (seconds): 120, TCP Time Wait (seconds): 15, App-ID Enabled: yes'. A red circle highlights the 'Expanded Coverage' section, which states 'web browsing, unknown top-level OpenText Enterprise Connect'. Another red circle highlights the 'Policy Review Recommendations' button at the bottom left. A red arrow points from the 'Policy Review Recommendations' button to the 'Expanded Coverage' section.

New 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.

For modified App-IDs, details 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.

STEP 4 | Based on your findings, click **Review Policies** to see how the new and modified App-IDs impact security policy enforcement: [See How New and Modified App-IDs Impact Your Security Policy](#).


See How New and Modified App-IDs Impact Your Security Policy

Newly-categorized and modified App-IDs can change the way the firewall enforces traffic. Perform a content update policy review to see how new and modified App-IDs impact your security policy, and to easily make any necessary adjustments. You can perform a content update policy review for both downloaded and installed content.

STEP 1 | Select **Device > Dynamic Updates**.

STEP 2 | [See the New and Modified App-IDs in a Content Release](#) to learn more about each App-ID that a content release introduces or modifies.

STEP 3 | For a downloaded or currently installed content release, click **Review Policies** in the Action column. The **Policy review based on candidate configuration** dialog allows you to filter by **Content Version** and view either new or modified App-IDs introduced in a specific release (you can also filter the policy impact of new App-IDs according to **Rulebase**, **Virtual System**, and **Application**).



STEP 4 | Select an App-ID from the **Application** drop-down to view policy rules that currently enforce the application. The rules displayed are based on the App-IDs 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 5 | Use the detail provided in the policy review to plan policy rule updates to take effect when the App-ID is installed, or if the content release version that included the App-ID is currently installed, the changes you make take effect immediately.

You can **Add app to selected policies** or **Remove app from selected policies**.

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. To mitigate any impact to security policy enforcement, you can use the **New App-ID** characteristic in a security policy rule so that the rule always enforces the most recently introduced App-IDs without requiring you to make configuration changes when new App-IDs are installed. 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 match only to the new App-IDs in that content release version.


You can choose to enforce all new App-IDs, or target the security policy rule to enforce certain types of new App-IDs that might have network-wide or critical impact (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 monthly, so a policy rule that allows the latest App-IDs 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.

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:

- [Monitor New App-IDs](#)—Monitor and get reports on new App-ID activity.
- [See the New and Modified App-IDs in a Content Release](#)—See how the newly-installed App-IDs impact your existing security policy rules.

Monitor New App-IDs

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. 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. Whether you're using it on the ACC or to generate reports (or to [Ensure Critical New App-IDs are Allowed](#)), 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 match only to the new App-IDs in that content release version.

- Generate a report with details specifically regarding new applications (applications introduced only in the latest content release).



- Use the ACC to monitor new application activity: select **ACC** and under **Global Filters**, select **Application > Application Characteristics > New App-ID**.



Disable and Enable App-IDs

You can disable all App-IDs introduced in a content release if you want to immediately benefit from the latest threat prevention, and plan to enable the App-IDs later, and you can 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 include application signatures that are 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 all App-IDs in a content release or for scheduled content updates.

While this option allows you to be protected against threats, by giving you the option to enable the App-ID at a later time, Palo Alto Networks recommends that instead of disabling App-IDs on a regular basis, you should instead configure a security policy rule to [Temporarily Allow New App-IDs](#). This rule will always allow the new App-IDs introduced in only the latest content release. Because content updates that include new App-IDs are released only once a month, this gives you time to assess the new App-IDs and adjust your security policy to cover the new App-IDs if needed, all the while ensuring that availability for critical applications is not affected.

- 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.
- 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**.

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.

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.

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.

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.


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.




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.




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, 0xffffffff00) 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 steps 4.c and 4.d 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.

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).

Safely Enable Applications on Default Ports

Applications running on unusual ports can indicate an attacker that is attempting to circumvent traditional port-based protections. Application-default is a feature of Palo Alto Networks firewalls that gives you an easy way to prevent this type of evasion and safely enable applications on their most commonly-used ports. Application-default is a best practice for application-based security policies—it reduces administrative overhead, and closes security gaps that port-based policy introduces:

- ❑ **Less overhead**—Write simple application-based security policy rules based on your business needs, instead of researching and maintaining application-to-port mappings. We've defined the default ports for [all applications with an App-ID](#).
- ❑ **Stronger security**—Enabling applications to run only on their default ports is a security best practice. Application-default helps you to make sure that critical applications are available without compromising security if an application is behaving in an unexpected way.

Additionally, the default ports an application uses can sometimes depend on whether the application is encrypted or cleartext. Port-based policy requires you to open all the default ports an application might use to account for encryption. Open ports introduce security gaps that an attacker can leverage to bypass your security policy. However, application-default differentiates between encrypted and clear-text application traffic. This means that it can enforce the default port for an application, regardless of whether it is encrypted or not.

For example, without application-default, you would need to open ports 80 and 443 to enable web-browsing traffic—you'd be allowing both cleartext and encrypted web-browsing traffic on both ports. With application-default turned on, the firewall strictly enforces cleartext web-browsing traffic only on port 80 and SSL-tunneled traffic only on port 443.

To see the ports that an application uses by default, you can visit [Applipedia](#) or select **Objects** > **Applications**. Application details include the application's standard port—the port it most commonly uses when in cleartext. For web-browsing, SMTP, FTP, LDAP, POP3, and IMAP details also include the application's secure port—the port the application uses when encrypted.



Select **Policy** > **Security** and add or a modify a rule to enforce applications only on their default port(s):



Using application-default as part of an application-based security policy and with SSL decryption is a best practice. Additionally, if you have existing security policy rules that control web-browsing traffic with the Service set to service-http and service-https, you should update those rules to use application-default instead.

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 [Applopedia](#).

The following table lists the applications for which the firewall has implicit support (as of [Content Update 595](#)).

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
clubcooe	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

Application	Implicitly Supports
fastviewer	http, ssl
forticlient-update	http
good-for-enterprise	http, ssl
google-cloud-print	http, ssl, jabber
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

Application	Implicitly Supports
ms-frs	msrpc
ms-rdp	t_120
ms-scheduler	msrpc
ms-service-controller	msrpc
nfs	rpc
oovoo	http, ssl
paloalto-updates	ssl
panos-global-protect	http
panos-web-interface	http
pastebin	http
pastebin-posting	http
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

Security Policy Rule Optimization

Policy Optimizer provides a simple workflow to migrate your legacy Security policy rulebase to an App-ID based rulebase, which improves your security by reducing the attack surface and gaining visibility into applications so you can safely enable them. Policy Optimizer identifies port-based rules so you can convert them to application-based allow rules or add applications from a port-based rule to an existing application-based rule without compromising application availability. It also identifies over-provisioned App-ID based rules (App-ID rules configured with unused applications). Policy Optimizer helps you prioritize which port-based rules to migrate first, identify application-based rules that allow applications you don't use, and analyze rule usage characteristics such as hit count.

Converting port-based rules to application-based rules improves your security posture because you select the applications you want to allow and deny all other applications, so you eliminate unwanted and potentially malicious traffic from your network. Combined with restricting application traffic to its default ports (set the Service to **application-default**), converting to application-based rules also prevents evasive applications from running on non-standard ports.

You can use this feature on:

- Firewalls that run PAN-OS version 9.0 and have App-ID enabled.
- Panorama running PAN-OS version 9.0. You don't have to upgrade firewalls that Panorama manages to use the **Policy Optimizer** capabilities. However, to use the **Rule Usage** capabilities ([Monitor Policy Rule Usage](#)), managed firewalls must run PAN-OS 8.1 or later. If managed firewalls connect to Log Collectors, those Log Collectors must also run PAN-OS version 9.0. Managed PA-7000 Series firewalls that have a Log Processing Card (LPC) can also run PAN-OS 8.1 (or later).



PA-7000 Series Firewalls support two logging cards, the PA-7000 Series Firewall Log Processing Card (LPC) and the high-performance PA-7000 Series Firewall Log Forwarding Card (LFC). Unlike the LPC, the LFC does not have disks to store logs locally. Instead, the LFC forwards all logs to one or more external logging systems, such as Panorama or a syslog server. If you use the LFC, the application usage information for Policy Optimizer does not display on the firewall because traffic logs aren't stored locally. If you use the LPC, the traffic logs are stored locally on the firewall, so the application usage information for Policy Optimizer displays on the firewall.

Use this feature to:

- **Migrate port-based rules to application-based rules**—Instead of combing through traffic logs and manually mapping applications to port-based rules, use Policy Optimizer to identify port-based rules and list the applications that matched each rule, so you can select the applications you want to allow and safely enable them. Converting your legacy port-based rules to application-based allow rules supports your business applications and enables you to block any applications associated with malicious activity.
- **Identify over-provisioned application-based rules**—Rules that are too broad allow applications you don't use on your network, which increases the attack surface and the risk of inadvertently allowing malicious traffic.



Remove unused applications from Security policy rules to reduce the attack surface and keep the rulebase clean. Don't allow applications that nobody uses on your network.



To migrate a configuration from a legacy firewall to a Palo Alto Networks device, see [Best Practices for Migrating to Application-Based Policy](#).

You can't sort Security policy rules in **Security > Policies** because sorting would change the rule order in the rulebase. However, under **Polices > Security > Policy Optimizer**, Policy Optimizer provides sorting options

that don't affect the rule order to help you prioritize which rules to convert or clean up first. You can sort rules by the amount of traffic during the past 30 days, the number of applications seen on the rule, the number of days with no new applications, and the number of applications allowed (for over-provisioned rules).

You can use Policy Optimizer in other ways as well, including validating pre-production rules and troubleshooting existing rules. Note that Policy Optimizer honors only **Log at Session End** and ignores **Log at Session Start** to avoid counting transient applications on rules.



Due to resource constraints, VM-50 Lite virtual firewalls don't support Policy Optimizer.

- [Policy Optimizer Concepts](#)
- [Migrate Port-Based to App-ID Based Security Policy Rules](#)
- [Rule Cloning Migration Use Case: Web Browsing and SSL Traffic](#)
- [Add Applications to an Existing Rule](#)
- [Identify Security Policy Rules with Unused Applications](#)
- [High Availability for Application Usage Statistics](#)
- [How to Disable Policy Optimizer](#)

Policy Optimizer Concepts

Review the following topics to learn more about this feature's support:

- [Sorting and Filtering Security Policy Rules](#)
- [Clear Application Usage Data](#)

Sorting and Filtering Security Policy Rules

You can filter Security policy rules to see all the port-based rules, which have no applications configured (**Policies > Security > Policy Optimizer > No App Specified**). You can also filter to see all the rules that have applications configured but traffic doesn't hit all of the applications (**Policies > Security > Policy Optimizer > Unused Apps**). You can sort the filtered policy rules based on different types of statistics to help prioritize which rules to convert from port-based to application-based rules or to clean up first.



*You can't filter or sort rules in **Policies > Security** because that would change the order of the policy rules in the rulebase. Filtering and sorting **Policies > Security > Policy Optimizer > No App Specified** and **Policies > Security > Policy Optimizer > Unused Apps** does not change the order of the rules in the rulebase.*

You can click several column headers to sort port-based rules (**No App Specified**) and rules with unused applications (**Unused Apps**) based on application usage statistics. In addition, you can [View Policy Rule Usage](#) to help identify and remove unused rules to reduce security risks and keep your policy rule base organized. Rule usage tracking allows you to quickly validate new rule additions and rule changes and to monitor rule usage for operations and troubleshooting tasks.

No App Specified
These are security policies that have Application configured as "any." Such policies present a security risk as they allow access to any applications or protocol on the configured service. To safely enable applications in your environment, Palo Alto Networks strongly recommends that you convert these policies to application-based security policies.

Name	Action	Traffic Bytes (30 days)	Apps Allowed	Apps Seen	Days with No New Apps	Log Status	Created	Last Updated
10 allow-esp	allow	3.0K B	any	36	0	Complete	2018-12-01 10:54:00	2018-12-01 10:54:00
11 traffic-to-internal	deny	441.0M B	any	36	0	Complete	2018-12-01 10:52:30	2018-12-01 10:52:30
12 allow-https	allow	452.0K B	any	1	10	Complete	2018-12-01 10:57:40	2018-12-01 10:57:40
13 any	deny	42.0K B	any	1	0	Complete	2018-12-01 10:59:20	2018-12-01 10:59:20

- **Traffic (Bytes, 30 days)**—The amount of traffic seen on the rule over the last 30 days. The 30-day window places rules that *currently* match the most traffic at the top of the list by default (a longer time frame places more emphasis on older rules that would remain at the top of the list because they have large cumulative totals even though they may no longer see much traffic). Click to reverse the order.
- **Apps Seen**—Place the rules with the most or least applications seen at the top. The firewall never automatically purges the application data.



The firewall updates Apps Seen approximately every hour. However, if there is a large volume of application traffic or a large number of rules, it may take longer than an hour to update. After you add an application to a rule, wait at least an hour before running Traffic logs to see the application's log information.

- **Days with No New Apps**—Place the rules with the most or least days since the last new application matched the rule at the top.
- **(Unused Apps only) Apps Allowed**—Place the rules with the most or least applications configured on the rule at the top.

Application usage statistics only count applications for rules that meet the following criteria:

- The rule's Action must be **Allow**.
- The rule's Log Setting must be **Log at Session End** (this is the default Log Setting). Rules that **Log at Session Start** are ignored to prevent counting transient applications.
- Valid traffic must match the rule. For example, if the session ends before enough traffic passes through the firewall to identify the application, it is not counted. The following traffic types are not valid and therefore don't count for Policy Optimizer statistics:
 - Insufficient-data
 - Not-applicable
 - Non-syn-tcp
 - Incomplete

You can filter the Traffic logs (**Monitor > Logs > Traffic**) to see traffic identified as one of these types. For example, to see all traffic identified as incomplete, use the filter (**app eq incomplete**).

If these criteria aren't met, the application isn't counted for statistics such as **Apps Seen**, doesn't affect statistics such as **Days with No New Apps**, and doesn't appear in lists of applications.



The firewall doesn't track application usage statistics for the interzone-default and intrazone-default Security policy rules.



If the UUID of a rule changes, the application usage statistics for that rule reset because the UUID change makes the firewall see the rule as a different (new) rule.

The screenshot displays the Palo Alto Networks Security Policy configuration interface. The left sidebar shows the navigation menu with 'Security' selected. The main content area is titled 'No App Specified' and contains a warning message. Below the warning is a table of security policies.

Name	Source	Destination	Traffic (Bytes, M Bytes)	App Allowed	App Name	Open with New Pane	Enforced	Modified	Created
allow-https	any	any	2,000 B	any	any	0	Enforced	2018-12-05 20:56:55	2018-12-05 20:56:55
traffic-to-internal	any	any	443,200 B	any	any	1	Enforced	2018-12-05 15:52:51	2018-11-05 16:52:51
web-external	any	any	452,100 B	any	any	15	Enforced	2018-12-04 19:07:46	2018-11-04 19:07:46
any-https	any	any	42,500 B	any	any	1	Enforced	2018-12-05 06:30:51	2018-12-05 06:30:51
any-https	any	any	42,500 B	any	any	1	Enforced	2018-12-05 06:30:51	2018-12-05 06:30:51

Applications & Usage - Traffic to Internet

Timeframe: Anytime

Apps on Rule: 79

Applications

Application	Subcategory	Rank	First Seen	Last Seen	Traffic (30 days)
all	encrypted-tunnel	4	2019-01-11	2019-05-02	3,462
google-base	internet-utility	8	2018-12-03	2019-05-02	236,944
slack-base	social-business	2	2019-01-11	2019-05-02	181,044
traps-management-service	management	1	2018-12-18	2019-05-02	138,294
crashplan	storage-backup	5	2018-12-18	2019-04-30	74,634
web-browsing	internet-utility	6	2019-01-11	2019-05-02	54,204
youtube-base	photo-video	6	2018-12-18	2019-05-02	34,384
google-docs-base	office-productivity	7	2019-03-25	2019-05-02	33,884

The last new app was discovered 1 days ago.

OK Cancel

- **Applications**—Alphabetical by application name. If you configure specific ports or port ranges for a rule's Service (the Service cannot be **any**), and there are standard (application default) ports for the application, and the configured ports don't match the application-default ports, then a yellow, triangular warning icon appears next to the application.
- **Subcategory**—Alphabetical by application subcategory, derived from the application content metadata.
- **Risk**—According to the risk rating of the application.
- **First Seen**—The first day the application was seen on the rule. The time stamp resolution is by the day only (not hourly).
- **Last Seen**—The last day the application was seen on the rule. The time stamp resolution is by the day only (not hourly).

- **Traffic (30 days)**—Traffic in bytes that matched the rule over the last 30 days is the default sorting method.

Set the **Timeframe** to display statistics for a particular time period—**Anytime**, the **Past 7 days**, the **Past 15 days**, or the **Past 30 days**.



Traffic (30 days) always displays only the last 30 days of traffic in bytes. Changing the Timeframe does not change the duration of the Traffic (30 days) bytes measurement.

Clicking the column header orders the display and clicking the same column again reverses the order. For example, click **Risk** to sort applications from low risk to high risk. Click **Risk** again to sort applications from high risk to low risk.

The firewall doesn't report application usage statistics in real time on **Policies > Security > Policy Optimizer > No App Specified**, **Policies > Security > Policy Optimizer > Unused Apps**, or on **Applications & Usage**, so this feature isn't a replacement for running reports.

- The firewall updates **Apps Allowed**, **Apps Seen**, and the applications listed in **Applications & Usage** approximately every hour, not in real time. If there is a large amount of traffic or a large number of rules, updates may take longer. After you add an application to a rule, wait at least an hour before running Traffic logs to see the application's log information.

The firewall updates **Apps Seen** approximately every hour. However, if there is a large volume of application traffic or a large number of rules, it may take longer than an hour to update. After you add an application to a rule, wait at least an hour before running Traffic logs to see the application's log information.

- The firewall updates **Days with No New Apps** and also **First Seen** and **Last Seen** on **Applications & Usage** once per day, at midnight device time.
- For rules with large numbers of applications seen, it may take longer to process application usage statistics.
- For Security policy rulebases with large numbers of rules that have many applications, it may take longer to process application usage statistics.
- For firewalls managed by Panorama, application usage data is visible only for rules Panorama pushes to the firewalls, not for rules configured locally on individual firewalls.


Clear Application Usage Data

You can use a CLI command to clear application usage data for an individual Security policy rule and reset **Apps Seen** and other application usage data.

STEP 1 | Find the UUID of the Security policy rule whose application usage data you want to clear.

There are two ways to find the UUID in the UI:

- In **Policies > Security**, copy the UUID from the **Rule UUID** column.
- In **Policies > Security**, select **Copy UUID** in the rule **Name** drop-down menu.



STEP 2 | Switch from the UI to the CLI.

Use the UUID you captured in the UI to clear the rule's application usage data:

```
admin@PA-VM>clear policy-app-usage-data ruleuuid <uuid-value>
```

Paste or type the rule's UUID as the value and execute the command to clear the rule's application usage data.

Migrate Port-Based to App-ID Based Security Policy Rules

When you transition from a legacy firewall to a Palo Alto Networks next-generation firewall, you inherit a large number of port-based rules that allow any application on the ports, which increases the attack surface because any application can use an open port. Policy Optimizer identifies all applications seen on any legacy port-based Security policy rule and provides an easy workflow for selecting the applications you want to allow on that rule. Migrate port-based rules to application-based allow rules to reduce the attack surface and safely enable applications on your network. Use Policy Optimizer to maintain the rulebase as you add new applications.



Migrate a few port-based rules at a time to application-based rules, in a prioritized manner. A gradual conversion is safer than migrating a large rulebase at one time and makes it easier to ensure that new application-based rules control the necessary applications. Use Policy Optimizer information to prioritize which rules to convert first.



To migrate a configuration from a legacy firewall to a Palo Alto Networks device, see [Best Practices for Migrating to Application-Based Policy](#).

STEP 1 | Identify port-based rules.

Port-based rules have no configured (allowed) applications. **Policies > Security > Policy Optimizer > No App Specified** displays all port-based rules (**Apps Allowed** is **any**).

Name	Source	Traffic (Bytes, 30 days)	Apps Seen	Apps Seen	Click with No New Apps	Compare	Hit Count	Last Hit	First Hit	Modified	Created
11	any	1,101 B	any	0	0	Consistent	82333	2018-12-09 11:30:45	2018-12-09 11:30:45	2018-12-09 10:54:05	2018-12-09 10:54:05
12	any	348,091 B	any	0	0	Consistent	30762	2018-12-09 11:30:51	2018-12-09 11:30:51	2018-12-09 10:54:05	2018-12-09 10:54:05
13	any	17,000 B	any	1	1	Consistent	280	2018-12-09 11:30:51	2018-12-09 11:30:51	2018-12-09 10:54:05	2018-12-09 10:54:05
14	any	823,000 B	any	1	1	Consistent	8	2018-12-09 11:30:51	2018-12-09 11:30:51	2018-12-09 10:54:05	2018-12-09 10:54:05

STEP 2 | Prioritize which port-based rules to convert first.

Policies > Security > Policy Optimizer > No App Specified enables you to [sort rules](#) without affecting their order in the rulebase and provides other information that helps you prioritize rules for conversion based on your business goals and risk tolerance.

- **Traffic (Bytes, 30 days)**—(Click to sort.) Rules that *currently* match the most traffic are at the top of the list. This is the default sorting order.
- **Apps Seen**—(Click to sort.) A large number of legitimate applications matching a port-based rule may indicate you should replace it with multiple application-based rules that tightly define the applications, users, and sources and destinations. For example, if a port-based rule controls traffic for multiple applications for different user groups on different sets of devices, create separate rules that pair applications with their legitimate users and devices to reduce the attack surface and increase visibility. (Clicking the **Apps Seen** number or **Compare** shows you the applications that have matched the rule.)



The firewall updates Apps Seen approximately every hour. However, if there is a large volume of application traffic or a large number of rules, it may take longer than an hour to update. After you add an application to a rule, wait at least an hour before running Traffic logs to see the application's log information.


- **Days with No New Apps**—(Click to sort.) When the applications seen on a port-based rule stabilize, you can be more confident the rule is mature, conversion won't accidentally exclude legitimate applications, and no more new applications will match the rule. The **Created** and **Modified** dates help you evaluate a rule's stability because older rules that have not been modified recently may also be more stable.
- **Hit Count**—Displays rules with the most matches over a selected time frame. You can exclude rules for which you reset the hit counter and specify the exclusion time period in days. Excluding rules with recently reset hit counters prevents misconceptions about rules that show fewer hits than you expect because you didn't know the counter was reset.



You can also use Hit Count to [View Policy Rule Usage](#) and help identify and remove unused rules to reduce security risks and keep your rulebase organized.

STEP 3 | Review the **Apps Seen** on port-based rules, starting with the highest priority rules.

On **No Apps Specified**, click **Compare** or the number in **Apps Seen** to open **Applications & Usage**, which lists applications that matched a port-based rule over a specified **Timeframe**, with each application's **Risk**, the date it was **First Seen**, the date it was **Last Seen**, and the amount of traffic over the last 30 days.



You can check **Applications seen** on port-based rules over the past 7, 15, or 30 days, or over the rule's lifetime (**Anytime**). For migrating rules, **Anytime** provides the most complete assessment of applications that matched the rule.

You can search and filter the **Apps Seen**, but keep in mind that it takes an hour or more to update **Apps Seen**. You can also order the **Apps Seen** by clicking the column headers. For example, you can click **Traffic (30 days)** to bring the applications with the most recent traffic to the top of the list, or click **Subcategory** to organize the applications by subcategory.



The granularity of measurement for First Seen and Last Seen data is one day, so on the day you define a rule, the dates in these two columns are the same. On the second day the firewall sees traffic on an application, you'll see a difference in the dates.

STEP 4 | Clone or add applications to the rule to specify the applications you want to allow on the rule.

On **Applications & Usage**, convert a port-based rule to an application-based rule by:

- **Cloning the rule**—Preserves the original port-based rule and places the cloned application-based rule directly above it in the rulebase.
- **Adding applications to the rule**—Replaces the original port-based rule with the new application-based rule and deletes the original rule.



If you have existing application-based rules and you want to migrate applications to them from port-based rules, you can [Add Applications to an Existing Rule](#).



Some applications appear on the network at intervals, for example, for quarterly or yearly events. These applications may not display on the Applications & Usage screen if the history isn't long enough to capture their latest activity.




When you clone a rule or add applications to a rule, nothing else about the original rule changes. The original rule's configuration remains the same except for the applications you added to the rule. For example, if the original rule's Service allowed Any application or specified a particular service, you need to change the Service to Application-Default to restrict the allowed applications to their default ports on the new rule.

Cloning is the safest way to migrate rules, especially when **Applications & Usage** shows more than a few well-known applications matching the rule ([Rule Cloning Migration Use Case: Web Browsing and SSL Traffic](#) provides an example of this). Cloning preserves the original port-based rule and places it below the cloned application-based rule, which eliminates the risk of losing application availability because traffic that doesn't match the cloned rule flows through to the port-based rule. When traffic from legitimate applications hasn't hit the port-based rule for a reasonable period of time, you can remove it to complete that rule's migration.

To **clone** a port-based rule:

1. In **Apps Seen**, click the check box next to each application you want in the cloned rule. Keep in mind that it takes an hour or more to update **Apps Seen**.
2. Click **Create Cloned Rule**. In the **Create Cloned Rule** dialog, **Name** the cloned rule ("slack" in this example) and add other applications in the same container and application dependencies, if required. For example, to clone a rule by selecting the slack-base application:



The green text is the selected application to clone. The container application (**slack**) is in the gray row. The applications listed in *italics* are applications that have not been seen on the rule but are in the same container as the selected application. Individual applications that have been seen on the rule are in normal font. All the applications are included in the cloned rule by default (**Add Container App**, which adds all the applications in the container, is selected by default) to help prevent the rule from breaking in the future.

3. If you want to allow all of the applications in the container, leave **Add container app** selected. This also "future proofs" the rule because when an application is added to the container app, it's automatically added to the rule.

If you want to constrain access to some of the individual applications in the container, uncheck the box next to each individual application you don't want users to access. This also unchecks the container app, so if you want to allow new applications in the container later, you have to add those applications individually.

If you uncheck the container app, all the apps are unchecked and you manually select the apps you want to include in the cloned rule.

4. Leave the application dependencies checked, in this example, **ssl** and **web-browsing** (these are applications that the selected application requires).
5. Click **OK** to add the new application-based rule directly above the port-based rule in the rulebase.
6. **Commit** the configuration.

When you clone a rule and **Commit** the configuration, the applications you select for the cloned rule are removed from the original port-based rule's **Apps Seen** list. For example, if a port-based rule has 16 **Apps Seen** and you select two individual applications and one dependent application for the cloned rule, after cloning, the port-based rule shows 13 **Apps Seen** because the three selected applications have been removed from the port-based rule ($16 - 3 = 13$). The cloned rule shows the three added applications in **Apps on Rule**.

Creating a cloned rule with a container app works a bit differently. For example, a port-based rule has 16 **Apps Seen** and you select one individual application and a container app for the cloned rule. The container app has five individual applications and has one dependent application. After cloning, the cloned rule shows seven **Apps on Rule**—the individual application, the five individual applications in the container app, and the dependent application for the container app. However, in the original port-based rule, **Apps Seen** shows 13 applications because only the individual application, the container app, and the container app's dependent application are removed from the port-based rule.


In contrast to cloning, adding applications to a port-based rule replaces the rule with the resulting application-based rule. Adding applications to a rule is simpler than cloning, but riskier because you may inadvertently miss applications that should be on the rule, and the original port-based rule is no longer in the rulebase to catch accidental omissions. However, adding applications to port-based rules that apply to only a few well-known applications migrates the rule quickly to an application-based rule. For example, for a port-based rule that only controls traffic to TCP port 22, the only legitimate application is SSH, so it's safe to add applications to the rule.



*Adding applications using the traditional Security policy rule's Application tab does not change Apps Seen or Apps on Rule. To preserve accurate application usage information, when replacing port-based rules with application-based rules, add applications using **Add to This Rule** or **Match Usage** (or create a cloned rule or add applications to an existing application-based rule instead) in Apps Seen.*

There are three ways to replace a port-based rule with an application-based rule by adding applications (**Add to This Rule** and **Match Usage** in **Apps Seen** and **Add** in **Apps on Rule**):

- **Add to This Rule** applications from **Apps Seen** (applications that matched the rule). Keep in mind that it takes an hour or more to update **Apps Seen**.
 1. Select applications from **Apps Seen** on the rule.
 2. Click **Add to Rule**. In the **Add to Rule** dialog, add other applications in the same container app and application dependencies, if required. For example, to add slack-base to a rule:



Similar to the **Create Cloned Rule** dialog, the green text in **Add to This Rule** is the selected application to add to the rule. The container app (**slack**) is in the gray row. The applications listed in *italics* are applications that have not been seen on the rule but are in the same container as the selected application. Individual applications that have been seen on the rule are in normal font. All the applications are included in the cloned rule by default (**Add Container App**, which adds all the applications in the container, is selected by default) to help prevent the rule from breaking in the future.

3. If you want to allow all of the applications in the container, leave **Add container app** selected. This also “future proofs” the rule because when an application is added to the container app, it’s automatically added to the rule.

If you want to constrain access to some of the individual applications in the container, uncheck the box next to each individual application you don’t want users to access. This also unchecks the container app, so if you want to allow new applications in the container later, you have to add those applications individually.

If you uncheck the container app, all the apps are unchecked and you manually select the apps you want to include in the cloned rule.

4. Leave the application dependencies checked, in this example, **ssl** and **web-browsing** (these are applications that the selected application requires).
5. Click **OK** to replace the port-based rule with the new application-based rule.

When you **Add to This Rule** and **Commit** the configuration, the applications you didn’t add are removed from **Apps Seen** because the new application-based rule no longer allows them. For example, if a rule has 16 **Apps Seen** and you **Add to This Rule** three applications, the resulting new rule shows only those three added applications in **Apps Seen**.

Add to This Rule with a container app works a bit differently. For example, a port-based rule has 16 **Apps Seen** and you select one individual application and a container app to add to the new rule. The container app has five individual applications and has one dependent application. After adding the applications to the rule, the new rule shows seven **Apps on Rule**—the individual application, the five individual applications in the container app, and the dependent application for the container app. However, **Apps Seen** shows 13 applications because the individual application, the container app, and the container app’s dependent application are removed from that list.

- Add all of the **Apps Seen** on the rule to the rule at one time with one click (**Match Usage**).



Port-based rules allow any application, so Apps Seen may include unneeded or unsafe applications. Use Match Usage to convert a rule only when the rule has seen a small number of well-known applications with legitimate business purposes. A good example is TCP port 22, which should only allow SSH traffic, so if SSH is the only application seen on a port-based rule that opens port 22, you can safely Match Usage.

1. In **Apps Seen**, click **Match Usage**. Keep in mind that it takes an hour or more to update **Apps Seen**. All the applications in **Apps Seen** are copied to **Apps on Rule**.
 2. Click **OK** to create the application-based rule and replace the port-based rule.
- If you know the applications you want on the rule, you can **Add** applications manually in **Apps on Rule**. However, this method is equivalent to using the traditional Security policy rule **Application** tab and does not change **Apps Seen** or **Apps on Rule**. To preserve accurate application usage information, convert rules using **Add to Rule**, **Create Cloned Rule**, or **Match Usage** in **Apps Seen**.
1. In **Apps on Rule**, **Add** (or **Browse**) and select applications to add to the rule. This is equivalent to adding applications on the **Application** tab.
 2. Click **OK** to add the applications to the rule and replace the port-based rule with the new application-based rule.



Because this method is equivalent to adding applications using the Application tab, the dialog to add application dependencies doesn't pop up.

STEP 5 | For each application-based rule, set the **Service** to **application-default**.



If business needs require you to allow applications (for example, internal custom applications) on non-standard ports between particular clients and servers, restrict the exception to only the required application, sources, and destinations. Consider rewriting custom applications so they use the application default port.

STEP 6 | **Commit** the configuration.

STEP 7 | Monitor the rules.

- **Cloned rules**—Monitor the original port-based rule to ensure the application-based rule matches the desired traffic. If applications you want to allow match the port-based rule, add them to the application-based rule or clone another application-based rule for them. When only applications that you don't want on your network match the port-based rule for a reasonable period of time, the cloned rule is robust (it catches all the application traffic you want to control) and you can safely remove it.
- **Rules with Added Applications**—Because you convert only port-based rules that have a few well-known applications directly to application-based rules, in most cases the rule is solid from the start. Monitor the converted rule to see if the expected traffic matches the rule—if there's less traffic than expected, the rule may not allow all of the necessary applications. If there's more traffic than expected, the rule may allow unwanted traffic. Listen to user feedback—if users can't access applications they need for business purposes, the rule (or another rule) may be too tight.

Rule Cloning Migration Use Case: Web Browsing and SSL Traffic

A port-based rule that allows web access on TCP ports 80 (HTTP web-browsing) and 443 (HTTPS SSL) provides no control over which applications use those open ports. There are many web applications, so a general rule that allows web traffic allows thousands of applications, many of which you don't want on your network.

This use case shows how to migrate a port-based policy that allows all web applications to an application-based policy that allows only the applications you want, so you can safely enable the applications you choose to allow. For rules that see a lot of applications, cloning the original port-based rule is safer than adding applications to the rule because adding replaces the port-based rule, so if you inadvertently forget to add a critical application, you affect application availability. And if you **Match Usage**, which also replaces the port-based rule, you allow all of the applications the rule has seen, which could be dangerous, especially with web browsing traffic.

Cloning the rule retains the original port-based rule and places the cloned rule directly above the port-based rule in the rulebase, so you can monitor the rules. Cloning also allows you to split rules that see a lot of different applications—such as a port-based web traffic rule—into multiple application-based rules so you can treat different groups of applications differently. When you're sure you're allowing all the applications you need to allow in the cloned rule (or rules), you can remove the port-based rule.

This example clones a port-based web traffic rule to create an application-based rule for social networking traffic (a subset of the application traffic seen on the port-based rule).

STEP 1 | Navigate to **Policies > Security > Policy Optimizer > No App Specified** to view the port-based rules.

STEP 2 | Click **Compare** for the rule you want to migrate.

In this example, the port-based rule that allows web access is named Traffic to internet.




				App Usage				
Name	Service	Traffic (Bytes, 30 days)	Apps Allowed	Apps Seen	Days with No New Apps	Compare	Modified	Created
11 Traffic to internet	web	5.7K	any	12	14	Compare	2018-12-18 13:32:33	2018-11-16 10:51:30
10 any	any	27.5M	any	3	14	Compare	2018-12-18 16:50:33	2018-12-18 16:50:33
2 sub-access	any	652.0K	any	1	29	Compare	2018-12-06 13:07:48	2018-11-14 15:59:25
13 rule1	any	18.4K	any	113	25	Compare	2018-12-13 10:54:59	2018-07-24 11:58:40

STEP 3 | Use the [sorting options](#) to review and select the applications you want to allow from **Apps Seen**.




The number of Apps Seen is updated approximately every hour, so if you don't see as many applications as you expect, check again after about an hour. Depending on the firewall's load, it may take longer than one hour for these fields to update.

For example, click **Subcategory** to sort the applications, scroll to the social-networking subcategory, and then select the applications you want to allow.



STEP 4 | Click **Create Cloned Rule**.

Create Cloned Rule shows the selected applications shaded green, the container apps shaded gray, individual applications in the container that haven't been seen on the rule in *italics*, and individual applications that have been seen on the rule in normal text font. Scrolling through **Applications** shows all the container apps and their individual applications.



Create Cloned Rule also shows the application dependencies of the selected applications. Because we selected social applications but not web-browsing or ssl, web-browsing and ssl are among the listed application dependencies.

STEP 5 | Name the cloned rule (in this example, the name will be Social Networking Apps).

STEP 6 | Select the applications you want in the cloned rule.

For applications you don't want to include, uncheck the corresponding box, which also unchecks the container app. If you don't include the container app, then when new apps are added to the container, they won't automatically be added to the rule.

If you uncheck the container app, all the individual applications in the container are unchecked and you must select the apps you want to add manually.

STEP 7 | Click OK to create the cloned rule.

STEP 8 | In Policies > Security, the cloned rule (Social Networking Apps) is inserted in the rulebase above the original port-based rule (Traffic to internet).



STEP 9 | Click the rule name to edit the cloned rule, which inherits the properties of the original port-based rule.

STEP 10 | On the Service/URL Category tab, delete service-http and service-https from **Service**.

This changes the **Service** to **application-default**, which prevents applications from using non-standard ports and further reduces the attack surface.



If business needs require you to allow applications (for example, internal custom applications) on non-standard ports between particular clients and servers, restrict the exception to only the required application, sources, and destinations. Consider rewriting custom applications so they use the application default port.

STEP 11 | On the Source, User, and Destination tabs, tighten the rule to apply to only the right users in only the right locations (zones, subnets).

For example, you may decide to limit social media activity to certain marketing, public relations, sales, and executive groups.

STEP 12 | Click OK.

STEP 13 | Commit the configuration.

STEP 14 | Repeat the process for other application categories in the port-based web access rule until your application-based rules allow only the applications you want to allow on your network.

When traffic you want to allow stops hitting the original port-based rule for a sufficient amount of time to be confident that the port-based rule is no longer needed, you can remove the port-based rule from the rulebase.


Add Applications to an Existing Rule

In some cases, you may want to add applications learned (seen) on a port-based rule to an application-based rule that already exists. For example, an administrator may create a cloned application-based rule for file-sharing applications from a port-based rule that allows internet access (a port 80/443 rule). A few days later, the administrator notices that the port-based internet access rule has seen more file-sharing applications and wants to add some or all of them to the cloned application-based rule because cloning another application-based rule for the same type of application would create an unnecessary rule and complicate the rulebase.


This example uses file-sharing applications to show you how to add applications to an existing rule.

STEP 1 | You have already taken the following steps to clone an application-based rule from the port-based internet access rule so you can control file-sharing apps:

1. Clicked **Compare** (or the number in **Apps Seen**) in **Policies > Security > Policy Optimizer > No App Specified** and filtered the file-sharing applications.




2. Selected the desired file-sharing applications and created a cloned rule.




3. Changed the Service from **service-http** and **service-https** to **application-default**.


STEP 2 | You check the port-based internet access rule later and discover that more file-sharing applications you need to allow for your business have been seen on the rule.




STEP 3 | Select the file-sharing apps you want to add to the existing rule.



STEP 4 | Click **Add to Existing Rule** and select the **Name** of the rule to which you want to add the applications, in this case, **file-sharing-apps**.



STEP 5 | Click **OK** to add the selected applications to the **file-sharing-apps** rule.



STEP 6 | The updated rule now controls the original cloned file-sharing applications and the applications you just added.



Identify Security Policy Rules with Unused Applications

If you have application-based Security policy rules that allow a large number of applications, you can remove unused applications (applications never seen on the rules) from those rules to tighten them so that they only allow applications actually seen in the rule's traffic. Identifying and removing unused applications from Security policy rules is a best practice that strengthens your security posture by reducing the attack surface.

STEP 1 | Identify Security policy rules that have unused applications.

Policies > Security > Policy Optimizer > Unused Apps displays all application-based rules configured with applications that aren't in use on the network—that is, rules configured to allow more applications than the number of applications actually seen on those rules.



The number of Apps Allowed and Apps Seen are updated approximately every hour, so if you configure applications on a rule and don't see as many Apps Allowed as you expect, check again after about an hour. Depending on the firewall's load, it may take longer than one hour for these fields to update.

STEP 2 | Prioritize which rules with unused applications to modify first.

Policies > Security > Policy Optimizer > Unused Apps enables you to [sort rules](#) without affecting their order in the rulebase and provides other information that helps you prioritize rules to clean up based on your business goals and risk tolerance.

- The difference between **Apps Allowed** (the number of applications the allow rule allows) and **Apps Seen** (the number of allowed applications actually seen on the rule) shows how many applications are configured on each rule but not actually seen on the rule, which indicates to what extent the rule is over-provisioned. Click **Apps Allowed** to sort by the number of applications allowed in a rule and click **Apps Seen** to sort by the number of applications actually seen on a rule.
- **Days with No New Apps** (click to sort) shows you the number of days since the last time a new application hit the rule. This indicates how likely it is that the rule is mature and won't see any applications that haven't already been seen. The longer the **Days with No New Apps**, the less likely that new applications will hit the rule and the more likely that you know all the applications the rule allows.
- **Created** and **Modified** dates also help determine whether a rule has matured enough to understand whether applications not seen on the rule may be seen at a later date or if the rule has seen all the applications expected to hit the rule. The longer the time since a rule was **Modified**, the more likely the rule is mature. (If **Created** and **Modified** are the same, the rule hasn't been modified.)
- **Hit Count**—Displays rules with the most matches over a selected time frame. You can exclude rules for which you reset the hit counter and specify the exclusion time period in days. Excluding rules with recently reset hit counters prevents misconceptions about rules that show fewer hits than you expect because you didn't know the counter was reset.




You can also use Hit Count to [View Policy Rule Usage](#).

You can also click **Traffic (Bytes, 30 days)** to sort by the amount of traffic a rule has seen over the last 30 days. Use this information to prioritize which rules to modify first. For example, you can prioritize rules with the largest difference between **Apps Allowed** and **Apps Seen** and that also have the most **Days with No New Apps**, because those rules have the greatest number of unused applications and are the most mature.

STEP 3 | Review the **Apps Seen** on the rule.

On **Unused Apps**, click **Compare** or the number in the **Apps Seen** column to open **Applications & Usage**, which shows the applications configured on the rule (**Apps on Rule**) and the **Apps Seen** on the rule.



- The number next to **Apps Seen** (12 in this example) is the number of applications actually seen on the rule. Keep in mind that it takes at least one hour to update **Apps Seen**.
- The number next to **Apps on Rule** (32 in this example) is how many applications are configured on the rule, which is calculated by counting each application in a container app (but not the container app itself—if you configure a container app on the rule, the rule allows the container app's individual applications). However, the **Applications** list shows only the applications manually configured on the rule, so when you configure a container app on a rule, **Applications** only shows the container app, not the individual applications in the container (unless you also manually configure the individual applications on the rule). For this reason, the number of **Apps on Rule** may not be the same as the number of applications in the **Applications** list.
- The pink-shaded **Apps on Rule** are either:
 - Applications configured on the rule but no matching traffic has been seen on the rule, so the application hasn't been seen on the rule.
 - Container apps, even when their individual applications have been seen on the rule. For example, if a rule configured with the container app active-directory sees traffic from the active-directory-base individual application, the container app is shaded pink even though one of the container app's individual applications was seen on the rule.

- The unshaded **Apps on Rule** are applications that have been seen on the rule.

Pink-shaded applications configured on the rule may be unused applications, although pink-shaded container apps may have individual applications in use. Consider removing pink-shaded applications to reduce the attack surface. Take into account periodically used applications, such as for quarterly or annual events, which may look unused if you don't examine a long enough time frame. **Timeframe** allows you to select the time frame for the **Apps Seen** on the rule. Select **Anytime** to see every application seen over the life of the rule. Depending on the **Created** or **Modified** date in the **No App Specified** dialog and the time between periodic events, the rule may not have been on the firewall long enough to see all periodically used applications.

STEP 4 | Remove unused applications from the rule.

Delete (or **Add**) applications in **Apps on Rule** to remove (or add) applications manually, or **Match Usage** to add the **Apps Seen** on the rule and delete applications for which no matching traffic has been seen on the rule with one click.

To remove applications from the rule manually, select the pink-shaded applications from **Apps on Rule** and **Delete** them. Ensure that none of the pink-shaded applications are required for periodic events before you remove them from the rule. (You can still add or delete applications on the Security policy rule's **Application** tab, but unused applications aren't shaded pink.)

Match Usage moves the **Apps Seen** on the rule to **Apps on Rule** and removes all unused (pink-shaded) applications from the rule.



You can clone rules from Policies > Security and from No App Specified to [Migrate Port-Based to App-ID Based Security Policy Rules](#). You can't clone a rule starting from Unused Apps.

STEP 5 | **Commit** the configuration.

STEP 6 | Monitor updated rules and listen to user feedback to ensure that updated rules allow the applications you want to allow and don't inadvertently block periodically used applications.



The number of Apps Allowed and Apps Seen are updated approximately every hour. After you remove all of the unused applications from a rule, the rule remains listed in Policies > Security > Policy Optimizer > Unused Apps until the firewall updates the display. When the firewall updates the display and the number of Apps Allowed is the same as the number of Apps Seen, the rule no longer displays in the Unused Apps screen. However, depending on the firewall's load, it may take longer than one hour for these fields to update.

High Availability for Application Usage Statistics

When you configure two firewalls as a High Availability (HA) pair, the application usage statistics are local to the firewall that generates the Traffic logs for the application. Where you can view application usage statistics also depends in part on the HA configuration:

- **Active/Passive**—The active device generates the application usage statistics. If a passive device has seen no user traffic, then only the active device displays the application usage statistics. If a passive device has seen traffic, then the passive device only displays the application usage statistics from the traffic that it has seen.

On a failover, the application usage statistics are based only on the Traffic logs generated on the newly active device (the device that was passive before the failover).

- **Active/Active**—The device that owns a session generates the Traffic logs for that session, so the application usage statistics for a session are only available on the device that owns the session. If one

active device owns a session, the other active device does not display that session's application usage statistics.

How to Disable Policy Optimizer

Policy Optimizer is enabled by default. Policy Optimizer provides many capabilities that make it easier to [Migrate Port-Based to App-ID Based Security Policy Rules](#) and to [Identify Security Policy Rules with Unused Applications](#) and remove the unused applications from the rules, but if you wish to disable the feature, you can.

STEP 1 | Navigate to **Device > Setup > Management > Policy Rulebase Settings**.

STEP 2 | Select the **Policy Application Usage** check box to enable the feature and deselect the check box to disable the feature.



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.



- *H.323 (H.225 and H.248) ALG is not supported in gatekeeper routed mode.*
- *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 \(ALG\)](#).*



When you use Dynamic IP and Port (DIPP) NAT, the Palo Alto Networks firewall ALG decoder needs a combination of IP and Port (Sent-by Address and Sent-by Port) under SIP headers (Contact and Via fields) to be able to translate the mentioned headers and open predict sessions based on them.

The following table lists IPv4, NAT, IPv6, NPTv6 and NAT64 ALGs and indicates with a check mark whether the ALG supports each protocol (such as SIP).

App-ID	IPv4	NAT	IPv6	NPTv6	NAT64
SIP	✓	✓	✓	—	—
SCCP	✓	✓	✓	—	—
MGCP	✓	✓	—	—	—
FTP	✓	✓	✓	✓	—
RTSP	✓	✓	✓	✓	—
MySQL	✓	✓	—	—	—
Oracle/SQLNet/ TNS	✓	✓	✓	✓	—
RPC	✓	✓	—	—	—

App-ID	IPv4	NAT	IPv6	NPTv6	NAT64
RSH	✓	✓	—	—	—
UNISlim	✓	✓	—	—	—
H.225	✓	✓	—	—	—
H.248	✓	✓	—	—	—

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.

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.

Use HTTP Headers to Manage SaaS Application Access

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.

You can use custom HTTP headers to disallow SaaS consumer accounts while allowing a specific enterprise account. Many SaaS applications allow or disallow access to applications based on information contained in specific HTTP headers. You can [Create HTTP Header Insertion Entries using Predefined Types](#) to manage access to popular SaaS applications, such as Google G Suite and Microsoft Office 365. Palo Alto Networks® uses content updates to maintain predefined rule sets specific to these applications, as well as to add new predefined rule sets.

You can also [Create Custom HTTP Header Insertion Entries](#) if you want to manage access to a SaaS application—that uses HTTP headers to limit service access—for which Palo Alto Networks has not provided a predefined set of rules.

Be aware that commercial SaaS applications always use SSL so decryption is necessary to perform HTTP header insertion. You can configure the firewall to decrypt traffic using SSL Forward Proxy decryption if traffic is not already decrypted by an upstream firewall.



You don't need a URL Filtering license to use this feature.

To understand how to use HTTP headers to manage SaaS applications, see the following:

- [Understand SaaS Custom Headers](#)
- [Domains used by the Predefined SaaS Application Types](#)
- [Create HTTP Header Insertion Entries using Predefined Types](#)
- [Create Custom HTTP Header Insertion Entries](#)

Understand SaaS Custom Headers

Before you begin, make sure you understand the custom HTTP headers you will use with the SaaS application you are managing. You need to understand what you can accomplish with these headers and the information you need to specify to accomplish your goals.

Be aware that SaaS applications that use custom headers do not always use them to control access to types of accounts. For example, Palo Alto Networks® provides predefined support for YouTube custom headers that determine whether network users can access restricted content.

You should also read the documentation for the SaaS application to which you want to control access so that you understand what headers you need to use for that application.



The following limits apply to HTTP header insertion:

- Header name character length: 100.
- Header value character length: 512.

Be aware that some SaaS applications might define custom header names, or assign values to their custom headers, that exceed these limits. These situations should be rare, but if a SaaS application does exceed one or both of these character length limits, then your next-generation firewall can not successfully manage access to that SaaS application.

The following table lists the headers that you can use for the SaaS applications for which Palo Alto Networks provides predefined support; each header also includes a link to more information specific to that header.

Application	Headers	For More Information
Dropbox	X-Dropbox-allowed-Team-Ids	https://www.dropbox.com/help/business/network-control You can allow access to sanctioned Enterprise Dropbox accounts. This header's value is the business account's team ID, which you can obtain from the network control section of the Dropbox admin console. You must also enable this functionality from the same location. For details on managing this header, as well as how to enable your Dropbox clients so that you can decrypt their traffic, contact your Dropbox account representative.
Google G Suite	X-GooGApps-Allowed-Domains	https://support.google.com/a/answer/1668854?hl=en You can allow access to specific Google accounts from your domain. The values that you give to this header are your domain and subdomains.
Microsoft Office 365	Restrict-Access-To-Tenants Restrict-Access-Context	https://docs.microsoft.com/en-us/azure/active-directory/active-directory-tenant-restrictions You provide Restrict-Access-To-Tenants with a list of tenants you want to allow your users to access. You can use any domain that is registered with a tenant to identify the tenant in this list. You provide Restrict-Access-Context with the directory ID that is setting the tenant restriction. You can find your directory ID in the Azure portal. Sign in as an administrator, select Azure Active Directory , then select Properties .
YouTube	YouTube-Restrict	https://support.google.com/a/answer/6214622?hl=en You provide this header with information on the type of videos you want your users to be able to view. You can specify either a Strict or Moderate setting. See https://support.google.com/a/answer/6212415 for details on these different settings.

Domains used by the Predefined SaaS Application Types

SaaS applications use HTTPS so, to insert custom headers into this traffic, custom headers must be decrypted. If you use the forward-proxy decryption available on the firewall to decrypt custom headers, you must identify the specific HTTPS traffic you want to decrypt by identifying the domains associated with the traffic. The following table identifies the relevant domains for each of the SaaS applications for which Palo Alto Networks® has provided predefined rules.

Application	Domains
Dropbox	*.dropbox.com
G Suite	*.google.com gmail.com
Microsoft Office 365	login.microsoftonline.com login.microsoft.com login.windows.net
YouTube	www.youtube.com m.youtube.com youtubei.googleapis.com youtube.googleapis.com www.youtube-nocookie.com

Create HTTP Header Insertion Entries using Predefined Types

STEP 1 | If there are no upstream devices already decrypting HTTPS traffic, configure [Decryption](#) using [Configure SSL Forward Proxy](#).



If you are configuring SSL decryption for Dropbox, then you must also configure your Dropbox clients to allow SSL traffic. These procedures are specific and private to Dropbox — to obtain these procedures, contact your Dropbox account representative.

1. **Add** a Custom URL Category for the SaaS application you are managing (**Objects > Custom Objects > URL Category**).
2. Specify a **Name** for the category.
3. **Add** the domains specific to the SaaS application you are managing. See [Domains used by the Predefined SaaS Application Types](#) for a list of the domains that you use for each of the predefined SaaS applications.
4. [Create a Decryption Policy Rule](#) and, as you follow this procedure, configure the following:
 - In the **Service/URL Category** tab, **Add** the **URL Category** that you created in the previous step.
 - In the **Options** tab, make sure the **Action** is set to **Decrypt** and that the **Type** is set to **SSL Forward Proxy**.

STEP 2 | Edit or add a [URL filtering profile](#).

STEP 3 | Select **HTTP Header Insertion** in the **URL Filtering Profile** dialog.

STEP 4 | Add an entry.

1. Specify a **Name** for this entry.
2. Select a predefined application **Type**.
This populates the **Domains** and **Headers** lists.
3. For each **Header**, enter a **Value**.
4. (Optional) Select **Log** to enable logging of insertion activity for the headers.
5. Click **OK** to save your changes.

STEP 5 | Add or edit a [Security Policy](#) rule (**Policies > Security**) that allows users to access the SaaS application for which you are configuring this header insertion rule.

1. Choose the URL filtering profile (**Actions > URL Filtering**) that you edited or created in Step 2.
2. Click **OK** to save and then **Commit** your changes.

STEP 6 | Verify that access to the SaaS application is working in the way you expect. From an endpoint:

1. Try to access an account or content that you expect to be able to access. If you cannot access the SaaS account or content, then the configuration is not working.
2. Try to access an account or content that you expect will be blocked. If you can access the SaaS account or content, then the configuration is not working.
3. If both of the previous steps work as expected, then you can [View Logs](#) (if you configured logging in step 4.4) and you should see the recorded HTTP header insertion activity.

Create Custom HTTP Header Insertion Entries

STEP 1 | If there are no upstream devices already decrypting HTTPS traffic, configure [Decryption](#) using [Configure SSL Forward ProxyDecryption](#).

1. **Add** a Custom URL Category for the SaaS application you are managing (**Objects > Custom Objects > URL Category**).
2. Specify a **Name** for the category.
3. **Add** the domains specific to the SaaS application you are managing.
4. [Create a Decryption Policy Rule](#) and, as you follow this procedure, configure the following:
 - In the **Service/URL Category** tab, **Add** the **URL Category** that you created in the previous step.
 - In the **Options** tab, make sure the **Action** is set to **Decrypt** and that the **Type** is set to **SSL Forward Proxy**.

STEP 2 | Edit or create a [URL filtering profile](#).

STEP 3 | Select **HTTP Header Insertion** in the **URL Filtering Profile** dialog.

STEP 4 | Add an entry.

1. Specify a **Name** for this entry.
2. Select **Custom** as the **Type**.
3. **Add** domains to the **Domains** list.

You can add up to 50 domains and each domain name can have up to 256 characters; wildcards are supported (for example, *.example.com).



HTTP header insertion occurs when a domain in this list matches the domain in the Host header of the HTTP request.

4. **Add** headers to the **Headers** list.

You can add up to 5 headers and each header can have up to 100 characters but cannot contain any spaces.

5. For each header **Value**.
6. (Optional) Select **Log** to enable logging of insertion activity for the headers.
7. Click **OK** to save your changes.

STEP 5 | Add or edit a [Security Policy](#) rule (**Policies > Security**) [Security Policy](#) that allows users to access the SaaS application for which you are configuring this header insertion rule.

1. Choose the URL filtering profile (**Actions > URL Filtering**) that you edited or created in Step 2.
2. Click **OK** to save and then **Commit** your changes.


STEP 6 | Verify that access to the SaaS application is working in the way you expect. From an endpoint that is connected to your network:

1. Try to access an account or content that you expect to be able to access. If you cannot access the SaaS account or content, then the configuration is not working.
2. Try to access an account or content that you expect will be blocked. If you can access the SaaS account or content, then the configuration is not working.
3. If both of the previous steps work as expected, then you can [View Logs](#) (if you configured logging in step 4.6) and you should see the recorded HTTP header insertion activity.

Maintain Custom Timeouts for Data Center Applications

Easily maintain custom timeouts for applications as you move from a port-based policy to an application-based policy. Use this method to maintain custom timeouts instead of overriding App-ID (losing application visibility) or creating a custom App-ID (expending time and research).

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.

The following steps describe how apply custom timeouts to applications; to apply custom timeouts to user groups, you can follow the same steps but just make sure to add the service object to the security policy rule that enforces the users to whom you want the timeout to apply.

STEP 1 | Select **Objects** > **Services** to add or modify a service object.

You can also create service objects as you are defining match criteria for a security policy rule: select **Policies** > **Security** > **Service/URL Category** and **Add** a new Service object to apply to the application traffic the rule governs.

STEP 2 | Select the protocol for the service to use (TCP or UDP).

STEP 3 | Enter the destination port number or a range of port numbers used by the service.

STEP 4 | Define the session timeout for the service.

- **Inherit from application** (default)—No service-based timeouts are applied; instead, apply the application timeout.
- **Override**—Define a custom session timeout for the service.

STEP 5 | If you chose to override the application timeout and define a custom session timeout, continue to:

- Enter a **TCP Timeout** value to set the Maximum length of time in seconds that a TCP session can remain open after data transmission has started. When this time expires, the session closes. The value range is 1 - 604800, and the default value is 3600 seconds.
- Enter a **TCP Half Closed** value to set the maximum length of time in seconds that a session remains in the session table between receiving the first FIN packet and receiving the second FIN packet or RST packet. If the timer expires, the session closes. The value range is 1 - 604800, and the default value is 120 seconds.
- Enter a **TCP Wait Time** value to set the maximum length of time in seconds that a session remains in the session table after receiving the second FIN packet or a RST packet. When the timer expires, the session closes. The value range is 1 - 600, and the default value is 15 seconds.

STEP 6 | Click **OK** to save the service object.

STEP 7 | Select **Policies > Security** and **Add** or modify a policy rule to govern the application traffic you want to control.

STEP 8 | Select **Service/URL Category** and **Add** the service object you just created to the security policy rule.

STEP 9 | Click **OK** and **Commit** your changes.

Threat Prevention

The Palo Alto Networks® next-generation firewall protects and defends your network from commodity threats and advanced persistent threats (APTs). The multi-pronged detection mechanisms of the firewall 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 antivirus, anti-spyware, and vulnerability protection features along with URL filtering and Application identification capabilities on the firewall.

Advanced threats are perpetuated by organized cyber adversaries who 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 together with WildFire™ and Panorama™ provide a comprehensive solution that intercepts and breaks the attack chain and provides visibility to prevent security infringement on your network infrastructure—both mobile and virtualized.



After you implement your threat prevention configurations, Export Configuration Table Data to create a PDF or CSV report of your configurations to use for internal review or for auditing.

- > Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions
- > Set Up Antivirus, Anti-Spyware, and Vulnerability Protection
- > DNS Security
- > Use DNS Queries to Identify Infected Hosts on the Network
- > Set Up Data Filtering
- > Predefined Data Filtering Patterns
- > Create a Data Filtering Profile
- > Set Up File Blocking
- > Prevent Brute Force Attacks
- > Customize the Action and Trigger Conditions for a Brute Force Signature
- > Enable Evasion Signatures
- > Prevent Credential Phishing
- > Monitor Blocked IP Addresses
- > Threat Signature Categories
- > Create Threat Exceptions
- > Custom Signatures
- > Learn More About and Assess Threats
- > Share Threat Intelligence with Palo Alto Networks
- > Threat Prevention Resources

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

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


- ❑ Upgrade to the most current PAN-OS software version and content release version to ensure that you have the latest security updates. See [Install Content and Software Updates](#).
- ❑ Set up the firewall to act as a DNS proxy and enable evasion signatures:

- [Configure a DNS Proxy Object](#).

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

- [Enable Evasion Signatures](#)

Evasion signatures that detect crafted HTTP or TLS requests can send alerts when clients connect to a domain other than the domain specified in the original DNS request. Make sure to configure DNS proxy before you enable evasion signatures. Without DNS proxy, evasion signatures can trigger alerts when a DNS server in the 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 allow only 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), create a new custom service that includes only port 587 and use that new service in your security policy rule instead of application-default. Additionally, make sure you restrict access to specific source and destinations zones and sets of IP addresses.
- ❑ Block all unknown applications and traffic using the Security policy. Typically, the only applications classified as unknown traffic are internal or custom applications on your network and potential threats. Unknown traffic can be either non-compliant applications or protocols that are anomalous or abnormal or it can be known applications that are using non-standard ports, both of which should be blocked. See [Manage Custom or Unknown Applications](#).

- ❑ **Set Up File Blocking** to block Portable Executable (PE) file types for internet-based SMB (Server Message Block) traffic from traversing trust to untrust zones (ms-ds-smb applications).



- ❑ 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**).



- Enable the drop **Mismatched overlapping TCP segment** option (**Packet Based Attack Protection > TCP Drop**).


By deliberately constructing connections with overlapping but different data in them, attackers attempt to cause misinterpretation of the intent of the connection and deliberately induce false positives or false negatives. Attackers also use IP spoofing and sequence number prediction to intercept a user's connection and inject their own data into that connection. Selecting the **Mismatched overlapping TCP segment** option specifies that PAN-OS discards frames with mismatched and overlapping data. Received segments are discarded when they are contained within another segment, when they overlap with part of another segment, or when they contain another complete segment.

- Enable the drop **TCP SYN with Data** and drop **TCP SYNACK with Data** options (**Packet Based Attack Protection > TCP Drop**).

Dropping SYN and SYN-ACK packets that contain data in the payload during a three-way handshake increases security by blocking malware contained in the payload and preventing it from extracting unauthorized data before the TCP handshake is completed.

- Strip TCP timestamps from SYN packets before the firewall forwards the packet (**Packet Based Attack Protection > TCP Drop**).

When you enable the **Strip TCP Options - TCP Timestamp** option, the TCP stack on both ends of the TCP connection will not support TCP timestamps. This prevents attacks that use different timestamps on multiple packets for the same sequence number.




- ❑ If you configure IPv6 addresses on your network hosts, be sure to enable support for IPv6 if not already enabled (**Network > Interfaces > Ethernet > IPv6**).

Enabling support for IPv6 allows access to IPv6 hosts and also filters IPv6 packets encapsulated in IPv4 packets, which 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**).



- ❑ Disable the options to **Forward datagrams exceeding UDP content inspection queue** and **Forward segments exceeding TCP content inspection queue** (**Device > Setup > Content-ID > Content-ID Settings**).

By default, when the TCP or UDP content inspection queues are full, the firewall skips content inspection for TCP segments or UDP datagrams that exceed the queue limit of 64. Disabling this option ensures content inspection for all TCP and UDP datagrams that the firewall allows. Only under specific

circumstances—for example, if the firewall platform is not sized appropriately to align with a use case—could disabling this setting impact performance.

- ❑ Disable the **Allow HTTP partial response** (**Device** > **Setup** > **Content-ID** > **Content-ID Settings**).

The HTTP partial response option allows a client to fetch only part of a file. When a next-generation firewall in the path of a transfer identifies and drops a malicious file, it terminates the TCP session with an RST packet. If the web browser implements the HTTP header range option, it can start a new session to fetch only the remaining part of the file, which prevents the firewall from triggering the same signature again due to the lack of context into the initial session and, at the same time, allows the web browser to reassemble the file and deliver the malicious content. Disabling this option prevents this from happening.



Disabling this option should not impact device performance. However, HTTP file transfer interruption recovery may be impaired.



- ❑ Create a Vulnerability Protection Profile that blocks protocol anomalies and all vulnerabilities with low and high severities.


A protocol anomaly occurs when a protocol behavior deviates from standard and compliant usage. For example, a malformed packet, poorly written application, or an application running on a non-standard port would all be considered protocol anomalies, and could be used as evasion tools.

If yours is a mission-critical network, where the business's highest priority is application availability, you should begin by alerting on protocol anomalies for a period of time to ensure that no critical internal applications are using established protocols in a non-standard way. If you find that certain critical applications trigger protocol anomaly signatures, you can then exclude those applications from protocol anomaly enforcement. To do this, add another rule to the Vulnerability Protection Profile that allows protocol anomalies and attach the profile to the security policy rule that enforces traffic to and from the critical applications.

Make sure that Vulnerability Protection Profile rules and security policy rules that allow protocol anomalies for critical internal applications are listed above rules that block protocol anomalies. Traffic is evaluated against security policy rules and associated Vulnerability Protection Profiles rules from top to bottom, and is enforced based on the first matching rule.


- Begin by alerting on protocol anomalies:

Create a Vulnerability Protection Profile rule with the **Action** set to Alert, the **Category** set to protocol-anomaly, and the **Severity** set to Any. Monitor your traffic to determine if any critical internal applications are using established protocols in non-standard ways. If you find this to be true, continue to allow protocol anomalies for those applications, and then block protocol anomalies for all other applications.




- Block protocol anomalies:

Create a Vulnerability Protection Profile rule with the **Category** set to protocol-anomaly, the rule **Action** set to Reset Both, and the **Severity** set to Any.



- Optionally allow protocol anomalies for critical applications that use established protocols in a non-standard way. To do this, create a Vulnerability Protection Profile rule that allows protocol anomalies: set the rule **Action** to Allow, the **Category** to protocol-anomaly, and the **Severity** to any. Attach the Vulnerability Protection Profile rule to the security policy rule that enforces traffic to and from critical applications.
- Add another rule to the Vulnerability Protection profile to block all vulnerabilities with low and higher severity. This rule must be listed after the rule that blocks protocol anomalies.



- ❑ Continue to attach the following security profiles to your Security policy rules to provide signature-based protection:
 - An Anti-Spyware profile to block all spyware with severity low and higher.
 - An Antivirus profile to block all content that matches an antivirus signature.

Set Up Antivirus, Anti-Spyware, and Vulnerability Protection

Every Palo Alto Networks next-generation firewall comes with predefined [Antivirus](#), [Anti-Spyware](#), and [Vulnerability Protection](#) profiles that you can attach to Security policy rules. 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.

Use the following workflow to set up the default Antivirus, Anti-Spyware, and Vulnerability Protection [Security Profiles](#).



*Palo Alto Networks defines a default action for all anti-spyware and vulnerability protection signatures. To see the default action, select **Objects > Security Profiles > Anti-Spyware** or **Objects > Security Profiles > Vulnerability Protection** and then select a profile. Click the **Exceptions** tab and then click **Show all signatures** to view the list of the signatures and the corresponding default Action. To change the default action, create a new profile and specify an Action, and/or add individual signature exceptions to **Exceptions** in the profile.*

STEP 1 | Verify that you have a Threat Prevention subscription.

The Threat Prevention subscription bundles the antivirus, anti-spyware, and vulnerability protection features in one license. To verify that you have an active Threat Prevention subscription, select **Device > Licenses** and verify that the **Threat Prevention** expiration date is in the future.

Threat Prevention	
Date Issued	February 03, 2016
Date Expires	February 03, 2021
Description	Antivirus, anti-spyware, vulnerability protection

STEP 2 | Download the latest content.

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** and install the latest Antivirus updates and then download and then **Install** the latest Applications and Threats updates.

STEP 3 | Schedule content updates.



Review the [Best Practices for Applications and Threats Content Updates](#) for important information on deploying updates.

1. Select **Device > Dynamic Updates** and then click **Schedule** to automatically retrieve signature updates for **Antivirus** and **Applications and Threats**.
2. Specify the frequency and timing for the updates:
 - **download-only**—The firewall automatically downloads the latest updates per the schedule you define but you must manually **Install** them.
 - **download-and-install**—The firewall automatically downloads and installs the updates per the schedule you define.
3. Click **OK** to save the update schedule; a commit is not required.
4. (Optional) Define a **Threshold** to indicate the minimum number of hours after an update becomes available before the firewall will download it. For example, setting the **Threshold** to **10** means the firewall will not download an update until it is at least 10 hours old regardless of the schedule.
5. (HA only) Decide whether to **Sync To Peer**, which enables peers to synchronize content updates after download and install (the update schedule does not sync across peers; you must manually configure the schedule on both peers).

There are additional considerations for deciding if and how to **Sync To Peer** depending on your HA deployment:

- **Active/Passive HA**—If the firewalls are using the MGT port for content updates, then schedule both firewalls to download and install updates independently. However, if the firewalls are using a data port for content updates, then the passive firewall will not download or install updates unless and until it becomes active. To keep the schedules in sync on both firewalls when using a data port for updates, schedule updates on both firewalls and then enable **Sync To Peer** so that whichever firewall is active downloads and installs the updates and also pushes the updates to the passive firewall.
- **Active/Active HA**—If the firewalls are using the MGT interface for content updates, then select **download-and-install** on both firewalls but do not enable **Sync To Peer**. However, if the firewalls are using a data port, then select **download-and-install** on both firewalls and enable **Sync To Peer** so that if one firewall goes into the active-secondary state, the active-primary firewall will download and install the updates and push them to the active-secondary firewall.

STEP 4 | (Optional) Create custom security profiles for antivirus, anti-spyware, and vulnerability protection.

Alternatively, you can use the predefined default or strict profiles.



Transition safely to best practice Security profiles for the best security posture.

- To create custom **Antivirus Profiles**, select **Objects > Security Profiles > Antivirus** and **Add** a new profile. Use the **Antivirus profile transition steps** to safely reach your goal.
- To create custom **Anti-Spyware Profiles**, select **Objects > Security Profiles > Anti-Spyware** and **Add** a new profile. Use the **Anti-Spyware profile transition steps** to safely reach your goal.
- To create custom **Vulnerability Protection Profiles**, select **Objects > Security Profiles > Vulnerability Protection** and **Add** a new profile. Use the **Vulnerability Protection profile transition steps** to safely reach your goal.

STEP 5 | Attach security profiles to your Security policy rules.



*When you configure the firewall with a Security policy rule that uses a Vulnerability Protection profile to block connections, the firewall automatically blocks that traffic in hardware (see **Monitor Blocked IP Addresses**).*

1. Select **Policies > Security** and select the rule you want to modify.
2. In the **Actions** tab, select **Profiles** as the **Profile Type**.
3. Select the security profiles you created for **Antivirus**, **Anti-Spyware**, and **Vulnerability Protection**.

The screenshot shows the 'Security Policy Rule' configuration window with the 'Actions' tab selected. The window is divided into several sections:

- Action Setting:** The 'Action' is set to 'Allow'. There is a checkbox for 'Send ICMP unreachable' which is currently unchecked.
- Profile Setting:** The 'Profile Type' is set to 'Profiles'. Below this, several security profiles are listed with dropdown menus: 'Antivirus' (default), 'Vulnerability Protection' (default), 'Anti-Spyware' (default), 'URL Filtering' (None), 'File Blocking' (None), 'Data Filtering' (None), and 'WildFire Analysis' (None).
- Log Setting:** There are checkboxes for 'Log at Session Start' and 'Log at Session End', both of which are checked. The 'Log Forwarding' is set to 'default'.
- Other Settings:** The 'Schedule' is set to 'None', and 'QoS Marking' is set to 'None'. There is an unchecked checkbox for 'Disable Server Response Inspection'.

At the bottom right of the window, there are 'Commit' and 'Cancel' buttons.

STEP 6 | Commit your changes.

Click **Commit**.

DNS Security

DNS Security is a continuously evolving threat prevention service designed to protect and defend your network from advanced threats using DNS. By leveraging advanced machine learning and predictive analytics, the service provides real-time DNS request analysis and rapidly produces and distributes DNS signatures that are specifically designed to defend against malware using DNS for C2 and data theft. Combined with an extensible cloud architecture, it provides access to a scalable threat intelligence system to keep your network protections up to date.


- [About DNS Security](#)
- [Domain Generation Algorithm \(DGA\) Detection](#)
- [DNS Tunneling Detection](#)
- [Cloud-Delivered DNS Signatures and Protections](#)
- [Enable DNS Security](#)

About DNS Security

With an active Threat Prevention license, customers can configure their firewalls to sinkhole DNS requests using a list of domains generated by Palo Alto Networks. These locally-accessed, customizable DNS signature lists are packaged with [antivirus and WildFire updates](#) and include the most relevant threats for policy enforcement and protection at the time of publication. For improved coverage against threats using DNS, the DNS Security subscription enables users to access real-time protections using advanced predictive analytics. Using techniques such as DGA/DNS tunneling detection and machine learning, threats hidden within DNS traffic can be proactively identified and shared through an infinitely scalable cloud service. Because the DNS signatures and protections are stored in a cloud-based architecture, you can access the full database of ever-expanding signatures that have been generated using a multitude of data sources. This allows you to defend against an array of threats using DNS in real-time against newly generated malicious domains. To combat future threats, updates to the analysis, detection, and prevention capabilities of the DNS Security service will be available through content releases.

To access the DNS Security service, you must have a valid Threat Prevention and DNS Security license.


The following workflow describes how the DNS Security service uses various data sources to generate DNS signatures:



Domain Generation Algorithm (DGA) Detection

Domain generation algorithms (DGAs) are used to auto-generate domains, typically in large numbers within the context of establishing a malicious command-and-control (C2) communications channel. DGA-based malware (such as Pushdo, BankPatch, and CryptoLocker) limit the number of domains from being blocked by hiding the location of their active C2 servers within a large number of possible suspects, and can be algorithmically generated based on factors such as time of day, cryptographic keys, or other unique values. While most domains generated by a DGA do not resolve as a valid domain, they must all be identified to fully defend against a given threat. DGA analysis determines whether a domain is likely to have been generated by a machine, rather than a person, by reverse-engineering and analyzing other frequently used techniques found in DGAs. Palo Alto Networks then uses these characteristics to identify and block previously unknown DGA-based threats in real-time.


You can analyze the sinkholed DNS queries by viewing the threat logs (**Monitor > Logs**, then select the log type from the list):



DNS Tunneling Detection


DNS tunneling can be used by attackers to encode data of non-DNS programs and protocols within DNS queries and responses. This provides attackers with an open back channel with which they can transfer files or remotely access the system. DNS tunnel detection uses machine learning to analyze the behavioral qualities of DNS queries, including n-gram frequency analysis of domains, entropy, query rate, and patterns to determine if the query is consistent with a DNS tunneling-based attack. Combined with the firewall's automated policy actions, this allows you to quickly detect C2 or data theft hidden in DNS tunnels and to automatically block it, based on your defined policy rules.

You can analyze the sinkholed DNS queries by viewing the threat logs (**Monitor > Logs**, then select the log type from the list):



Cloud-Delivered DNS Signatures and Protections

As a cloud-based service, DNS Security allows you to access an infinitely scalable DNS signature and protections source to defend your organization from malicious domains. Domain signatures and protections generated by Palo Alto Networks are derived from a multitude of sources, including WildFire traffic analysis, passive DNS, active web crawling & malicious web content analysis, URL sandbox analysis, HoneyNet, DGA reverse engineering, telemetry data, whois, the Unit 42 research organization, and third party data sources such as the [Cyber Threat Alliance](#). This on-demand cloud database provides users with access to the complete Palo Alto Network's DNS signature set, including signatures generated using advanced analysis techniques, as well as real-time DNS request analysis. Locally available, downloadable DNS signature sets (packaged with the [antivirus and WildFire updates](#)) come with a hard-coded capacity limitation of 100k signatures and do not include signatures generated through advanced analysis. To better accommodate the influx of new DNS signatures being produced on a daily basis, the cloud-based signature database provides users with instant access to newly added DNS signatures without the need to download updates. If network connectivity goes down or is otherwise unavailable, the firewall uses the onbox DNS signature set.



Enable DNS Security


To enable DNS sinkholing for domain queries using DNS security, you must activate your DNS Security subscription, create (or modify) an Anti-Spyware policy to reference the DNS Security service, enable the sinkhole action, and attach the profile to a security policy rule.

STEP 1 | [Activate Subscription Licenses.](#)

STEP 2 | Configure DNS signature policy settings to send malware DNS queries to the defined sinkhole.

1. Select **Objects > Security Profiles > Anti-Spyware**.
2. Create or modify an existing profile, or select one of the existing default profiles and clone it.
3. **Name** the profile and, optionally, provide a description.
4. Select the **DNS Signatures > Policies & Settings** tab.
5. If the **Palo Alto Networks DNS Security** source is not present, click **Add** and select it from the list.

6. Select an action to be taken when DNS lookups are made to known malware sites for the DNS Security signature source. The options are alert, allow, block, or sinkhole. Verify that the action is set to sinkhole.
 7. (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.
 8. In the **DNS Sinkhole Settings** section, verify that **Sinkhole** is enabled. For your convenience, the default Sinkhole address (sinkhole.paloaltonetworks.com) is set to access a Palo Alto Networks server. Palo Alto Networks can automatically refresh this 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. Select or create a **Security Policy Rule**.
3. On the **Actions** tab, select the **Log at Session End** check box to enable logging.
4. 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 or modified profile.
5. Click **OK** to save the policy rule.

STEP 4 | Test that the policy action is enforced.

1. Access the following test domains to verify that the policy action for a given threat type is being enforced:
 - Malware— test-malware.testpanw.com
 - C2— test-c2.testpanw.com
 - DGA— test-dga.testpanw.com
 - DNS Tunneling— test-dnstun.testpanw.com
2. To monitor the activity on the firewall:
 1. Select **ACC** and add a URL Domain as a global filter to view the Threat Activity and Blocked Activity for the domain you accessed.


Monitor > Logs > Threat

2. Select and filter by (`action eq sinkhole`) to view logs on sinkholed domains.

STEP 5 | Identify Infected Traffic Hosts in the Traffic Logs

STEP 6 | (Optional) Add domain signature exceptions in cases where false-positives occur.

1. Select **Objects > Security Profiles > Anti-Spyware**.
2. Select a profile to modify.
3. **Add** or modify the Anti-Spyware profile from which you want to exclude the threat signature, and select **DNS Signatures > Exceptions**.
4. Search for a DNS signature to exclude by entering the name or FQDN.
5. Select the **DNS Threat ID** for the DNS signature that you want to exclude from enforcement.
6. Click **OK** to save your new or modified Anti-Spyware profile.



STEP 7 | (Optional) Configure the DNS signature lookup timeout setting. If the firewall is unable to retrieve a signature verdict in the allotted time due to connectivity issues, the request, including all subsequent DNS responses, are passed through. You can check the average latency to verify that the requests fall within the configured period. If the average latency exceeds the configured period, consider updating the setting to a value that is higher than the average latency to prevent requests from timing out.

1. In the CLI, issue the following command to view the average latency.

```
show dns-proxy dns-signature
counters
```

The default timeout is 100 milliseconds.

2. Scroll down through the output to the latency section under the Signature query API heading and verify that the average latency falls within the defined timeout period. This latency indicates the amount of time it takes, on average, to retrieve a signature verdict from the DNS security service. Additional latency statistics for various latency periods can be found below the averages.

Signature query API:

•
•
•

```
[latency    ] :
    max    1870 (ms)  min     16(ms)  avg     27(ms)
```

```
50 or less : 47246
100 or less : 113
200 or less : 25
400 or less : 15
else : 21
```

3. If the average latency is consistency above the default timeout value, you can raise the setting so that the requests fall within a given period. Select **Device > Content-ID** and update the **Realtime Signature Lookup** setting.
4. Commit the changes.

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. A compromised host might initiate communication with a command-and-control (C2) server—once the connection is made, an attacker can remotely control the infected host, in order to further infiltrate the network or exfiltrate data.

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 firewall has two sources of DNS signatures that it can use to identify malicious and C2 domains:


- (Requires Threat Prevention) Local DNS signatures—This is a limited, on-box set of DNS signatures that the firewall can use to identify malicious domains. The firewall gets new DNS signatures as part of daily antivirus updates.
- (Requires DNS Security) [DNS Security](#) signatures—The firewall accesses the Palo Alto Networks DNS Security cloud service to check for malicious domains against the complete database of DNS signatures. Certain signatures—that only DNS Security provides—can uniquely detect C2 attacks that use machine learning techniques, like domain generation algorithms (DGAs) and DNS tunneling.

DNS queries to domains in the local DNS signature set or the DNS Security signature set are redirected to a Palo Alto Networks server, and the host is unable to access the malicious domain. The following topics provide details on how to enable DNS sinkholing so that you can identify infected hosts.

- Learn [How DNS Sinkholing Works](#).
- [Configure DNS Sinkholing](#).
- [Configure DNS Sinkholing for a List of Custom Domains](#).
- [Enable DNS Security](#) to sinkhole C2 domains.
- [Configure the Sinkhole IP Address to a Local Server on Your Network](#).
- [See Infected Hosts that Attempted to Connect to a Malicious Domain](#).

How DNS Sinkholing Works

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 IP address that you define if you choose to [Configure DNS Sinkholing for a List of Custom Domains](#)). Infected hosts can then be easily identified in the traffic logs.



Configure DNS Sinkholing

To enable DNS sinkholing, 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 signature source that you specify are resolved to the default Palo Alto Networks sinkhole IP address. The IP addresses currently are IPv4—sinkhole.paloaltonetworks.com and a loopback address IPv6 address—::1. These address are subject to change and can be updated with content updates.

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. Verify that **Palo Alto Network Content DNS Signatures** is present in the **DNS Signature Source**.
5. (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.

STEP 2 | Verify the sinkholing settings on the Anti-Spyware profile.

1. On the **DNS Signatures** tab, verify that the **Action on DNS Queries** is **sinkhole**.
2. 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](#).

3. Click **OK** to save the Anti-Spyware profile.

STEP 3 | Attach the Anti-Spyware profile to a Security policy rule.

1. Select **Policies > Security** and select a security policy rule.

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 by monitoring the activity on the firewall.

1. Select **ACC** and add a URL Domain as a global filter to view the Threat Activity and Blocked Activity for the domain you accessed.
2. Select **Monitor > Logs > Threat** and filter by `(action eq sinkhole)` to view logs on sinkholed domains.

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 model supports a maximum of 50,000 domain names total in one or more external dynamic lists but no maximum limit is enforced for any one list.

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. Click **Add** and select **External Dynamic Lists** in the drop-down.




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 the external dynamic list from the Anti-Spyware profile (see [Configure the Firewall to Access an External Dynamic List](#)). The **Type** is preset to **Domain List**.
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.

STEP 2 | Verify the sinkholing settings on the Anti-Spyware profile.

1. On the **DNS Signatures** tab, verify that the **Action on DNS Queries** is **sinkhole**.
2. 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](#).



3. Click **OK** to save the Anti-Spyware profile.

STEP 3 | Attach the Anti-Spyware profile to a Security policy rule.

1. Select **Policies > Security** and select a security policy rule.
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. [View External Dynamic List Entries](#) that belong to the domain list, and access a domain from the list.
2. To monitor the activity on the firewall:
 1. Select **ACC** and add a URL Domain as a global filter to view the Threat Activity and Blocked Activity for the domain you accessed.
 2. Select **Monitor > Logs > Threat** and filter by `(action eq sinkhole)` to view logs on sinkholed domains.

STEP 5 | Verify whether entries in the external dynamic list are ignored or skipped.

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>
```



As an alternative, you can use the firewall interface to [Retrieve an External Dynamic List from the Web Server](#).

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

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 [Verify the sinkholing settings on the Anti-Spyware profile](#). 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 policy 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 policy rule and then **Commit**.

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.

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.

1. 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.
2. In the second column of the release note, locate a line item with a domain extension (for example, .com, .edu, or .net). The left column will 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:tbsbana 1
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:5:5:5:510.15.0.20
```

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.

See Infected Hosts that Attempted to Connect to a Malicious Domain

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.

- 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.



- 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.

Source	Source Host Name	Destination	Destination Host Name
192.168.2.50	192.168.2.10	10.15.0.20	10.15.0.20

5. To view scheduled reports that have run, select **Monitor > Reports**.

Set Up Data Filtering

Use [Data Filtering Profiles](#) to prevent sensitive, confidential, and proprietary information from leaving your network. Predefined patterns, built-in settings, and customizable options make it easy for you to protect files that contain certain file properties (such as a document title or author), credit card numbers, regulated information from different countries (like social security numbers), and third-party data loss prevention (DLP) labels.

- **Predefined Data Patterns**—Easily filter common patterns, including credit card numbers. Predefined data filtering patterns also identify specific (regulated) information from different countries of the world, such as social security numbers (United States), INSEE Identification numbers (France), and New Zealand Internal Revenue Department Identification Numbers. Many of the predefined data filtering patterns enable compliance for standards such as HIPAA, GDPR, Gramm-Leach-Bliley Act.
- **Built-In Support for Azure Information Protection and Titus Data Classification**—Predefined file properties allow you to filter content based on [Azure Information Protection](#) and Titus labels. Azure Information Protection labels are stored in metadata, so make sure that you [know the GUID of the Azure Information Protect label](#) that you want the firewall to filter.
- **Custom Data Patterns for Data Loss Prevention (DLP) Solutions**—If you're using a third-party, endpoint DLP solution that populates file properties to indicate sensitive content, you can create a custom data pattern to identify the file properties and values tagged by your DLP solution and then log or block the files that your Data Filtering profile detects based on that pattern.

Create a Data Filtering Profile

[Data Filtering](#) profiles can keep sensitive information from leaving your network.

To get started, you'll first create a data pattern that specifies the information types and fields that you want the firewall to filter. Then, you attach that pattern to a data filtering profile, which specifies how you want to enforce the content that the firewall filters. Add the data filtering profile to a security policy rule to start filtering traffic matching the rule.

STEP 1 | Define a new data pattern object to detect the information you want to filter.


1. Select **Objects > Custom Objects > Data Patterns** and **Add** a new object.
2. Provide a descriptive **Name** for the new object.
3. (Optional) Select **Shared** if you want the data pattern to be available to:
 - **Every virtual system (vsys) on a multi-vsys firewall**—If cleared (disabled), the data pattern is available only to the Virtual System selected in the **Objects** tab.
 - **Every device group on Panorama**—If cleared (disabled), the data pattern is available only to the Device Group selected in the **Objects** tab.
4. (Optional—Panorama only) Select **Disable override** to prevent administrators from overriding the settings of this data pattern object in device groups that inherit the object. This selection is cleared by default, which means administrators can override the settings for any device group that inherits the object.
5. (Optional—Panorama only) Select **Data Capture** to automatically collect the data that is blocked by the filter.



Specify a password for Manage Data Protection on the Settings page to view your captured data (Device > Setup > Content-ID > Manage Data Protection).

6. Set the **Pattern Type** to one of the following:

- **Predefined Pattern**—Filter for credit card, social security numbers, and personally identifiable information for several compliance standards including HIPAA, GDPR, Gramm-Leach-Bliley Act.
 - **Regular Expression**—Filter for custom data patterns.
 - **File Properties**—Filter based on file properties and the associated values.
7. **Add** a new rule to the data pattern object.
8. Specify the data pattern according to the **Pattern Type** you selected for this object:
- **Predefined**—Select the **Name** and choose the predefined data pattern on which to filter.
 - **Regular Expression**—Specify a descriptive **Name**, select the **File Type** (or types) you want to scan, and then enter the specific **Data Pattern** you want the firewall to detect.
 - **File Properties**—Specify a descriptive **Name**, select the **File Type** and **File Property** you want to scan, and enter the specific **Property Value** that you want the firewall to detect.
 - **To filter Titus classified documents:** Select one of the non-AIP protected file types, and set the **File Property** to TITUS GUID. Enter the Titus label GUID as the **Property Value**.
 - **For Azure Information Protection labeled documents:** Select any **File Type** except Rich Text Format. For the file type you choose, set the **File Property** to Microsoft MIP Label, and enter the [Azure Information Protection label GUID](#) as the **Property Value**.



9. Click **OK** to save the data pattern.

STEP 2 | Add the data pattern object to a data filtering profile.

1. Select **Objects > Security Profiles > Data Filtering** and **Add** or modify a data filtering profile.
2. Provide a descriptive **Name** for the new profile.
3. **Add** a new profile rule and select the Data Pattern you created in Step .
4. Specify **Applications**, **File Types**, and what **Direction** of traffic (upload or download) you want to filter based on the data pattern.



The file type you select must be the same file type you defined for the data pattern earlier, or it must be a file type that includes the data pattern file type. For example, you could define both the data pattern object and the data filtering profile to scan all Microsoft Office documents. Or, you could define the data pattern object to match to only Microsoft PowerPoint Presentations while the data filtering profile scans all Microsoft Office documents.

If a data pattern object is attached to a data filtering profile and the configured file types do not align between the two, the profile will not correctly filter documents matched to the data pattern object.

5. Set the **Alert Threshold** to specify the number of times the data pattern must be detected in a file to trigger an alert.
6. Set the **Block Threshold** to block files that contain at least this many instances of the data pattern.
7. Set the **Log Severity** recorded for files that match this rule.
8. Click **OK** to save the data filtering profile.

STEP 3 | Apply the data filtering settings to traffic.

1. Select **Policies > Security** and **Add** or modify a security policy rule.
2. Select **Actions** and set the Profile Type to **Profiles**.
3. Attach the Data Filtering profile you created in Step 2 to the security policy rule.
4. Click **OK**.

STEP 4 | (Recommended) Prevent web browsers from resuming sessions that the firewall has terminated.



This option ensures that when the firewall detects and then drops a sensitive file, a web browser cannot resume the session in an attempt to retrieve the file.

1. Select **Device > Setup > Content-ID** and edit Content-ID Settings.
2. Clear the **Allow HTTP partial response**.
3. Click **OK**.


STEP 5 | Monitor files that the firewall is filtering.

Select **Monitor > Data Filtering** to view the files that the firewall has detected and blocked based on your data filtering settings.

Predefined Data Filtering Patterns

To comply with standards such as HIPAA, GDPR, and the Gramm-Leach-Bliley Act, the firewall provides predefined data patterns. You can use these patterns to prevent common types of sensitive information, like credit cards and social security numbers, from leaving your network.

You can find predefined data patterns by selecting **Objects > Custom Objects > Data Patterns** and clicking **Add** a new object. Then, set the **Pattern Type** to **Predefined Pattern** and **Add** a new rule to the data pattern object. Select a data pattern from the list that appears under **Name**.





If the type of information you want to protect is not covered in the list of predefined patterns, you can use [regular expressions](#) to create custom patterns.

The following is a list of available data patterns:

Pattern	Description
Credit Card Numbers	16-digit credit card numbers
Social Security Numbers	9-digit social security numbers with dashes
Social Security Numbers (without dash separator)	9-digit social security numbers without dashes
ABA Routing Number	The American Banking Association Routing Number
AHV Identification Number	Swiss Alters und Hinterlassenenversicherungsnummer
Codice Fiscale Identification Number	Italian Fiscal Tax Code Card Identification Number
CorporateNumber Identification Number	Japanese National Tax Agency Corporate Number
CUSIP Identification Number	Committee on Uniform Security Identification Procedures Identification Number
DEA Registration Number	U.S. Drug Enforcement Administration Registration Number
DNI Identification Number	Spanish Documento nacional de identidad Identification Number number
HK Identification Number	Hong Kong Residents Identification Number
INSEE Identification Number	French National Institute of Statistics and Economic Studies identification number
IRD Identification Number	New Zealand Internal Revenue Department Identification Number
MyKad Identification Number	Malaysia MyKad Identity Card Identification Number
MyNumber Identification Number	Japanese Social Security and Tax Number System Identification Number
NHI Identification Number	New Zealand National Health Index Number
NIF Identification Number	Spanish Tax Identification Number
NIN Identification Number	Taiwan Identification Card Number

Pattern	Description
NRIC Identification Number	Singapore National Registration Identity Card Identification Number
Permanent Account Identification Number	India Permanent Account Number of Indian nationals
PRC Identification Number	People's Republic of China Resident Identification Number
PRN Identification Number	Republic of South Korea Resident Registration Number
Republic of South Korea Resident Registration	Republic of South Korea Resident Registration Number

Set Up File Blocking

[File Blocking Profiles](#) allow you to identify specific file types that you want to block or monitor. For most traffic (including traffic on your internal network), 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 notice that the browser is attempting to download something they were not aware of. For policy rules that allow general web browsing, be stricter with your file blocking because the risk of users unknowingly downloading malicious files is much higher. For this type of traffic, attach a more strict file blocking profile that also blocks portable executable (PE) files.

You can define your own custom File Blocking profiles or choose one of the following predefined profiles when applying file blocking to a Security policy rule. You can clone and edit the predefined profiles, which are available with content release version 653 and later, and then follow [File Blocking profile safe transition steps](#) to preserve application availability as you transition to [best practice file blocking](#) settings:

- **basic file blocking**—Attach this profile to the Security policy rules that allow traffic to and from less sensitive applications to block files that are commonly included in malware attack campaigns or that have no real use case for upload/download. This profile blocks upload and download of PE files (.scr, .cpl, .dll, .ocx, .pif, .exe), Java files (.class, .jar), Help files (.chm, .hlp) and other potentially malicious file types, including .vbe, .hta, .wsf, .torrent, .7z, .rar, .bat. Additionally, it prompts users to acknowledge when they attempt to download encrypted-rar or encrypted-zip files. This rule alerts on all other file types to give you complete visibility into all file types coming in and out of your network.
- **strict file blocking**—Use this stricter profile on the Security policy rules that allow access to your most sensitive applications. This profile blocks the same file types as the other profile, and additionally blocks flash, .tar, multi-level encoding, .cab, .msi, encrypted-rar, and encrypted-zip files.

These predefined profiles are designed to provide the most secure posture for your network. However, if you have business-critical applications that rely on some of the applications that are blocked in these default profiles, you can clone the profiles and modify them as necessary. Make sure you only use the modified profiles for those users who need to upload and/or download a risky file type. Additionally, to reduce your attack surface, make sure you are using other security measures to ensure that the files your users are uploading and downloading do not pose a threat to your organization. For example, if you must allow download of PE files, make sure you are [sending all unknown PE files to WildFire for analysis](#). Additionally, maintain a strict URL filtering policy to ensure that users cannot download content from web sites that have been known to host malicious content.

STEP 1 | Create the file blocking profile.

1. Select **Objects > Security Profiles > File Blocking** and **Add** a profile.
2. Enter a **Name** for the file blocking profile such as **Block_EXE**.
3. (Optional) Enter a **Description**, such as **Block users from downloading exe files from websites**.
4. (Optional) Specify that the profile is **Shared** with:
 - **Every virtual system (vsys) on a multi-vsys firewall**—If cleared (disabled), the profile is available only to the Virtual System selected in the **Objects** tab.
 - **Every device group on Panorama**—If cleared (disabled), the profile is available only to the Device Group selected in the **Objects** tab.
5. (Optional—Panorama only) Select **Disable override** to prevent administrators from overriding the settings of this file blocking profile in device groups that inherit the profile. This selection is cleared

by default, which means administrators can override the settings for any device group that inherits the profile.

STEP 2 | Configure the file blocking options.

1. **Add** and define a rule for the profile.
2. Enter a **Name** for the rule, such as **BlockEXE**.
3. Select **Any** or specify one or more specific **Applications** for filtering, such as **web-browsing**.



Only web browsers can display the response page (continue prompt) that allows users to confirm their Choosing any other application results in blocked traffic for those applications because there is no prompt displayed to allow users to continue.

4. Select **Any** or specify one or more specific **File Types**, such as **exe**.
5. Specify the **Direction**, such as **download**.
6. Specify the **Action** (**alert**, **block**, or **continue**). For example, select **continue** to prompt users for confirmation before they are allowed to download an executable (.exe) file. Alternatively, you could **block** the specified files or you could configure the firewall to simply trigger an **alert** when a user downloads an executable file.
7. Click **OK** to save the profile.

STEP 3 | Apply the file blocking profile to a security policy rule.

1. Select **Policies > Security** and either select an existing policy rule or **Add** a new rule as described in [Set Up a Basic Security Policy](#).
2. On the **Actions** tab, select the file blocking profile you configured in the previous step. In this example, the profile name is **Block_EXE**.
3. **Commit** your configuration.

STEP 4 | To test your file blocking configuration, access an endpoint PC in the trust zone of the firewall and attempt to download an executable file from a website in the untrust zone; a response page should display. Click **Continue** to confirm that you can download the file. You can also set other actions, such as **alert** or **block**, which do not provide an option for the user to continue the download. 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 continue action, you can choose only the web-browsing application. If you choose any other application, traffic that matches the security policy will not flow through the firewall because users are not prompted with an option to continue. Additionally, you need to configure and enable a decryption policy for HTTPS websites.



Check your logs to determine the application used when you test 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`. (It can help 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, and 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 Protection profile to a Security policy 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](#).

Customize the Action and Trigger Conditions for a Brute Force Signature

The firewall includes two types of predefined brute force signatures—parent signatures and child signatures. 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 specified time interval and that matches the traffic pattern defined in the child signature.

Typically, the default action for a child signature is *allow* because a single event is not indicative of an attack. This ensures that legitimate traffic is not blocked and avoids generating threat logs for non-noteworthy events. Palo Alto Networks recommends that you do not change the default action without careful consideration.

In most cases, the brute force signature is a noteworthy event due to its recurrent pattern. If needed, you can do one of the following to customize the action for a brute-force signature:

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

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



To effectively mitigate an attack, specify the block-ip address action instead of the drop or reset action for most brute force signatures.

STEP 1 | Create a new Vulnerability Protection profile.

1. Select **Objects > Security Profiles > Vulnerability Protection** and **Add** a profile.
2. Enter a **Name** for the Vulnerability Protection profile.
3. (Optional) Enter a **Description**.
4. (Optional) Specify that the profile is **Shared** with:
 - **Every virtual system (vsys) on a multi-vsyz firewall**—If cleared (disabled), the profile is available only to the Virtual System selected in the **Objects** tab.
 - **Every device group on Panorama**—If cleared (disabled), the profile is available only to the Device Group selected in the **Objects** tab.
5. (Optional—Panorama only) Select **Disable override** to prevent administrators from overriding the settings of this Vulnerability Protection profile in device groups that inherit the profile. This selection is cleared by default, which means administrators can override the settings for any device group that inherits the profile.


STEP 2 | Create a rule that defines the action for all signatures in a category.

1. On the **Rules** tab, **Add** and enter a **Rule Name** for a new rule.
2. (Optional) Specify a specific threat name (default is **any**).
3. Set the **Action**. In this example, it is set to **Block IP**.



If you set a Vulnerability Protection profile to Block IP, the firewall first uses hardware to block IP addresses. If attack traffic exceeds the blocking capacity of the hardware, the firewall then uses software blocking mechanisms to block the remaining IP addresses.

4. Set **Category** to **brute-force**.
5. (Optional) If blocking, specify the **Host Type** on which to block: **server** or **client** (default is **any**).
6. See Step 3 to customize the action for a specific signature.
7. See Step 4 to customize the trigger threshold for a parent signature.




8. Click **OK** to save the rule and the profile.

STEP 3 | (Optional) Customize the action for a specific signature.

1. On the **Exceptions** tab, **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: **Allow**, **Alert**, **Block Ip**, or **Drop**. If you select **Block Ip**, complete these additional tasks:
 1. Specify the **Time** period (in seconds) after which to trigger the action.
 2. Specify whether to **Track By** and block the IP address using the **IP source** or the **IP source and destination**.
4. Click **OK**.
5. For each modified signature, select the check box in the **Enable** column.
6. Click **OK**.

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. Edit () the time attribute and the aggregation criteria for the signature.
 2. To modify the trigger threshold, specify the **Number of Hits** per number of **seconds**.
 3. Specify whether to aggregate the number of hits (**Aggregation Criteria**) by **source**, **destination**, or **source-and-destination**.
 4. Click **OK**.

STEP 5 | Attach this new profile to a Security policy rule.

1. Select **Policies > Security** and **Add** or modify a Security policy rule.
2. On the **Actions** tab, select **Profiles** as the **Profile Type** for the Profile Setting.
3. Select your **Vulnerability Protection** profile.
4. Click **OK**.

STEP 6 | Commit your changes.

1. Click **Commit**.

Enable Evasion Signatures

Palo Alto Networks evasion signatures detect crafted HTTP or TLS requests, and can alert to instances where a client connects to a domain other than the domain specified in a DNS query. Evasion signatures are effective only when the firewall is also enabled to act as a DNS proxy and resolve domain name queries. As a best practice, take the following steps to enable evasion signatures.

STEP 1 | Enable a firewall intermediate to clients and servers to act as a DNS proxy.

Configure a [DNS Proxy Object](#), including:

- Specify the interfaces on which you want the firewall to listen for DNS queries.
- Define the DNS servers with which the firewall communicates to resolve DNS requests.
- Set up static FQDN-to-IP address entries that the firewall can resolve locally, without reaching out to DNS servers.
- Enable caching for resolved hostname-to-IP-address mappings.

STEP 2 | Get the latest Applications and Threats content version (at least content version 579 or later).

1. Select **Device > Dynamic Updates**.
2. **Check Now** to get the latest Applications and Threats content update.
3. Download and Install Applications and Threats content version 579 (or later).

STEP 3 | Define how the firewall should enforce traffic matched to evasion signatures.

1. Select **Objects > Security Profiles > Anti-Spyware** and **Add** or modify an [Anti-spyware profile](#).
2. Select **Exceptions** and select **Show all signatures**.
3. Filter signatures based on the keyword **evasion**.
4. 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** for signature IDs 14978 and 14984 to **alert** or **drop**.
5. Click **OK** to save the updated Anti-spyware profile.
6. 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 4 | Commit your changes.

Click **Commit**.

Prevent Credential Phishing

Phishing sites are sites that attackers disguise as legitimate websites with the aim to steal user information, especially the credentials that provide access to your network. When a phishing email enters a network, it takes just a single user to click the link and enter credentials to set a breach into motion. You can detect and prevent in-progress phishing attacks, thereby preventing credential theft, by controlling sites to which users can submit corporate credentials based on the site's URL category. This allows you to block users from submitting credentials to untrusted sites while allowing users to continue to submit credentials to corporate and sanctioned sites.

Credential phishing prevention works by scanning username and password submissions to websites and comparing those submissions against valid corporate credentials. You can choose what websites you want to either allow or block corporate credential submissions to based on the URL category of the website. When the firewall detects a user attempting to submit credentials to a site in a category you have restricted, it either displays a block response page that prevents the user from submitting credentials, or presents a continue page that warns users against submitting credentials to sites classified in certain URL categories, but still allows them to continue with the credential submission. You can customize these block pages to educate users against reusing corporate credentials, even on legitimate, non-phishing sites.

To enable Credential phishing prevention you must configure both [User-ID](#) to detect when users submit valid corporate credentials to a site (as opposed to personal credentials) and [URL Filtering](#) to specify the URL categories in which you want to prevent users from entering their corporate credentials. The following topics describe the different methods you can use to detect credential submissions and provide instructions for configuring credential phishing protection.

- [Methods to Check for Corporate Credential Submissions](#)
- [Configure Credential Detection with the Windows-based User-ID Agent](#)
- [Set Up Credential Phishing Prevention](#)

Methods to Check for Corporate Credential Submissions

Before you [Set Up Credential Phishing Prevention](#), decide which method you want the firewall to use to check if credentials submitted to a web page are valid, corporate credentials.

Method to Check Submitted Credentials	User-ID Configuration Requirements	How does this method detect corporate usernames and/or passwords as users submit them to websites?
Group Mapping	Group Mapping configuration on the firewall	The firewall determines if the username a user submits to a restricted site matches any valid corporate username. To do this, the firewall matches the submitted username to the list of usernames in its user-to-group mapping table to detect when users submit a corporate usernames to a site in a restricted category. This method only checks for corporate username submissions based on LDAP group membership, which makes it simple to configure, but more prone to false positives.
IP User Mapping	IP-address-to-username mappings	The firewall determines if the username a user submits to a restricted site maps to the IP address of the logged-in user.

Method to Check Submitted Credentials	User-ID Configuration Requirements	How does this method detect corporate usernames and/or passwords as users submit them to websites?
	identified through User Mapping , GlobalProtect , or Authentication Policy and Captive Portal .	To do this, the firewall matches the IP address of the logged in user and the username submitted to a web site to its IP-address-to-user mapping table to detect when users submit their corporate usernames to a site in a restricted category. Because this method matches the IP address of the logged-in user associated with the session against the IP-address-to-username mapping table, it is an effective method for detecting corporate username submissions, but it does not detect corporate password submission. If you want to detect corporate username and password submission, you must use the Domain Credential Filter method.
Domain Credential Filter	Windows-based User-ID agent configured with the User-ID credential service add-on - AND - IP-address-to- username mappings identified through User Mapping, GlobalProtect, or Authentication Policy and Captive Portal.	The firewall determines if the username and password a user submits matches the same user's corporate username and password. To do this, the firewall must able to match credential submissions to valid corporate usernames and passwords and verify that the username submitted maps to the IP address of the logged in user as follows: • To detect corporate usernames and passwords—The firewall retrieves a secure bit mask, called a <i>bloom filter</i>, from a Windows-based User-ID agent equipped with the User-ID credential service add-on. This add-on service scans your directory for usernames and password hashes and deconstructs them into a secure bit mask—the bloom filter—and delivers it to the Windows agent. The firewall retrieves the bloom filter from the Windows agent at regular intervals and, whenever it detects a user submitting credentials to a restricted category, it reconstructs the bloom filter and looks for a matching username and password hash. The firewall can only connect to one Windows-based User-ID agent running the User-ID credential service add-on. • To verify that the credentials belong to the logged-in user—The firewall looks for a mapping between the IP address of the logged-in user and the detected username in its IP-address-to-username mapping table. To learn more how the domain credential method works, and the requirements for enabling this type of detection, see Configure Credential Detection with the Windows-based User-ID Agent.

Configure Credential Detection with the Windows-based User-ID Agent

[Domain Credential Filter](#) detection enables the firewall to detect passwords submitted to web pages. This credential detection method requires the Windows-based User-ID agent and the User-ID credential service, an add-on to the User-ID agent, to be installed on a *read-only domain controller (RODC)*.



The Domain Credential Filter detection method is supported with the Windows-based User-ID agent only. You cannot use the PAN-OS integrated User-ID agent to configure this method of credential detection.

An RODC is a Microsoft Windows server that maintains a read-only copy of an Active Directory database that a domain controller hosts. When the domain controller is located at a corporate headquarters, for example, RODCs can be deployed in remote network locations to provide local authentication services. Installing the User-ID agent on an RODC can be useful for a few reasons: access to the domain controller directory is not required to enable credential detection and you can support credential detection for a limited or targeted set of users. Because the directory the RODC hosts is read-only, the directory contents remain secure on the domain controller.



Because you must install the Windows-based User-ID agent on the RODC for credential detection, as a best practice deploy a separate agent for this purpose. Do not use the User-ID agent installed on the RODC to map IP addresses to users.

After you install the User-ID agent on an RODC, the User-ID credential service runs in the background and scans the directory for the usernames and password hashes of group members that are listed in the RODC password replication policy (PRP)—you can define who you want to be on this list. The User-ID credential service then takes the collected usernames and password hashes and deconstructs the data into a type of bit mask called a *bloom filter*. Bloom filters are compact data structures that provide a secure method to check if an element (a username or a password hash) is a member of a set of elements (the sets of credentials you have approved for replication to the RODC). The User-ID credential service forwards the bloom filter to the Windows-based User-ID agent; the firewall retrieves the latest bloom filter from the User-ID agent at regular intervals and uses it to detect usernames and password hash submissions. Depending on your settings, the firewall then blocks, alerts, or allows on valid password submissions to web pages, or displays a response page to users warning them of the dangers of phishing, but allowing them to continue with the submission.

Throughout this process, the User-ID agent does not store or expose any password hashes, nor does it forward password hashes to the firewall. Once the password hashes are deconstructed into a bloom filter, there is no way to recover them.

STEP 1 | Configure User Mapping Using the Windows User-ID Agent.



To enable credential detection, you must install the Windows-based User-ID agent on an RODC (requires Microsoft Windows 2012 or 2012 R2). Install a separate User-ID agent for this purpose.


Important items to remember when setting up User-ID to enable [Domain Credential Filter](#) detection:

- Because the effectiveness of credential phishing detection is dependent on your RODC setup, make sure that you also review best practices and recommendations for [RODC Administration](#).
- Download User-ID [software updates](#):
 - User-ID Agent Windows installer—UalInstall-x.x.x-x.msi.


- User-ID Agent Credential Service Windows installer—UaCredInstall64-x.x.x-x.msi.
- Install the User-ID agent and the User Agent Credential service on an RODC using an account that has privileges to read Active Directory via LDAP (the User-ID agent also requires this privilege).
- The User-ID Agent Credential Service requires permission to log on with the local system account. For more information, refer to [Create a Dedicated Service Account for the User-ID Agent](#).
- The service account must be a member of the local administrator group on the RODC. For more information, refer to the following [link](#).

STEP 2 | Enable the User-ID agent and the User Agent Credential service (which runs in the background to scan permitted credentials) to share information.

1. On the RODC server, launch the User-ID Agent.
2. Select **Setup** and edit the Setup section.



3. Select the **Credentials** tab. This tab only displays if you have already installed the User-ID Agent Credential Service.



4. Select **Import from User-ID Credential Agent**. This enables the User-ID agent to import the bloom filter that the User-ID credential agent creates to represent users and the corresponding password hashes.
5. Click **OK**, **Save** your settings, and **Commit**.

STEP 3 | In the RODC directory, define the group of users for which you want to support credential submission detection.

- Confirm that the groups that should receive credential submission enforcement are added to the Allowed RODC Password Replication Group.
- Check that none of the groups in the Allowed RODC Password Replication Group are also in the Denied RODC Password Replication Group by default. Groups listed in both will not be subject to credential phishing enforcement.

STEP 4 | Continue to the next task.

[Set Up Credential Phishing Prevention](#) on the firewall.

Set Up Credential Phishing Prevention

After you have decided which of the [Methods to Check for Corporate Credential Submissions](#) you want to use, take the following steps to enable the firewall to detect when users submit corporate credentials to web pages and either alert on this action, block the credential submission, or require users to acknowledge the dangers of phishing before continuing with credential submission.

STEP 1 | If you have not done so already, [Enable User-ID](#).

Each of the [Methods to Check for Corporate Credential Submissions](#) requires a different User-ID configuration to check for corporate credential submissions:

- If you plan to use the group mapping method, which detects whether a user is submitting a valid corporate username, [Map Users to Groups](#).
- If you plan to use the IP user mapping method, which detects whether a user is submitting a valid corporate username and that username belongs to the logged-in user, [Map IP Addresses to Users](#).
- If you plan to use the domain credential filter method, which detects whether a user is submitting a valid username and password and that those credentials belong to the logged in user, [Configure Credential Detection with the Windows-based User-ID Agent](#) and [Map IP Addresses to Users](#).

STEP 2 | If you have not done so already, configure a [best practice URL Filtering profile](#) to ensure protection against URLs that have been observed hosting malware or exploitive content.

1. Select **Objects > Security Profiles > URL Filtering** and **Add** or modify a URL Filtering profile.
2. Block access to all known dangerous URL categories: malware, phishing, dynamic-dns, unknown, command-and-control, extremism, copyright-infringement, proxy-avoidance-and-anonymizers, newly-registered-domain, grayware, and parked.

STEP 3 | Configure the URL Filtering profile to detect corporate credential submissions to websites that are in allowed URL categories.



The firewall does not check credential submissions for trusted sites, even if you enable the checks for the URL categories for these sites, to provide best performance. The trusted sites represent sites where Palo Alto Networks has not observed any malicious or phishing attacks. Updates for this trusted site list are delivered through Application and Threat content updates. For a list of App-IDs that are exempt from credential detection, see [Trusted App-IDs That Skip Credential Submission Detection on live.paloaltonetworks.com](#).

1. Select **User Credential Detection**.
2. Select one of the [Methods to Check for Corporate Credential Submissions](#) to web pages from the **User Credential Detection** drop-down:



Confirm that the format for the primary username is the same as the username format that the User-ID source provides.

- **Use IP User Mapping**—Checks for valid corporate username submissions and verifies that the login username maps to the source IP address of the session. To do this, the firewall matches the submitted username and source IP address of the session against its IP-address-to-username mapping table. To use this method you can use any of the user mapping methods described in [Map IP Addresses to Users](#).

- **Use Domain Credential Filter**—Checks for valid corporate usernames and password submissions and verifies that the username maps to the IP address of the logged in user. See [Configure Credential Detection with the Windows-based User-ID Agent](#) for instructions on how to set up User-ID to enable this method.
- **Use Group Mapping**—Checks for valid username submissions based on the user-to-group mapping table populated when you configure the firewall to [Map Users to Groups](#).

With group mapping, you can apply credential detection to **any** part of the directory, or specific group, such as groups like IT that have access to your most sensitive applications.



This method is prone to false positives in environments that do not have uniquely structured usernames. Because of this, you should only use this method to protect your high-value user accounts.

3. Set the **Valid Username Detected Log Severity** the firewall uses to log detection of corporate credential submissions. By default, the firewall logs these events as medium severity.

STEP 4 | Block (or alert) on credential submissions to allowed sites.

1. Select **Categories**.
2. For each Category to which **Site Access** is allowed, select how you want to treat **User Credential Submissions**:
 - **alert**—Allow users to submit credentials to the website, but generate a URL Filtering log each time a user submits credentials to sites in this URL category.
 - **allow**—(default) Allow users to submit credentials to the website.
 - **block**—Block users from submitting credentials to the website. When a user tries to submit credentials, the firewall displays the [Anti Phishing Block Page](#), preventing the credential submission.
 - **continue**—Present the [Anti Phishing Continue Page](#) response page to users when they attempt to submit credentials. Users must select Continue on the response page to continue with the submission.
3. Select **OK** to save the URL Filtering profile.

STEP 5 | Apply the URL Filtering profile with the credential detection settings to your Security policy rules.

1. Select **Policies > Security** and **Add** or modify a Security policy rule.
2. On the **Actions** tab, set the **Profile Type** to **Profiles**.
3. Select the new or updated **URL Filtering** profile to attach it to the Security policy rule.
4. Select **OK** to save the Security policy rule.

STEP 6 | Commit the configuration.

STEP 7 | Monitor credential submissions the firewall detects.



Select ACC > Hosts Visiting Malicious URLs to see the number of users who have visited malware and phishing sites.


Select **Monitor > Logs > URL Filtering**.

The new **Credential Detected** column indicates events where the firewall detected a HTTP post request that included a valid credential:

Category	Application	Action	Credential Detected
unknown	web-browsing	Block-url	yes
EDL: shared-URL	web-browsing	Block-url	yes
EDL: shared-URL	web-browsing	Block-url	yes
malware	web-browsing	Block-url	yes
EDL: shared-URL	web-browsing	Block-url	yes
malware	web-browsing	Block-url	yes

(To display this column, hover over any column header and click the arrow to select the columns you'd like to display).

Log entry details also indicate credential submissions:



STEP 8 | Validate and troubleshoot credential submission detection.

- Use the following CLI command to view credential detection statistics:

```
> show user credential-filter statistics
```

The output for this command varies depending on the method configured for the firewall to detect credential submissions. For example, if the [Domain Credential Filter](#) method is configured in any URL Filtering profile, a list of User-ID agents that have forwarded a bloom filter to the firewall is displayed, along with the number of credentials contained in the bloom filter.

- ([Group Mapping method only](#)) Use the following CLI command to view group mapping information, including the number of URL Filtering profiles with Group Mapping credential detection enabled and the usernames of group members that have attempted to submit credentials to a restricted site.

```
> show user group-mapping statistics
```

- ([Domain Credential Filter method only](#)) Use the following CLI command to see all Windows-based User-ID agents that are sending mappings to the firewall:

```
> show user user-id-agent state all
```

The command output now displays bloom filter counts that include the number of bloom filter updates the firewall has received from each agent, if any bloom filter updates failed to process, and how many seconds have passed since the last bloom filter update.

- ([Domain Credential Filter method only](#)) The Windows-based User-ID agent displays log messages that reference BF (bloom filter) pushes to the firewall. In the User-ID agent interface, select **Monitoring > Logs**.

Monitor Blocked IP Addresses

The firewall maintains a block list of source IP addresses that it's blocking. When the firewall blocks a source IP address, such as when you configure either of the following policy rules, the firewall blocks that traffic in hardware before those packets use CPU or packet buffer resources:

- A classified DoS Protection policy rule with the action to **Protect** (a classified DoS Protection policy specifies that incoming connections match a source IP address, destination IP address, or source and destination IP address pair, and is associated with a Classified DoS Protection profile, as described in [DoS Protection Against Flooding of New Sessions](#)).
- A [Security Policy](#) rule that uses a Vulnerability Protection profile

Hardware IP address blocking is supported on PA-3200 Series, PA-5200 Series, and PA-7000 Series firewalls.

You can view the block list, get detailed information about an IP address on the block list, or view counts of addresses that hardware and software are blocking. You can delete an IP address from the list if you think it shouldn't be blocked. You can change the source of detailed information about addresses on the list. You can also change how long hardware blocks IP addresses.

- View block list entries.

1. Select **Monitor > Block IP List**.

Entries on the block list indicate in the Type column whether they were blocked by hardware (hw) or software (sw).


2. View at the bottom of the screen:

- Count of **Total Blocked IPs** out of the number of blocked IP addresses the firewall supports.
- Percentage of the block list the firewall has used.

3. To filter the entries displayed, select a value in a column (which creates a filter in the **Filters** field) and Apply Filter (). Otherwise, the firewall displays the first 1,000 entries.

4. Enter a **Page** number or click the arrows at the bottom of the screen to advance through pages of entries.

5. To view details about an address on the block list, hover over a Source IP address and click the down arrow link. Click the **Who Is** link, which displays [Network Solutions Who Is](#) information about the address.



Block Time	Type	Source IP Address	Ingress Zone	Time Remaining	Block Source
09/08 11:57:52	hw	192.168.2.28	L2_trust	0	test_block
09/08 11:57:54	sw	192.168.2.28	L2_trust	0	test_block

- Delete block list entries.



Delete an entry if you determine the IP address shouldn't be blocked. Then revise the policy rule that caused the firewall to block the address.

1. Select **Monitor > Block IP List**.
2. Select one or more entries and click **Delete**.

3. (Optional) Select **Clear All** to remove all entries from the list.

- Disable or re-enable hardware IP address blocking for troubleshooting purposes.



While hardware IP address blocking is disabled, the firewall still performs any software IP address blocking you have configured.

```
> set system setting hardware-acl-blocking [enable | disable]
```



To conserve CPU and packet buffer resources, leave hardware IP address blocking enabled unless Palo Alto Networks technical support asks you to disable it, for example, if they are debugging a traffic flow.

- Tune the number of seconds that IP addresses blocked by hardware remain on the block list (range is 1-3,600; default is 1).

```
> set system setting hardware-acl-blocking duration <seconds>
```



Maintain a shorter duration for hardware block list entries than software block list entries to reduce the likelihood of exceeding the blocking capacity of the hardware.

- Change the default website for finding more information about an IP address from [Network Solutions Who Is](#) to a different website.

```
# set deviceconfig system ip-address-lookup-url <url>
```

- View counts of source IP addresses blocked by hardware and software, for example to see the rate of an attack.

View the total sum of IP address entries on the hardware block table and block list (blocked by hardware and software):

```
> show counter global name flow_dos_blk_num_entries
```

View the count of IP address entries on the hardware block table that were blocked by hardware:

```
> show counter global name flow_dos_blk_hw_entries
```

View the count of IP address entries on the block list that were blocked by software:

```
> show counter global name flow_dos_blk_sw_entries
```

- View block list information per slot on a PA-7000 Series firewall.

```
> show dos-block-table software filter slot <slot-number>
```

Threat Signature Categories

There are three types of Palo Alto Networks threat signatures, each designed to detect different types of threats as the firewall scans network traffic:

- Antivirus signatures—Detect viruses and malware found in executables and file types.
- Anti-spyware signatures—Detects command-and-control (C2) activity, where spyware on an infected client is collecting data without the user's consent and/or communicating with a remote attacker.
- Vulnerability signatures—Detects system flaws that an attacker might otherwise attempt to exploit.

A signature's severity indicates the risk of the detected event, and a signature's default action (for example, block or alert) is how Palo Alto Networks recommends that you enforce matching traffic.

You must [Set Up Antivirus, Anti-Spyware, and Vulnerability Protection](#) to tell the firewall what action to take when it detects a threat, and you can easily use the default security profiles to start blocking threats based on Palo Alto Networks recommendations. For each signature type, category, and even specific signatures you can continue to modify or create new profiles to more granularly enforce potential threats.

The following table lists all possible signature categories by type—Antivirus, Spyware, and Vulnerability—and includes the content update (Applications and Threats, Antivirus, or WildFire) that provides the signatures in each category. You can also go to the Palo Alto Networks [Threat Vault](#) to [Learn More About Threat Signatures](#).

Threat Category	Content Update that Provides These Signatures	Description
Antivirus Signatures		
apk	Antivirus WildFire or WildFire Private	Malicious Android Application (APK) files.
dmg	Antivirus WildFire or WildFire Private	Malicious Apple disk image (DMG) files, that are used with Mac OS X.
flash	Antivirus Wildfire or WildFire Private	Adobe Flash applets and Flash content embedded in web pages.
java-class	Antivirus	Java applets (JAR/class file types).
macho	Antivirus Wildfire or WildFire Private	Mach object files (Mach-O) are executables, libraries, and object code that are native to Mac OS X.
office	Antivirus Wildfire or WildFire Private	Microsoft Office files, including documents (DOC, DOCX, RTF), workbooks (XLS, XLSX), and PowerPoint presentations (PPT, PPTX).

Threat Category	Content Update that Provides These Signatures	Description
openoffice	Antivirus Wildfire or WildFire Private	Office Open XML (OOXML) 2007+ documents.
pdf	Antivirus Wildfire or WildFire Private	Portable Document Format (PDF) files.
pe	Antivirus Wildfire or WildFire Private	Portable executable (PE) files can automatically execute on a Microsoft Windows system and should be only allowed when authorized. These files types include: • Object code. • Fonts (FONs). • System files (SYS). • Driver files (DRV). • Windows control panel items (CPLs). • DLLs (dynamic-link libraries). • OCXs (libraries for OLE custom controls, or ActiveX controls). • SCRs (scripts that can be used to execute other files). • Extensible Firmware Interface (EFI) files, which run between an OS and firmware in order to facilitate device updates and boot operations. • Program information files (PIFs).
pkg	Antivirus Wildfire or WildFire Private	Apple software installer packages (PKGs), used with Mac OS X.
Spyware Signatures		
adware	Applications and Threats	Detects programs that display potentially unwanted advertisements. Some adware modifies browsers to highlight and hyperlink the most frequently searched keywords on web pages-these links redirect users to advertising websites. Adware can also retrieve updates from a command-and-control (C2) server and install those updates in a browser or onto a client system. Newly-released protections in this category are rare.
autogen	Antivirus	These payload-based signatures detect command-and-control (C2) traffic and are automatically-generated. Importantly, autogen signatures can detect C2 traffic even when the C2 host is unknown or changes rapidly.

Threat Category	Content Update that Provides These Signatures	Description
backdoor	Applications and Threats	Detects a program that allows an attacker to gain unauthorized remote access to a system.
botnet	Applications and Threats	Indicates botnet activity. A botnet is a network of malware-infected computers ("bots") that an attacker controls. The attacker can centrally command every computer in a botnet to simultaneously carry out a coordinated action (like launching a DoS attack, for example).
browser-hijack	Applications and Threats	Detects a plugin or software that is modifying browser settings. A browser hijacker might take over auto search or track users' web activity and send this information to a C2 server. Newly-released protections in this category are rare.
cryptominer	Applications and Threats	(Sometimes known as cryptojacking or miners) Detects the download attempt or network traffic generated from malicious programs designed to use computing resources to mine cryptocurrencies without the user's knowledge. Cryptominer binaries are frequently delivered by a shell script downloader that attempts to determine system architecture and kill other miner processes on the system. Some miners execute within other processes, such as a web browser rendering a malicious web page.
data-theft	Applications and Threats	Detects a system sending information to a known C2 server. Newly-released protections in this category are rare.
dns	Antivirus	Detects DNS requests to connect to malicious domains. dns and dns-wildfire signatures detect the same malicious domains; however, dns signatures are included in the daily Antivirus content update and dns-wildfire signatures are included in the WildFire updates that release protections every 5 minutes.
dns-security	Antivirus	Detects DNS requests to connect to malicious domains. dns-security includes signatures from dns and dns-wildfire in addition to the unique signatures generated by the DNS Security service.
dns-wildfire	Wildfire or WildFire Private	Detects DNS requests to connect to malicious domains. dns and dns-wildfire signatures detect the same malicious domains; however, dns signatures are included in the daily Antivirus content update and dns-wildfire signatures are included in the WildFire updates that release protections every 5 minutes.

Threat Category	Content Update that Provides These Signatures	Description
downloader	Applications and Threats	(Also known as droppers, stagers, or loaders) Detects programs that use an internet connection to connect to a remote server to download and execute malware on the compromised system. The most common use case is for a downloader to be deployed as the culmination of <i>stage one</i> of a cyber attack, where the downloader's fetched payload execution is considered <i>second stage</i> . Shell scripts (Bash, PowerShell, etc.), trojans, and malicious lure documents (also known as maldocs) such as PDFs and Word files are common downloader types.
fraud	Applications and Threats	(Including form-jacking, phishing, and scams) Detects access to compromised websites that have been determined to be injected with malicious JavaScript code to collect sensitive user information. (for example, Name, address, email, credit card number, CVV, expiration date) from payment forms that are captured on the checkout pages of e-commerce websites.
hacktool	Applications and Threats	Detects traffic generated by software tools that are used by malicious actors to conduct reconnaissance, attack or gain access to vulnerable systems, exfiltrate data, or create a command and control channel to surreptitiously control a computer system without authorization. These programs are strongly associated with malware and cyber attacks. Hacking tools might be deployed in a benign manner when used in Red and Blue Team operations, penetration tests, and R&D. The use or possession of these tools may be illegal in some countries, regardless of intent.
keylogger	Applications and Threats	Detects programs that allow attackers to secretly track user activity, by logging keystrokes and capturing screenshots. Keyloggers use various C2 methods to periodically sends logs and reports to a predefined e-mail address or a C2 server. Through keylogger surveillance, an attacker could retrieve credentials that would enable network access.
networm	Applications and Threats	Detects a program that self-replicates and spreads from system to system. Net-worms might use shared resources or leverage security failures to access target systems.
phishing-kit	Applications and Threats	Detects when a user attempts to connect to a phishing kit landing page (likely after receiving an email with a link to the malicious site). A phishing website tricks users into submitting credentials that an attacker can steal to gain access to the network.  <i>In addition to blocking access to phishing kit landing pages, enable Multi-Factor Authentication and Credential Phishing</i>

Threat Category	Content Update that Provides These Signatures	Description
		Prevention to prevent phishing attacks at all stages.
post-exploitation	Applications and Threats	Detects activity that indicates the post-exploitation phase of an attack, where an attacker attempts to assess the value of a compromised system. This might include evaluating the sensitivity of the data stored on the system, and the system's usefulness in further compromising the network.
webshell	Applications and Threats	Detects web shells and web shell traffic, including implant detection and command and control interaction. Web shells must first be implanted by a malicious actor onto the compromised host, most often targeting a web server or framework. Subsequent communication with the web shell file frequently enables a malicious actor to establish a foothold in the system, conduct service and network enumeration, data exfiltration, and remote code execution in the context of the web server user. The most common web shell types are PHP, .NET, and Perl markup scripts. Attackers can also use web shell-infected web servers (the web servers can be both internet-facing or internal systems) to target other internal systems.
spyware	Applications and Threats	Detect outbound C2 communication. These signatures are either auto-generated or are manually created by Palo Alto Networks researchers.  <i>Spyware and autogen signatures both detect outbound C2 communication; however, autogen signatures are payload-based and can uniquely detect C2 communications with C2 hosts that are unknown or change rapidly.</i>

Vulnerability Signatures

brute force	Applications and Threats	A brute-force signature detects multiple occurrences of a condition in a particular time frame. While the activity in isolation might be benign, the brute-force signature indicates that the frequency and rate at which the activity occurred is suspect. For example, a single FTP login failure does not indicate malicious activity. However, many failed FTP logins in a short period likely indicate an attacker attempting password combinations to access an FTP server. You can tune the action and trigger conditions for brute force signatures.

Threat Category	Content Update that Provides These Signatures	Description
code execution	Applications and Threats	Detects a code execution vulnerability that an attacker can leverage to run code on a system with the privileges of the logged-in user.
code-obfuscation	Applications and Threats	Detects code that has been transformed to conceal certain data while retaining its function. Obfuscated code is difficult or impossible to read, so it's not apparent what commands the code is executing or with which programs its designed to interact. Most commonly, malicious actors obfuscate code to conceal malware. More rarely, legitimate developers might obfuscate code to protect privacy, intellectual property, or to improve user experience. For example, certain types of obfuscation (like minification) reduce file size, which decreases website load times and bandwidth usage.
dos	Applications and Threats	Detects a denial-of-service (DoS) attack, where an attacker attempts to render a targeted system unavailable, temporarily disrupting the system and dependent applications and services. To perform a DoS attack, an attacker might flood a targeted system with traffic or send information that causes it to fail. DoS attacks deprive legitimate users (like employees, members, and account holders) of the service or resource to which they expect access.
exploit-kit	Applications and Threats	Detects an exploit kit landing page. Exploit kit landing pages often contain several exploits that target one or many common vulnerabilities and exposures (CVEs), for multiple browsers and plugins. Because the targeted CVEs change quickly, exploit-kit signatures trigger based on the exploit kit landing page, and not the CVEs. When a user visits a website with an exploit kit, the exploit kit scans for the targeted CVEs and attempts to silently deliver a malicious payload to the victim's computer.
info-leak	Applications and Threats	Detects a software vulnerability that an attacker could exploit to steal sensitive or proprietary information. Often, an info-leak might exist because comprehensive checks do not exist to guard the data, and attackers can exploit info-leaks by sending crafted requests.
insecure-credentials	Applications and Threats	Detects the use of weak, compromised, and manufacturer default passwords for software, network appliances, and IoT devices.
overflow	Applications and Threats	Detects an overflow vulnerability, where a lack of proper checks on requests could be exploited by an attacker. A successful attack could lead to remote code execution with the privileges of the application, server or operating system.

Threat Category	Content Update that Provides These Signatures	Description
phishing	Applications and Threats	Detects when a user attempts to connect to a phishing kit landing page (likely after receiving an email with a link to the malicious site). A phishing website tricks users into submitting credentials that an attacker can steal to gain access to the network.  <i>In addition to blocking access to phishing kit landing pages, enable Multi-Factor Authentication and Credential Phishing Prevention to prevent phishing attacks at all stages.</i>
protocol-anomaly	Applications and Threats	Detects protocol anomalies, where a protocol behavior deviates from standard and compliant usage. For example, a malformed packet, poorly-written application, or an application running on a non-standard port would all be considered protocol anomalies, and could be used as evasion tools. It is a best practice to block protocol anomalies of any severity.
sql-injection	Applications and Threats	Detects a common hacking technique where an attacker inserts SQL queries into an application's requests, in order to read from or modify a database. This type of technique is often used on websites that do not comprehensively sanitize user input.

Create Threat Exceptions

Palo Alto Networks defines a recommended default action (such as block or alert) for threat signatures. You can use a threat ID to exclude a threat signature from enforcement or modify the action the firewall enforces for that threat signature. For example, you can modify the action for threat signatures that are triggering false positives on your network.

Configure threat exceptions for antivirus, vulnerability, spyware, and DNS signatures to change firewall enforcement for a threat. However, before you begin, make sure the firewall is detecting and enforcing threats based on the default signature settings:

- [Get the latest](#) Antivirus, Threats and Applications, and WildFire signature updates.
- [Set Up Antivirus, Anti-Spyware, and Vulnerability Protection](#) and apply these security profiles to your security policy.

STEP 1 | Exclude antivirus signatures from enforcement.



While you can use an Antivirus profile to exclude antivirus signatures from enforcement, you cannot change the action the firewall enforces for a specific antivirus signature. However, you can define the action for the firewall to enforce for viruses found in different types of traffic by editing the Decoders (Objects > Security Profiles > Antivirus > <antivirus-profile> > Antivirus).


1. Select **Objects > Security Profiles > Antivirus**.
2. **Add** or modify an existing Antivirus profile from which you want to exclude a threat signature and select **Virus Exception**.
3. **Add** the **Threat ID** for the threat signature you want to exclude from enforcement.



4. Click **OK** to save the Antivirus profile.

STEP 2 | Modify enforcement for vulnerability and spyware signatures (except DNS signatures; skip to the next option to modify enforcement for DNS signatures, which are a type of spyware signature).

1. Select **Objects > Security Profiles > Anti-Spyware** or **Objects > Security Profiles > Vulnerability Protection**.
2. **Add** or modify an existing Anti-Spyware or Vulnerability Protection profile from which you want to exclude the threat signature and then select **Exceptions**.
3. **Show all signatures** and then filter to select the signature for which you want to modify enforcement rules.
4. Check the box under the **Enable** column for the signature whose enforcement you want to modify.
5. Select the **Action** you want the firewall to enforce for this threat signature.




For signatures that you want to exclude from enforcement because they trigger false positives, set the **Action** to **Allow**.

6. Click **OK** to save your new or modified Anti-Spyware or Vulnerability Protection profile.

STEP 3 | Modify enforcement for DNS signatures.

By default, the DNS lookups to malicious hostnames that DNS signatures are detect are sinkholed.

1. Select **Objects > Security Profiles > Anti-Spyware**.
2. **Add** or modify the Anti-Spyware profile from which you want to exclude the threat signature, and select **DNS Signatures**.
3. **Add** the **DNS Threat ID** for the DNS signature that you want to exclude from enforcement:



4. Click **OK** to save your new or modified Anti-Spyware profile.

Custom Signatures

You can create custom threat signatures to detect and block very specific traffic. The following resources will help you get started:

- [Creating Custom Threat Signatures from Snort Signatures](#)—This tech note provides guidance on how to create custom signatures, by demonstrating how to create one based on a Snort signature (Snort is a free and open source IPS).
- [Creating Custom Threat and Application Signatures](#)—Walks through custom signature examples, and describes how to use regex and contexts to create custom signatures (with details on each available context).

Monitor and Get Threat Reports

Features of [Threat Vault](#) and [AutoFocus](#) are integrated into the firewall to provide visibility into the nature of the threats the firewall detects and to give a more complete picture of how an artifact fits into your organization's network traffic (an artifact is property, activity, or behavior associated with a file, email link, or session). You can get immediate, contextual information about a threat or to seamlessly shift your threat investigation from the firewall to the Threat Vault and AutoFocus.




Additionally, you can use [Threat Signature Categories](#)—which classify types of threat events—to narrow your view into a certain type of threat activity or to build custom reports.

- [Monitor Activity and Create Custom Reports Based on Threat Categories](#)
- [Learn More About Threat Signatures](#)
- [AutoFocus Threat Intelligence for Network Traffic](#)


Monitor Activity and Create Custom Reports Based on Threat Categories

Threat categories classify different types of threat signatures to help you understand and draw connections between events threat signatures detect. Threat categories are subsets of the more broad threat signature types: spyware, vulnerability, antivirus, and DNS signatures. Threat log entries display the **Threat Category** for each recorded event.


- Filter Threat logs by threat category.
 1. Select **Monitor > Logs > Threat**.
 2. Add the Threat Category column so you can view the Threat Category for each log entry:




3. To filter based on Threat Category:
 - Use the log query builder to add a filter with the **Attribute** Threat Category and in the **Value** field, enter a Threat Category.
 - Select the Threat Category of any log entry to add that category to the filter:



- Filter ACC activity by threat category.
 1. Select **ACC** and add Threat Category as a global filter:



2. Select the Threat Category to filter all ACC tabs.



- Create custom reports based on threat categories to receive information about specific types of threats that the firewall has detected.
 1. Select **Monitor > Manage Custom** reports to [add a new custom report or modify an existing one](#).
 2. Choose the **Database** to use as the source for the custom report—in this case, select **Threat** from either of the two types of database sources, [summary databases](#) and [Detailed logs](#). Summary database data is condensed to allow a faster response time when generating reports. Detailed logs take longer to generate but provide an itemized and complete set of data for each log entry.
 3. In the Query Builder, add a report filter with the Attribute **Threat Category** and in the Value field, select a threat category on which to base your report.
 4. To test the new report settings, click **Run Now**.
 5. Click **OK** to save the report.

Learn More About Threat Signatures

Firewall Threat logs record all threats the firewall detects based on threat signatures ([Set Up Antivirus](#), [Anti-Spyware](#), and [Vulnerability Protection](#)) and the ACC displays an overview of the top threats on your network. Each event the firewall records includes an ID that identifies the associated threat signature.

You can use the threat ID found with a Threat log or ACC entry to:

-
- Easily check if a threat signature is configured as an exception to your security policy ([Create Threat Exceptions](#)).
 - Find the latest Threat Vault information about a specific threat. Because the Threat Vault is integrated with the firewall, you can view threat details directly in the firewall context or launch a Threat Vault search in a new browser window for a threat the firewall logged.



If a signature has been disabled, the signature UTID might be reused for a new signature.

Review the content update release notes for notifications regarding new and disabled signatures. Signatures might be disabled in cases where: the activity the signature detects has fallen out of use by attackers, the signature generated significant false positives, or the signature was consolidated with other like signatures into a single signature (signature optimization).

STEP 1 | Confirm the firewall is connected to the Threat Vault.


Select **Device > Setup > Management** and edit the **Logging and Reporting** setting to **Enable Threat Vault Access**. Threat vault access is enabled by default.

STEP 2 | Find the threat ID for threats the firewall detects.

- To see each threat event the firewall detects based on threat signatures, select **Monitor > Logs > Threat**. You can find the ID for a threat entry listed in the ID column, or select the log entry to view log details, including the Threat ID.
- To see an overview of top threats on the network, select **ACC > Threat Activity** and take a look at the Threat Activity widget. The ID column displays the threat ID for each threat displayed.
- To see details for threats that you can configure as threat exceptions (meaning, the firewall enforces the threat differently than the default action defined for the threat signature), select **Objects > Security Profiles > Anti-Spyware/Vulnerability Protection**. Add or modify a profile and click the **Exceptions** tab to view configured exceptions. If no exceptions are configured, you can filter for threat signatures or select **Show all signatures**.

STEP 3 | Hover over a **Threat Name** or the threat **ID** to open the drop-down, and click **Exception** to review both the threat details and how the firewall is configured to enforce the threat.

For example, find out more about a top threat charted on the ACC:



STEP 4 | Review the latest **Threat Details** for the threat and launch a Threat Vault search based on the threat ID.

- Threat details displayed include the latest Threat Vault information for the threat, resources you can use to learn more about the threat, and CVEs associated with the threat.
- Select **View in Threat Vault** to open a Threat Vault search in a new window and look up the latest information the Palo Alto Networks threat database has for this threat signature.

STEP 5 | Check if a threat signature is configured as an exception to your security policy.

- If the **Used in current security rule** column is clear, the firewall is enforcing the threat based on the recommended default signature action (for example, block or alert).
- A checkmark anywhere in the **Used in current security rule** column indicates that a security policy rule is configured to enforce a non-default action for the threat (for example, allow), based on the associated **Exempt Profiles** settings.



The Used in security rule column does not indicate if the Security policy rule is enabled, only if the Security policy rule is configured with the threat exception. Select Policies > Security to check if an indicated security policy rule is enabled.

STEP 6 | Add an IP address on which to filter the threat exception or view existing **Exempt IP Addresses**.

Configure an exempt IP address to enforce a threat exception only when the associated session has either a matching source or destination IP address; for all other sessions, the threat is enforced based on the default signature action.

AutoFocus Threat Intelligence for Network Traffic

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:

- View an AutoFocus intelligence summary for session artifacts recorded in the firewall logs.
- 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.

- [AutoFocus Intelligence Summary](#)
- [Enable AutoFocus Threat Intelligence](#)
- [View and Act on AutoFocus Intelligence Summary Data](#)




You can also enforce policy based on AutoFocus findings:

- [Export AutoFocus artifacts \(IP addresses, URLs, and domains\) and use them in an external dynamic list.](#)
- [Use an AutoFocus miner as an external dynamic list source.](#)

AutoFocus Intelligence Summary

The AutoFocus Intelligence Summary offers a centralized view of information about an artifact that AutoFocus has extracted from threat intelligence gathered from other AutoFocus users, WildFire, the PAN-DB URL filtering database, Unit 42, and open-source intelligence.



AutoFocus Intelligence Summary

Analysis Information	The Analysis Information tab displays the following information: • Sessions—The number of sessions logged in your firewall(s) in which the firewall detected samples associated with the artifact. • Samples—A comparison of organization and global samples associated with the artifact and grouped by WildFire verdict (benign, malware, or grayware). <i>Global</i> refers to samples from all WildFire submissions, while <i>organization</i> refers only to samples submitted to WildFire by your organization. • Matching Tags—The AutoFocus tags matched to the artifact. AutoFocus Tags indicate whether an artifact is linked to malware or targeted attacks.
Passive DNS	The Passive DNS tab displays passive DNS history that includes the artifact. This 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 to which the DNS request resolved (private IP addresses are not displayed) • The number of times the request was made • The date and time the request was first seen and last seen
Matching Hashes	The Matching Hashes tab displays the 5 most recently detected matching samples. Sample information includes: • The SHA256 hash of the sample • The sample file type • The date and time that WildFire analyzed a sample and assigned a WildFire verdict to it • The WildFire verdict for the sample • The date and time that WildFire updated the WildFire verdict for the sample (if applicable)

Enable AutoFocus Threat Intelligence

Activate the AutoFocus license, and enable the firewall to communicate with AutoFocus. Once you're set up, you'll be able to view the [AutoFocus Intelligence Summary](#) for a log or ACC artifact, to assess its pervasiveness in your network and any associated threats.

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 show the license, [Activate Subscription Licenses](#).

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 [Assess Firewall Artifacts with AutoFocus](#).

View and Act on AutoFocus Intelligence Summary Data

Interact with the AutoFocus Intelligence Summary to display more information about an artifact or extend your artifact research to AutoFocus. AutoFocus tags reveal if the artifact is associated with certain types of malware or malicious behavior.

STEP 1 | Confirm that the firewall is connected to AutoFocus.

[Enable AutoFocus Threat Intelligence](#) on the firewall (active AutoFocus subscription required).

STEP 2 | Find artifacts to investigate.

You can view an AutoFocus Intelligence Summary for artifacts when you:

- [View Logs](#) (Traffic, Threat, URL Filtering, WildFire Submissions, Data Filtering, and Unified logs only).
- [View External Dynamic List Entries](#).

STEP 3 | Hover over an artifact to open the drop-down, and click **AutoFocus**.



The AutoFocus Intelligence Summary is only available for the following types of artifacts:

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

STEP 4 | Launch an AutoFocus search for the artifact for which you opened the AutoFocus Intelligence Summary.

Click the **Search AutoFocus for...** link at the top of the AutoFocus Intelligence Summary window. The search results include all samples associated with the artifact. Toggle between the **My Samples** and **All Samples** tabs and compare the number of samples to determine the pervasiveness of the artifact in your organization.



STEP 5 | Launch an AutoFocus search for other artifacts in the AutoFocus Intelligence Summary.

Click on the following artifacts to determine their pervasiveness in your organization:

- WildFire verdicts in the Analysis Information tab
- URLs and IP addresses in the Passive DNS tab
- The SHA256 hashes in the Matching Hashes tab

STEP 6 | View the number of sessions associated with the artifact in your organization per month.

Hover over the session bars.




STEP 7 | View the number of samples associated with the artifact by scope and WildFire verdict.

Hover over the samples bars.



STEP 8 | View more details about matching AutoFocus. tags.

Hover over a matching tag to view the tag description and other tag details.



STEP 9 | View other samples associated with a matching tag.

Click a matching tag to launch an AutoFocus search for that tag. The search results include all samples matched to the tag.

Unit 42 tags identify threats and campaigns that pose a direct security risk. Click on a Unit 42 matching tag to see how many samples in your network are associated with the threat the tag identifies.

STEP 10 | Find more matching tags for an artifact.

Click the ellipsis (...) to launch an AutoFocus search for the artifact. The Tags column in the search results displays more matching tags for the artifact, which give you an idea of other malware, malicious behavior, threat actors, exploits, or campaigns where the artifact is commonly detected.



Share Threat Intelligence with Palo Alto Networks

Telemetry is the process of collecting and transmitting data for analysis. When you enable telemetry on the firewall, the firewall periodically collects and sends information that includes applications, threats, and device health to Palo Alto Networks. Sharing threat intelligence provides the following benefits:

- Enhanced vulnerability and spyware signatures delivered to you and other customers worldwide. For example, when a threat event triggers vulnerability or spyware signatures, the firewall shares the URLs associated with the threat with the Palo Alto Networks threat research team, so they can properly classify the URLs as malicious.
- Rapid testing and evaluation of experimental threat signatures with no impact to your network, so that critical threat prevention signatures can be released to all Palo Alto Networks customers faster.
- Improved accuracy and malware detection abilities within PAN-DB URL filtering, DNS-based command-and-control (C2) signatures, and WildFire.

Palo Alto Networks uses the threat intelligence extracted from telemetry to deliver these benefits to you and other Palo Alto Networks users. All Palo Alto Networks users benefit from the telemetry data shared by each user, making telemetry a community-driven approach to threat prevention. Palo Alto Networks does not share your telemetry data with other customers or third-party organizations.

- [What Telemetry Data Does the Firewall Collect?](#)
- [Passive DNS Monitoring](#)
- [Enable Telemetry](#)

What Telemetry Data Does the Firewall Collect?

The firewall collects and forwards different sets of telemetry data to Palo Alto Networks based on the Telemetry settings you enable. The firewall collects the data from fields in your log entries (see [Log Types and Severity Levels](#)); the log type and combination of fields vary based on the setting. Review the following table before you [Enable Telemetry](#).

Setting	Description
Application Reports	The number and size of known applications by destination port, unknown applications by destination port, and unknown applications by destination IP address. The firewall generates these reports from Traffic logs and forwards them every 4 hours.
Threat Prevention Reports	Attacker information, the number of threats for each source country and destination port, and the correlation objects that threat events triggered. The firewall generates these reports from Threat logs and forwards them every 4 hours.
URL Reports	URLs with the following PAN-DB URL categories: malware, phishing, dynamic DNS, proxy-avoidance, questionable, parked, and unknown (URLs that PAN-DB has not yet categorized). The firewall generates these reports from URL Filtering logs. URL Reports also include PAN-DB statistics such as the version of the URL filtering database on the firewall and on the PAN-DB cloud, the number of URLs in those databases, and the number of URLs that the firewall

Setting	Description
	categorized. These statistics are based on the time that the firewall forwarded the URL Reports. The firewall forwards URL Reports every 4 hours.
File Type Identification Reports	Information about files that the firewall has blocked or allowed based on data filtering and file blocking settings. The firewall generates these reports from Data Filtering logs and forwards them every 4 hours.
Threat Prevention Data	Log data from threat events that triggered signatures that Palo Alto Networks is evaluating for efficacy. Threat Prevention Data provides Palo Alto Networks more visibility into your network traffic than other telemetry settings. When enabled, the firewall may collect information such as source or victim IP addresses. Enabling Threat Prevention Data also allows unreleased signatures that Palo Alto Networks is currently testing to run in the background. These signatures do not affect your security policy rules and firewall logs, and have no impact to your firewall performance. The firewall forwards Threat Prevention Data every 5 minutes.
Threat Prevention Packet Captures	Packet captures (if you have enabled your firewall to Take a Threat Packet Capture) of threat events that triggered signatures that Palo Alto Networks is evaluating for efficacy. Threat Prevention Packet Captures provide Palo Alto Networks more visibility into your network traffic than other telemetry settings. When enabled, the firewall may collect information such as source or victim IP addresses. The firewall forwards Threat Prevention Packet Captures every 5 minutes.
Product Usage Statistics	Back traces of firewall processes that have failed, as well as information about the firewall status. Back traces outline the execution history of the failed processes. These reports include details about the firewall model and the PAN-OS and content release versions installed on your firewall. The firewall forwards Product Usage Statistics every 5 minutes.
Passive DNS Monitoring	Domain-to-IP address mappings based on firewall traffic. When you enable Passive DNS Monitoring, the firewall acts as a passive DNS sensor and send DNS information to Palo Alto Networks for analysis. The firewall forwards data from Passive DNS Monitoring in 1 MB batches.

Passive DNS Monitoring

Passive DNS monitoring enables the firewall to act as a passive DNS sensor and send DNS information to Palo Alto Networks for analysis to improve threat intelligence and threat prevention capabilities. The data collected includes non-recursive DNS query (that is, the web browser sends a query to a DNS server to translate a domain to an IP address, and the server returns a response without querying other DNS servers) and response packet payloads. See [DNS Overview](#) for more background information about DNS.

The threat intelligence that the firewall collects from passive DNS monitoring consists solely of domain-to-IP address mappings. 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

passive DNS 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 PAN-DB URL category and DNS-based C2 signature accuracy and WildFire malware detection.

The firewall forwards DNS responses only 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

Enable Telemetry

When you enable telemetry, you define what data the firewall collects and shares with Palo Alto Networks. For some telemetry settings, you can preview what the data that your firewall sends will look like before committing. The firewall uses the **Palo Alto Networks Services** service route to send the data you share from telemetry to Palo Alto Networks.

STEP 1 | Select **Device > Setup > Telemetry**, and edit the Telemetry settings.

STEP 2 | Select the telemetry data you want to share with Palo Alto Networks. For more specific descriptions of this data, see [What Telemetry Data Does the Firewall Collect?](#) By default, all telemetry settings are disabled.




To enable Threat Prevention Packet Captures, you must also enable Threat Prevention Data.

STEP 3 | Open a report sample () to view the type of data that the firewall collects for Application Reports, Threat Prevention Reports, URL Reports, and File Type Identification Reports.

The report sample, formatted in XML, is based on your firewall activity in the first 4 hours since you first viewed the report sample. A report sample does not display any entries if the firewall did not find any matching traffic for the report. The firewall only collects new information for a report sample when you restart the firewall and open a report sample.

The figure below shows a report sample for Threat Prevention Reports:



Application Reports, Threat Prevention Reports, URL Reports, and File Type Identification Reports each consist of multiple reports. In the report sample, **Type** describes the name of a report. **Aggregate** lists the log fields that the firewall collects for the report (refer to [Syslog Field Descriptions](#) to determine the name of the fields as they appear in the firewall logs). **Values** indicates the units of measure used in the report (for example, the value **count** for the **Attackers (threat)** report refers to the number of times the firewall detected a threat associated with a particular threat ID).

STEP 4 | View the type of data that the firewall collects for Product Usage Statistics.

Enter the following operational CLI command: **show system info**

STEP 5 | Click **OK** and **Commit** your changes.

STEP 6 | If you enabled Threat Prevention Data and Threat Prevention Packet Captures, view the data that the firewall collected.

1. Edit the Telemetry settings.
2. Click **Download Threat Prevention Data** () to download a tarball file (.tar.gz) with the most recent 100 folders of data that the firewall collected for Threat Prevention Data and Threat Prevention Packet Captures. If you never enabled these settings or if you enabled them but no threat events have matched the conditions for these settings, the firewall does not generate a file and instead returns an error message.

There is currently no way to view the DNS information that the firewall collects through passive DNS monitoring.

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.

Decryption

Palo Alto Networks firewalls can decrypt and inspect traffic to provide visibility into threats and to control protocols, certificate verification, and failure handling. Decryption can enforce policies on encrypted traffic so that the firewall handles encrypted traffic according to your configured security settings. Decrypt traffic to prevent malicious encrypted content from entering your network and sensitive content from leaving your network concealed as encrypted traffic. Enabling decryption can include preparing the keys and certificates required for decryption, creating decryption profiles and policies, and configuring decryption port mirroring.

- > [Decryption Overview](#)
- > [Decryption Concepts](#)
- > [Prepare to Deploy Decryption](#)
- > [Define Traffic to Decrypt](#)
- > [Configure SSL Forward Proxy](#)
- > [Configure SSL Inbound Inspection](#)
- > [Configure SSH Proxy](#)
- > [Configure Server Certificate Verification for Undecrypted Traffic](#)
- > [Decryption Exclusions](#)
- > [Enable Users to Opt Out of SSL Decryption](#)
- > [Temporarily Disable SSL Decryption](#)
- > [Configure Decryption Port Mirroring](#)
- > [Verify Decryption](#)
- > [Decryption Broker](#)
- > [Activate Free Licenses for Decryption Features](#)

Decryption Overview

The Secure Sockets Layer (SSL) and Secure Shell (SSH) encryption protocols 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 certificates to affirm trust between the devices and the keys to decode the data. Decrypt SSL and SSH traffic to:

- Prevent malware concealed as encrypted traffic from being introduced into your network. For example, an attacker compromises a website that uses SSL encryption. Employees visit that website and unknowingly download an exploit or malware. The malware then uses the infected employee endpoint to move laterally through the network and compromise other systems.
- Prevent sensitive information from moving outside the network.
- Ensure the appropriate applications are running on a secure network.
- Selectively decrypt traffic; for example, create a Decryption policy and profile to exclude traffic for financial or healthcare sites from decryption.

Palo Alto Networks firewall decryption is policy-based, and can decrypt, inspect, and control inbound and outbound SSL and SSH connections. A Decryption policy enables you to specify traffic to decrypt by destination, source, service, or URL category, and to block, restrict, or forward the specified traffic according to the security settings in the associated Decryption profile. A Decryption profile controls SSL protocols, certificate verification, and failure checks to prevent traffic that uses weak algorithms or unsupported modes from accessing the network. The firewall uses certificates and keys to decrypt traffic to plaintext, and then enforces App-ID and security settings on the plaintext traffic, including Decryption, Antivirus, Vulnerability, Anti-Spyware, URL Filtering, WildFire, and File-Blocking profiles. After decrypting and inspecting traffic, the firewall re-encrypts the plaintext traffic as it exits the firewall to ensure privacy and security.

The firewall provides three types of Decryption policy rules: [SSL Forward Proxy](#) to control outbound SSL traffic, [SSL Inbound Inspection](#) to control inbound SSL traffic, and [SSH Proxy](#) to control tunneled SSH traffic. You can attach a Decryption profile to a policy rule to apply granular access settings to traffic, such as checks for server certificates, unsupported modes, and failures.

SSL decryption (both forward proxy and inbound inspection) requires certificates to establish the firewall as a trusted third party, and to establish trust between a client and a server to secure an SSL/TLS connection. You can also use certificates when excluding servers from SSL decryption for technical reasons (the site breaks decryption for reasons such as certificate pinning, unsupported ciphers, or mutual authentication). SSH decryption does not require certificates.



Use the [Decryption Best Practices Checklist](#) to plan, implement, and maintain your decryption deployment.

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](#).

You can also use [Decryption Mirroring](#) to forward decrypted traffic as plaintext to a third party solution for additional analysis and archiving.



If you enable Decryption mirroring, be aware of local laws and regulations about what traffic you can mirror and where and how you can store the traffic, because all mirrored traffic, including sensitive information, is forwarded in cleartext.

Decryption Concepts

Review the following topics to learn more about decryption features and support:

- [Keys and Certificates for Decryption Policies](#)
- [SSL Forward Proxy](#)
- [SSL Forward Proxy Decryption Profile](#)
- [SSL Inbound Inspection](#)
- [SSL Inbound Inspection Decryption Profile](#)
- [SSL Protocol Settings Decryption Profile](#)
- [SSH Proxy](#)
- [SSH Proxy Decryption Profile](#)
- [Decryption Profile for No Decryption](#)
- [SSL Decryption for Elliptical Curve Cryptography \(ECC\) Certificates](#)
- [Perfect Forward Secrecy \(PFS\) Support for SSL Decryption](#)
- [SSL Decryption and Subject Alternative Names \(SANs\)](#)
- [High Availability Support for Decrypted Sessions](#)
- [Decryption Mirroring](#)
- [Decryption Broker](#)

Keys and Certificates for Decryption Policies

Keys are strings of numbers typically generated using a mathematical operation involving random numbers and large primes. Keys transform strings—such as passwords and shared secrets—from unencrypted plaintext to encrypted ciphertext and from encrypted ciphertext to unencrypted plaintext. 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 establish trust between a client and a server 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 the 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. A certificate authority (CA) must issue all certificates. After the CA verifies a client or server, the CA issues the certificate and signs it with a private key.



If you have two CAs (Device > Certificate Management > Device Certificates) with the same subject and key, and one CA expires, delete (custom) or disable (predefined) the expired CA. If you do not delete or disable an expired CA, the firewall can build a chain to the expired CA if it is enabled in the trusted chain resulting in a Block page.

When you apply a decryption policy to traffic, 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.

The following table describes the different certificates Palo Alto Networks firewalls use for decryption.

Certificates Used With Decryption	Description
Forward Trust (Used for SSL Forward Proxy decryption)	The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate signed by a CA that the firewall trusts. To configure a Forward Trust certificate on the firewall to present to clients when the server certificate is signed by a trusted CA, see Configure SSL Forward Proxy. 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 Configure the Key Size for SSL Proxy Server certificates. For added security, consider storing the private key associated with the Forward Trust certificate on a hardware security module (see Store Private Keys on an HSM).  <i>Back up the private key associated with the firewall's Forward Trust CA certificate (not the firewall's master key) in a secure repository so that if an issue occurs with the firewall, you can still access the Forward Trust CA certificate. For added security, consider storing the private key associated with the Forward Trust certificate on a hardware security module (see Store Private Keys on an HSM).</i>
Forward Untrust (Used for SSL Forward Proxy decryption)	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 Configure SSL Forward Proxy.
SSL Inbound Inspection	The certificates of the servers on your network for which you want to perform SSL Inbound Inspection of traffic destined for those servers. Import the server certificates onto the firewall.  <i>Beginning in PAN-OS 8.0, firewalls use the Elliptic-Curve Diffie-Hellman Ephemeral (ECDHE) algorithm to perform strict certificate checking. This means that if the firewall uses an intermediate certificate, you must reimport the certificate from your web server to the firewall after you upgrade to a PAN-OS 8.0 or later release and combine the server certificate with the intermediate certificate (install a chained certificate). Otherwise, SSL Inbound Inspection sessions that have an intermediate certificate in the chain will fail. To install a chained certificate:</i> 1. <i>Open each certificate (.cer) file in a plain-text editor such as Notepad.</i> 2. <i>Paste each certificate end-to-end with the Server Certificate at the top with each signer included below.</i>


Certificates Used With Decryption	Description
	3. Save the file as a text (.txt) or certificate (.cer) file (the name of the file cannot contain blank spaces). 4. Import the combined (chained) certificate into the firewall.

SSL Forward Proxy

When you configure the firewall to decrypt SSL traffic going to external sites, it functions as an [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 into your corporate network by decrypting the traffic so that the firewall can apply decryption profiles and security policies and profiles to the traffic.

In SSL Forward Proxy decryption, the firewall is a man-in-the-middle between the internal client and the external server. The firewall uses certificates to transparently represent the client to the server and to transparently represent the server to the client, so that the client believes it is communicating directly with the server (even though the client session is with the firewall), and server believes it is communicating directly with the client (even though the server session is also with the firewall). The firewall uses certificates to establish itself as a trusted third party (man-in-the-middle) for the client-server session (for details on certificates, see [Keys and Certificates for Decryption Policies](#)).

The following figure shows this process in detail. See [Configure SSL Forward Proxy](#) for details on configuring SSL Forward Proxy.




1. The internal client on your network attempts to initiate a TLS session with an external server.
2. The firewall intercepts the client's SSL certificate request. For the client, the firewall acts as the external server, even though the secure session being established is with the firewall, not with the actual server.

3. The firewall then forwards the client's SSL certificate request to the server to initiate a separate session with the server. To the server, the firewall looks like the client, the server doesn't know there's a man-in-the-middle, and the server verifies the certificate.
4. The server sends the firewall a signed certificate intended for the client.
5. The firewall analyzes the server certificate. If the server certificate is signed by a CA that the firewall trusts and meets the policies and profiles you configure, the firewall generates an SSL Forward Trust copy of the server certificate and sends it to the client. If the server certificate is signed by a CA that the firewall does not trust, the firewall generates an SSL Forward Untrust copy of the server certificate and sends it to the client. The certificate copy the firewall generates and sends to the client contains extensions from the original server certificate and is called an *impersonation* certificate because it is not the server's actual certificate. If the firewall does not trust the server, the client sees a block page warning message that the site they're attempting to connect to is not trusted, and if you [Enable Users to Opt Out of SSL Decryption](#), the client can choose to proceed or terminate the session.
6. The client verifies the firewall's impersonation certificate. The client then initiates a session key exchange with the server, which the firewall proxies in the same manner as it proxies the certificates. The firewall forwards the client key to the server, and makes an impersonation copy of the server key for the client, so that firewall remains an "invisible" proxy, the client and server believe their session is with each other, but there are still two separate sessions, one between the client and the firewall, and the other between the firewall and the server. Now all parties have the certificates and keys required and the firewall can decrypt the traffic.
7. All SSL session traffic between goes through the firewall transparently between the client and the server. The firewall decrypts the SSL traffic, applies security policies and profiles and decryption profiles to the traffic, re-encrypts the traffic, and then forwards it.


SSL Forward Proxy Decryption Profile

The SSL Forward Proxy Decryption profile (**Objects > Decryption Profile > SSL Decryption > SSL Forward Proxy**) controls the server verification, session mode checks, and failure checks for outbound traffic defined in SSL Forward Proxy Decryption policies to which you attach the profile. The following figure shows the general best practice recommendations for SSL Forward Proxy Decryption profile settings, but the settings you use also depend on your company's security compliance rules and local laws and regulations. There are also specific best practices for perimeter [internet gateway decryption profiles](#) and for [data center decryption profiles](#).



Server Certificate Verification:

- **Block sessions with expired certificates**—Always check this box to block sessions with servers that have expired certificates and prevent access to potentially insecure sites. If you don't check this box, users can connect with and transact with potentially malicious sites and see warning messages when they attempt to connect, but the connection is not prevented.
- **Block sessions with untrusted issuers**—Always check this box to block sessions with servers that have untrusted certificate issuers. An untrusted issuer may indicate a [man-in-the-middle attack](#), a [replay attack](#), or other attack.
- **Block sessions with unknown certificate status**—Blocks the SSL session when a the certificate revocation status of the server returns with the status “unknown”. Because certificate status may be unknown for multiple reasons, for general decryption security, checking this box usually tightens security too much. However, in higher-security areas of the network such as the data center, checking this box makes sense.
- **Block sessions on certificate status check timeout**—Whether to block sessions if the status check times out depends on your company's security compliance stance because it's a tradeoff between tighter security and a better user experience. Certificate status verification examines the Certificate Revocation List (CRL) on a revocation server or uses Online Certificate Status Protocol (OCSP) to find out if the issuing CA has revoked the certificate and the certificate should not be trusted. However, revocation servers can be slow to respond, which can cause the session to timeout and the firewall to block the session even though the certificate may be valid. If you **Block sessions on certificate status check timeout** and the revocation server is slow to respond, you can use **Device > Setup > Session > Decryption Settings** and click **Certificate Revocation Checking** to change the default timeout value of five seconds to another value. For example, you could increase the timeout value to eight seconds, as shown in the following figure. Enable both CRL and OCSP [certificate revocation checking](#) because server certificates can contain the CRL URL in the CRL Distribution Point (CDP) extension or the OCSP URL in the Authority Information Access (AIA) certificate extension.



- **Restrict certificate extensions**—Checking this box limits the certificate extensions in the server certificate to key usage and extended key usage and blocks certificates with other extensions. However, in certain deployments, some other certificate extensions may be necessary, so only check this box if your deployment requires no other certificate extensions.
- **Append certificate's CN value to SAN extension**—Checking this box ensures that when a browser requires a server certificate to use a Subject Alternative Name (SAN) and doesn't support certificate matching based on the Common Name (CN), if the certificate doesn't have a SAN extension, users can still access the requested web resources because the firewall adds the SAN extension (based on the CN) to the impersonation certificate.

Unsupported Mode Checks. If you don't block sessions with unsupported modes, users receive a warning message if they connect with potentially unsafe servers, and they can click through that message and reach the potentially dangerous site. Blocking these sessions protects you from servers that use weak, risky protocol versions and algorithms:

- **Block sessions with unsupported versions**—When you configure the [SSL Protocol Settings Decryption Profile](#), you specify the minimum version of SSL protocol to allow on your network to reduce the attack

surface by blocking weak protocols. Always check this box to block sessions with the weak SSL protocol versions that you have chosen not to support.

- **Block sessions with unsupported cipher suites**—Always check this box to block sessions if the firewall doesn't support the cipher suite specified in the SSL handshake. You configure which algorithms the firewall supports on the **SSL Protocol Settings** tab of the Decryption profile.
- **Block sessions with client authentication**—If you have no critical applications that require client authentication, block it because firewall can't decrypt sessions that require client authentication. The firewall needs both the client and the server certificates to perform bi-directional decryption, but the firewall only knows the server certificate. This breaks decryption for client authentication sessions. When you check this box, the firewall blocks all sessions with client authentication except sessions from sites on the SSL Decryption Exclusion list (**Device > Certificate Management > SSL Decryption Exclusion**).

If you don't **Block sessions with client authentication**, when the firewall attempts to decrypt a session that uses client authentication, the firewall allows the session and adds an entry that contains the server URL/IP address, the application, and the Decryption profile to its Local SSL Decryption Exclusion Cache. Entries remain in the local cache for 12 hours and then age out. If the same user or a different user attempts to access the server within 12 hours using client authentication, the firewall matches the session to the SSL Decryption Exclusion Cache entry, does not attempt to decrypt the traffic, and allows the encrypted session.

If the SSL Decryption Exclusion Cache becomes full, the firewall purges the oldest entries as new entries arrive. If you change the Decryption policy or profile, the firewall flushes the local exclusion cache because changing the policy or profile can change the classification outcome of the session.



You may need to allow traffic on your network from other sites that use client authentication in addition to the Predefined sites on the SSL Decryption Exclusion list. Create a Decryption profile that allows sessions with client authentication. Add it to a Decryption policy rule that applies only to the server(s) that house the application. To increase security even more, you can require Multi-Factor Authentication to complete the user login process.

Failure Checks:

- **Block sessions if resources not available**—If you don't block sessions when firewall processing resources aren't available, then encrypted traffic that you want to decrypt enters the network still encrypted, risking allowing potentially dangerous connections. However, blocking sessions when firewall processing resources aren't available may affect the user experience by making sites that users normally can reach temporarily unreachable. Whether to implement failure checks depends on your company's security compliance stance and the importance to your business of the user experience, weighed against tighter security. Alternatively, consider using firewall models with more processing power so that you can decrypt more traffic.
- **Block sessions if HSM not available**—If you use a Hardware Security Module (HSM) to store your private keys, whether you use one depends on your compliance rules about where the private key must come from and how you want to handle encrypted traffic if the HSM isn't available. For example, if your company mandates the use of an HSM for private key signing, then block sessions if the HSM isn't available. However, if your company is less strict about this, then you can consider not blocking sessions if the HSM isn't available. (If the HSM is down, the firewall can process decryption for sites for which it has cached the response from the HSM, but not for other sites.) The best practice in this case depends on your company's policies. If the HSM is critical to your business, run the HSM in a high-availability (HA) pair (PAN-OS 8.1 supports two members in an HSM HA pair).

SSL Inbound Inspection

Use SSL Inbound Inspection to decrypt and inspect inbound SSL/TLS traffic from a client to a targeted network server (any server you have the certificate for and can import it onto the firewall) and block

suspicious sessions. 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 ensure that the sensitive data does not move outside the secure company network by blocking or restricting the session.

On the firewall, you must [install the certificate](#) and private key for each server for which you want to perform SSL inbound inspection. You must also install the public key certificate as well as the private key on each firewall that performs SSL inbound inspection. The way the firewall performs SSL inbound inspection depends on the type of key negotiated, Rivest, Shamir, Adleman (RSA) or Perfect Forward Secrecy (PFS).

For RSA keys, the firewall performs SSL inbound inspection without terminating the connection. As the encrypted session flows through the firewall, the firewall transparently makes a copy of it and decrypts it so that the firewall can apply the appropriate policy to the traffic.




When you configure the [SSL Protocol Settings Decryption Profile](#) for SSL Inbound Inspection traffic, create separate profiles for servers with different security capabilities. For example, if one set of servers supports only RSA, the SSL Protocol Settings only need to support RSA. However, the SSL Protocol Settings for servers that support PFS should support PFS. Configure SSL Protocol Settings for the highest level of security that the server supports, but check performance to ensure that the firewall resources can handle the higher processing load that higher security protocols and algorithms require.

For PFS keys using the Diffie-Hellman exchange (DHE) or Elliptic Curve Diffie-Hellman exchange (ECDHE), the firewall acts as a man-in-the-middle proxy between the external client and the internal server. Because PFS generates a new key with every session, the firewall can't simply copy and decrypt the inbound SSL flow as it passes through, the firewall must act as a proxy device.

The following figure shows how SSL Inbound Inspection works when the key exchange algorithm is RSA. When the key exchange algorithm is PFS, the firewall functions as a proxy (creates a secure session between the client and the firewall and another secure session between the firewall and the server) and must generate a new session key for each secure session.


See [Configure SSL Inbound Inspection](#) for details on enabling this feature.



SSL Inbound Inspection Decryption Profile

The SSL Inbound Inspection Decryption profile (**Objects > Decryption Profile > SSL Decryption > SSL Inbound Inspection**) controls the session mode checks and failure checks for inbound traffic defined in the

SSL Inbound Inspection Decryption policies to which you attach the profile. The following figure shows the general best practice recommendations for SSL Inbound Inspection Decryption profile settings, but the settings you use also depend on your company's security compliance rules and local laws and regulations.



Unsupported Mode Checks. If you don't block sessions with unsupported modes, users receive a warning message if they connect with potentially unsafe servers, and they can click through that message and reach the potentially dangerous site. Blocking these sessions protects you from servers that use weak, risky protocol versions and algorithms:

1. **Block sessions with unsupported versions**—When you configure the [SSL Protocol Settings Decryption Profile](#), you specify the minimum version of SSL protocol to allow on your network to reduce the attack surface by blocking weak protocols. Always check this box to block sessions with the weak SSL protocol versions that you have chosen not to support.
2. **Block sessions with unsupported cipher suites**—Always check this box to block sessions if the firewall doesn't support the cipher suite specified in the SSL handshake. You configure which algorithms the firewall supports on the **SSL Protocol Settings** tab of the Decryption profile.

Failure Checks:

- **Block sessions if resources not available**—If you don't block sessions when firewall processing resources aren't available, then encrypted traffic that you want to decrypt enters the network still encrypted, risking allowing potentially dangerous connections. However, blocking sessions when firewall processing resources aren't available may affect the user experience by making sites that users normally can reach temporarily unreachable. Whether to implement failure checks depends on your company's security compliance stance and the importance to your business of the user experience, weighed against tighter security. Alternatively, consider using firewall models with more processing power so that you can decrypt more traffic.
- **Block sessions if HSM not available**—If you use a Hardware Security Module (HSM) to store your private keys, whether you use one depends on your compliance rules about where the private key must come from and how you want to handle encrypted traffic if the HSM isn't available. For example, if your company mandates the use of an HSM for private key signing, then block sessions if the HSM isn't available. However, if your company is less strict about this, then you can consider not blocking sessions if the HSM isn't available. (If the HSM is down, the firewall can process decryption for sites for which it has cached the response from the HSM, but not for other sites.) The best practice in this case depends on your company's policies. If the HSM is critical to your business, run the HSM in a high-availability (HA) pair (PAN-OS 8.1 supports two members in an HSM HA pair).

SSL Protocol Settings Decryption Profile

The SSL Protocol Settings (**Objects > Decryption Profile > SSL Decryption > SSL Protocol Settings**) control whether you allow vulnerable SSL/TLS protocol versions, weak encryption algorithms, and weak authentication algorithms. SSL Protocol Settings apply to outbound SSL Forward Proxy and inbound SSL Inbound Inspection traffic. These settings don't apply to SSH Proxy traffic or to traffic that you don't decrypt.

The following figure shows the general best practice recommendations for SSL Protocol Settings. There are also specific best practices for perimeter [internet gateway decryption profiles](#) and for [data center decryption profiles](#).



When you configure SSL Protocol Settings for SSL Inbound Inspection traffic, create separate profiles for servers with different security capabilities. For example, if one set of servers supports only RSA, the SSL Protocol Settings only need to support RSA. However, the SSL Protocol Settings for servers that support PFS should support PFS. Configure SSL Protocol Settings for the highest level of security that the target server you are protecting supports, but check performance to ensure that the firewall resources can handle the higher processing load that higher security protocols and algorithms require.

The screenshot shows the 'Decryption Profile' configuration window with the 'SSL Protocol Settings' tab selected. The 'Name' field is 'best-practice-ssl-decryption'. The 'SSL Decryption' section has 'No Decryption' and 'SSH Proxy' buttons. The 'SSL Forward Proxy' and 'SSL Inbound Inspection' buttons are also present. The 'Protocol Versions' section has 'Min Version' set to 'TLSv1.2' and 'Max Version' set to 'Max'. The 'Key Exchange Algorithms' section has checkboxes for RSA, DH, and ECDH, all of which are checked. The 'Encryption Algorithms' section has checkboxes for 3DES, RC4, AES128-CBC, AES256-CBC, AES128-GCM, and AES256-GCM. The 3DES and RC4 checkboxes are unchecked, while the others are checked. The 'Authentication Algorithms' section has checkboxes for MD5, SHA1, SHA256, and SHA384, all of which are checked. A note at the bottom states: 'Note: For unsupported modes and features, the session information is cached for 12 hours, so future sessions between the same host and server pair are not decrypted. Check boxes to Start. Decryption enabled.' The 'OK' and 'Cancel' buttons are at the bottom right.

Protocol Versions:

- Set the **Min Version** to **TLSv1.2** to provide the strongest security—business sites that value security support TLSv1.2. If a site (or a category of sites) only supports weaker ciphers, review the site and determine if it really houses a legitimate business application. If it does, make an exception for only that site by configuring a Decryption profile with a **Min Version** that matches the strongest cipher the site supports and then applying the profile to a Decryption policy rule that limits allowing the weak cipher to only the site or sites in question. If the site doesn't house a legitimate business application, don't weaken your security posture to support the site—weak protocols (and ciphers) contain known vulnerabilities that attackers can exploit. If the site belongs to a category of sites that you don't need for business purposes, use [URL Filtering](#) to block access to the entire category. Don't support weak encryption or authentication algorithms unless you must do so to support important legacy sites, and when you make exceptions, create a separate Decryption profile that allows the weaker protocol just for those sites. Don't downgrade the main Decryption profile that you apply to most sites to TLSv1.1 just to accommodate a few exceptions.



Qualys SSL Labs [SSL Pulse](#) web page provides up-to-date statistics on the percentages of different ciphers and protocols in use on the 150,000 most popular sites in the world so you can see trends and understand how widespread worldwide support is for more secure ciphers and protocols.

- Set the **Max Version** to **Max** rather than to a particular version so that as the protocols improve, the firewall automatically supports the newest and best protocols. Whether you intend to attach a Decryption profile to a Decryption policy rule that governs inbound (SSL Inbound Inspection) or outbound (SSL Forward Proxy) traffic, avoid allowing weak algorithms.

Key Exchange Algorithms: Leave all three boxes checked (default) to support both RSA and **PFS** (DHE and ECDHE) key exchanges.



To support HTTP/2 traffic, you must leave the ECDHE box checked.

Encryption Algorithms: When you set the protocol version to TLSv1.2, the older, weaker 3DES and RC4 algorithms are automatically unchecked (blocked). For any traffic for which you must allow a weaker TLS protocol, create a separate Decryption profile and apply it only to traffic for that site, and uncheck the 3DES and RC4 boxes. Do not allow traffic that uses the 3DES or RC4 algorithms. If unchecking the 3DES or RC4 boxes prevents you from accessing a site that you must use for business, create a separate Decryption profile for that site. Don't weaken decryption for any other sites.

Authentication Algorithms: The older, weaker MD5 algorithm is automatically unchecked (blocked). Do not allow MD5 authenticated traffic on your network; SHA1 is the weakest authentication algorithm you should allow. If no necessary sites use SHA1, uncheck the box and block traffic to further reduce the attack surface.

SSH Proxy

In an SSH Proxy configuration, the firewall resides between a client and a server. SSH Proxy enables the firewall to decrypt inbound and outbound SSH connections and ensures that attackers don't use SSH to tunnel unwanted applications and content. SSH decryption does not require certificates and the firewall automatically generates the key used for SSH decryption when the firewall boots up. During the boot up process, the firewall checks if there is an existing key. If not, the firewall generates a key. The firewall uses the key to decrypt SSH sessions for all virtual systems configured on the firewall and all SSH v2 sessions.

SSH allows tunneling, which can hide malicious traffic from decryption. The firewall can't decrypt traffic inside an SSH tunnel. You can block all SSH tunnel traffic by configuring a Security policy rule for the application **ssh-tunnel** with the **Action** set to **Deny** (along with a Security policy rule to allow traffic from the **ssh** application).

SSH tunneling sessions can tunnel X11 Windows packets and TCP packets. One SSH connection may contain multiple channels. When you apply an SSH Decryption profile to traffic, for each channel in the connection, the firewall examines the App-ID of the traffic and identifies the channel type. The channel type can be:


- session
- X11
- forwarded-tcpip
- direct-tcpip

When the channel type is session, the firewall identifies the traffic as allowed SSH traffic such as SFTP or SCP. When the channel type is X11, forwarded-tcpip, or direct-tcpip, the firewall identifies the traffic as SSH tunneling traffic and blocks it.



Limit SSH use to administrators who need to manage network devices, log all SSH traffic, and consider configuring [Multi-Factor Authentication](#) to help ensure that only legitimate users can use SSH to access devices, which reduces the attack surface.

The following figure shows how SSH Proxy decryption works. See [Configure SSH Proxy](#) for how to enable SSH Proxy decryption.



When the client sends an SSH request to the server to initiate a session, the firewall intercepts the request and forwards it to the server. The firewall then intercepts the server response and forwards it to the client. This establishes two separate SSH tunnels, one between the firewall and the client and one between the firewall and the server, with firewall functioning as a proxy. As traffic flows between the client and the server, the firewall checks whether the SSH traffic is being routed normally or if it is using SSH tunneling (port forwarding). The firewall doesn't perform content and threat inspection on SSH tunnels; however, if the firewall identifies SSH tunnels, it blocks the SSH tunneled traffic and restricts the traffic according to configured security policies.

SSH Proxy Decryption Profile

The SSH Proxy Decryption profile (**Objects > Decryption Profile > SSH Proxy**) controls the session mode checks and failure checks for SSH traffic defined in the SSH Proxy Decryption policies to which you attach the profile. The following figure shows the general best practice recommendations for SSH Proxy Decryption profile settings, but the settings you use also depend on your company's security compliance rules and local laws and regulations.



The firewall doesn't perform content and threat inspection on SSH tunnels (port forwarding). However, the firewall distinguishes between the SSH application and the SSH-tunnel application. If the firewall identifies SSH tunnels, it blocks the SSH tunneled traffic and restricts the traffic according to configured security policies.



Unsupported Mode Checks. The firewall supports SSHv2. If you don't block sessions with unsupported modes, users receive a warning message if they connect with potentially unsafe servers, and they can click through that message and reach the potentially dangerous site. Blocking these sessions protects you from servers that use weak, risky protocol versions and algorithms:

1. **Block sessions with unsupported versions**—The firewall has a set of predefined supported versions. Checking this box blocks traffic with weak versions. Always check this box to block sessions with the weak protocol versions to reduce the attack surface.
2. **Block sessions with unsupported algorithms**—The firewall has a set of predefined supported algorithms. Checking this box blocks traffic with weak algorithms. Always check this box to block sessions with unsupported algorithms to reduce the attack surface.

Failure Checks:

- **Block sessions on SSH errors**—Checking this box terminates the session if SSH errors occur.
- **Block sessions if resources not available**—If you don't block sessions when firewall processing resources aren't available, then encrypted traffic that you want to decrypt enters the network still encrypted, risking allowing potentially dangerous connections. However, blocking sessions when firewall processing resources aren't available may affect the user experience by making sites that users normally can reach temporarily unreachable. Whether to implement failure checks depends on your company's security compliance stance and the importance to your business of the user experience, weighed against tighter security. Alternatively, consider using firewall models with more processing power so that you can decrypt more traffic.

Decryption Profile for No Decryption

The No Decryption profile (**Objects > Decryption Profile > No Decryption**) controls server verification checks for traffic that you choose not to decrypt as defined in "No Decryption" Decryption policies to which you attach the profile. (Don't exclude traffic that you can't decrypt because a site breaks decryption for technical reasons such as a pinned certificate or mutual authentication by policy. Instead, add the hostname to the [Decryption Exclusion List](#).) The following figure shows the general best practice recommendations for the No Decryption profile settings, but the settings you use also depend on your company's security compliance rules and local laws and regulations.



- **Block sessions with expired certificates**—Always check this box to block sessions with servers that have expired certificates and prevent access to potentially insecure sites. If you don't check this box, users can connect with and transact with potentially malicious sites and see warning messages when they attempt to connect, but the connection is not prevented.
- **Block sessions with untrusted issuers**—Always check this box to block sessions with servers that have untrusted certificate issuers. An untrusted issuer may indicate a [man-in-the-middle attack](#), a [replay attack](#), or other attack.



(Applies to PAN-OS 9.0.9 and later for TLSv1.2 and earlier) If you choose to allow sessions with untrusted issuers (not recommended) and only Block sessions with expired certificates, there is a scenario in which a session with a trusted, expired issuer may be blocked inadvertently. When the firewall's certificate store contains a valid, self-signed Trusted CA and the server sends an expired CA in the certificate chain, the firewall does not check its certificate store. Instead, the firewall blocks the session based on the expired CA when it should find the trusted, valid alternative trust anchor and allow the session based on that trusted self-signed certificate.

To avoid this scenario, in addition to Block sessions with expired certificates, enable Block sessions with untrusted issuers. This forces the firewall to check its certificate store, find the self-signed Trusted CA, and allow the session.

SSL Decryption for Elliptical Curve Cryptography (ECC) Certificates

The firewall automatically decrypts SSL traffic from websites and applications using ECC certificates, including Elliptical Curve Digital Signature Algorithm (ECDSA) certificates. As organizations transition to using ECC certificates to benefit from the strong keys and small certificate size, you can continue to maintain visibility into and safely enable ECC-secured application and website traffic.



Decryption for websites and applications using ECC certificates is not supported for traffic that is mirrored to the firewall; encrypted traffic using ECC certificates must pass through the firewall directly for the firewall to decrypt it.

You cannot use a [hardware security module \(HSM\)](#) to store the private keys associated with ECDSA certificates.

Perfect Forward Secrecy (PFS) Support for SSL Decryption


PFS is a secure communication protocol that prevents the compromise of one encrypted session from leading to the compromise of multiple encrypted sessions. With PFS, a server generates unique private keys for each secure session it establishes with a client. If a server private key is compromised, only the single session established with that key is vulnerable—an attacker cannot retrieve data from past and

future sessions because the server establishes each connected with a uniquely generated key. The firewall decrypts SSL sessions established with PFS key exchange algorithms, and preserves PFS protection for past and future sessions.

Support for Diffie-Hellman (DHE)-based PFS and elliptical curve Diffie-Hellman (ECDHE)-based PFS is enabled by default (**Objects > Decryption Profile > SSL Decryption > SSL Protocol Settings**).




If you use the DHE or ECDHE key exchange algorithms to enable PFS support for SSL decryption, you can use a [hardware security module \(HSM\)](#) to store the private keys for SSL Inbound Inspection.



SSL Decryption and Subject Alternative Names (SANs)

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. 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.

To enable SAN support for decrypted SSL traffic, update the decryption profile attached to the relevant decryption policy: select **Objects > Decryption Profile > SSL Decryption > SSL Forward Proxy > Append Certificate's CN Value to SAN Extension**).



High Availability Support for Decrypted Sessions

High Availability (HA) syncs are supported for inbound, decrypted SSL sessions, if the sessions were established using non-PFS key exchange algorithms. When a failover occurs, the passive device continues to inspect and enforce the decrypted traffic.

HA syncs are not supported for:

- decrypted SSL sessions (both inbound and outbound) that were established using PFS key exchange algorithms
- decrypted, outbound SSL sessions using non-PFS key exchange algorithms

In these cases, when a failover occurs, the passive device allows transferred sessions without decrypting them. New sessions will then continue to be decrypted based on your decryption policy.

The following table details HA support for decrypted sessions:

	PFS Protected Session	Non-PFS Protected Session
Inbound SSL Session (Inbound Inspection Decryption)	— No HA Sync	✓ HA Sync
Outbound SSL Sessions (SSL Forward Proxy Decryption)	— No HA Sync	— No HA Sync

Decryption Mirroring

Decryption mirroring creates a copy of decrypted traffic from a firewall and sends it to a traffic collection tool such as NetWitness or Solera, which can receive raw packet captures for archiving and analysis. Organizations that require comprehensive data capture for forensic and historical purposes or for data leak prevention (DLP) can install a free license to enable the feature.

After you install the license, connect the traffic collection tool directly to an Ethernet interface on the firewall and set the **Interface Type** to **Decrypt Mirror**. The firewall simulates a TCP handshake with the collection tool and then sends every data packet through that interface, decrypted (as cleartext).



Decryption port mirroring is not available on the VM-Series for public cloud platforms (AWS, Azure, Google Cloud Platform) and VMware NSX.

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.

The following graphic shows the process for mirroring decrypted traffic and the section [Configure Decryption Port Mirroring](#) describes how to license and enable this feature.



Prepare to Deploy Decryption

The most time-consuming part of deploying decryption isn't configuring the decryption policies and profiles, it's preparing for the deployment by working with stakeholders to decide what traffic to decrypt and not to decrypt, educating your user population about changes to website access, developing a private key infrastructure (PKI) strategy, and planning a staged, prioritized rollout.

Set goals for decryption and review [Decryption planning best practices checklist](#) to ensure that you understand the recommended best practices. The best practice goal is to decrypt as much traffic as your firewall resources permit and decrypt the most important traffic first.



Migrate from port-based to application-based [Security](#) policy rules before you create and deploy Decryption policy rules. If you create Decryption rules based on port-based Security policy and then migrate to application-based Security policy, the change could cause the Decryption rules to block traffic that you intend to allow because Security policy rules are likely to use application default ports to prevent application traffic from using non-standard ports. For example, traffic identified as web-browsing application traffic (default port 80) may have underlying applications that have different default ports, such as HTTPS traffic (default port 443). The application-default rule blocks the HTTPS traffic because it sees the decrypted traffic using a "non-standard" port (443 instead of 80). Migrating to App-ID based rules before deploying decryption means that when you test your decryption deployment in POCs, you'll discover Security policy misconfiguration and fix it before rolling it out to the general user population.

To prepare to deploy Decryption:

- [Work with Stakeholders to Develop a Decryption Deployment Strategy](#)
- [Develop a PKI Rollout Plan](#)
- [Size the Decryption Firewall Deployment](#)
- [Plan a Staged, Prioritized Deployment](#)

Work with Stakeholders to Develop a Decryption Deployment Strategy

Work with stakeholders such as legal, finance, HR, executives, security, and IT/support to develop a decryption deployment strategy. Start by getting the required approvals to decrypt traffic to secure the corporation. Decrypting traffic involves understanding how legal regulations and business needs affect what you can and can't decrypt.

Identify and prioritize the traffic you want to decrypt. The best practice is to decrypt as much traffic as you can to gain visibility into potential threats in encrypted traffic and prevent those threats. If incorrect firewall sizing prevents you from decrypting all of the traffic you want to decrypt, prioritize the most critical servers, the highest-risk traffic categories, and less trusted segments and IP subnets. To help prioritize, ask yourself questions such as, "What happens if this server is compromised?" and "How much risk am I willing to take in relation to the level of performance I want to achieve?"

Next, identify traffic that you can't decrypt because the traffic breaks decryption for technical reasons such as a pinned certificate, an incomplete certificate chain, unsupported ciphers, or mutual authentication. Decrypting sites that break decryption technically results in blocking that traffic. Evaluate the websites that break decryption technically and ask yourself if you need access to those sites for business reasons. If you don't need access to those sites, allow decryption to block them. If you need access to any of those sites for business purposes, add them to the SSL Decryption [Exclusion](#) list to except them from decryption. The SSL Decryption Exclusion list is exclusively for sites that break decryption technically.

Identify sensitive traffic that you *choose* not to decrypt for legal, regulatory, personal, or other reasons, such as financial, health, or government traffic, or the traffic of certain executives. This is not traffic that breaks decryption technically, so you don't use the SSL Decryption Exclusion list to except this traffic from decryption. Instead, you [Create a Policy-Based Decryption Exclusion](#) to identify and control traffic you choose not to decrypt and apply the No Decryption decryption profile to the policy to prevent servers with certificate issues from accessing the network. Policy-based decryption exclusions are only for traffic you choose not to decrypt.

When you plan decryption policy, consider your company's security compliance rules, computer usage policy, and your business goals. Extremely strict controls can impact the user experience by preventing access to non-business sites the user used to access, but may be required for government or financial institutions. There is always a tradeoff between usability, management overhead, and security. The tighter the decryption policy, the greater the chance that a website will become unreachable, which may result in user complaints and possibly modifying the rulebase.



Although a tight decryption policy may initially cause a few user complaints, those complaints can draw your attention to unsanctioned or undesirable websites that are blocked because they use weak algorithms or have certificate issues. Use complaints as a tool to better understand the traffic on your network.

Different groups of users and even individual users may require different decryption policies, or you may want to apply the same decryption policy to all users. For example, executives may be exempted from decryption policies that apply to other employees. And you may want to apply different decryption policies to employee groups, contractors, partners, and guests. Prepare updated legal and HR computer usage policies to distribute to all employees, contractors, partners, guests, and any other network users so that when you roll out decryption, users understand their data can be decrypted and scanned for threats.



How you handle guest users depends on the access they require. Isolate guests from the rest of your network by placing them on a separate VLAN and on a separate SSID for wireless access. If guests don't need to access your corporate network, don't let them on it and there will be no need to decrypt their traffic. If guests need to access your corporate network, decrypt their traffic:

- *Enterprises don't control guest devices. Decrypt guest traffic and subject it to your guest Security policy so the firewall can inspect the traffic and prevent threats. To do this, redirect guest users through a captive portal, instruct them how to download and install the CA certificate, and clearly notify guests that their traffic will be decrypted. Include the process in your company's privacy and computer usage policy.*
- *Create separate Decryption [policy](#) rules and Security policy rules to tightly control guest access so that guests can only access the areas of your network that they need to access.*

Similarly to different groups of users, decide which devices to decrypt and which applications to decrypt. Today's networks support not only corporate devices, but BYOD, mobile, remote-user and other devices, including contractor, partner, and guest devices. Today's users attempt to access many sites, both sanctioned and unsanctioned, and you should decide how much of that traffic you want to decrypt.



Enterprises don't control BYOD devices. If you allow BYOD devices on your network, decrypt their traffic and subject it to the same Security policy that you apply to other network traffic so the firewall can inspect the traffic and prevent threats. To do this, redirect BYOD users through a captive portal, instruct them how to download and install the CA certificate, and clearly notify users that their traffic will be decrypted. Educate BYOD users about the process and include it in your company's privacy and computer usage policy.

Decide what traffic you want to log and investigate what traffic you can log. Be aware of local laws regarding what types of data you can log and store, and where you can log and store the data. For example, local laws may prevent logging and storing personal information such as health and financial data.

Decide how to handle bad certificates. For example, will you block or allow sessions for which the certificate status is unknown? Understanding how you want to handle bad certificates determines how you configure the decryption profiles that you attach to decryption policies to control which sessions you allow based on the server certificate verification status.

Develop a PKI Rollout Plan

Plan how to roll out your [public key infrastructure](#) (PKI). Network devices need an SSL Forward Trust CA certificate for trusted sites and an SSL Forward Untrust CA certificate for untrusted sites. Generate separate Forward Trust and Forward Untrust certificates (do not sign the Forward Untrust certificate with the Enterprise Root CA because you want the Untrust certificate to warn users that they are trying to access potentially unsafe sites). Palo Alto Networks next-generation firewalls have two methods of generating CA certificates for SSL decryption:

- **Generate the SSL CA certificates from your Enterprise Root CA as subordinate certificates**—If you have an existing Enterprise PKI, this is the best practice. Generating a subordinate certificate from your Enterprise Root CA makes the rollout easier and smoother because network devices already trust the Enterprise Root CA, so you avoid any certificate issues when you begin the deployment phase. If you don't have an Enterprise Root CA, consider getting one.
- **Generate a self-signed Root CA certificate on the firewall and create subordinate CA certificates on that firewall**—If you don't have an Enterprise Root CA, this method provides a self-signed Root CA certificate and the subordinate Forward Trust and Untrust CA certificates. With this method, you need to install the self-signed certificates on all of your network devices so that those devices recognize the firewall's self-signed certificates. Because the certificates must be deployed to all devices, this method is better for small deployments and proof-of-concept (POC) trials than for large deployments.



Do not export the Forward Untrust certificate to the Certificate Trust Lists of your network devices! This is critical because installing the Untrust certificate in the Trust List results in devices trusting websites that the firewall does not trust. In addition, users won't see certificate warnings for untrusted sites, so they won't know the sites are untrusted and may access those sites, which could expose your network to threats.



Regardless of whether you generate Forward Trust certificates from your Enterprise Root CA or use a self-signed certificate generated on the firewall, generate a separate subordinate Forward Trust CA certificate for each firewall. The flexibility of using separate subordinate CAs enables you to [revoke](#) one certificate when you decommission a device (or device pair) without affecting the rest of the deployment and reduces the impact in any situation in which you need to revoke a certificate. Separate Forward Trust CAs on each firewall also helps troubleshoot issues because the CA error message the user sees includes information about the firewall the traffic is traversing. If you use the same Forward Trust CA on every firewall, you lose the granularity of that information.

There is no benefit to using different Forward Untrust certificates on different firewalls, so you can use the same Forward Untrust certificate on all firewalls. If you need additional security for your private keys, consider [storing them on an HSM](#).

You may need to make special accommodations for guest users. If guest users don't need access to your corporate network, don't allow access, and then you won't have to decrypt their traffic or create infrastructure to support guest access. If you need to support guest users, discuss with your legal department whether you can decrypt guest traffic.

If you can decrypt guest traffic, treat guests similarly to the way you treat BYOD devices. Decrypt guest traffic and subject it to the same Security policy that you apply to other network traffic. Do this by redirecting guest users through a captive portal, instruct them how to download and install the CA certificate, and clearly notify users that their traffic will be decrypted. Include the process in your company's privacy and computer usage policy. In addition, restrict guest traffic to only the areas guests need to access.

If you can't decrypt guest traffic for legal reasons, then isolate guest traffic and prevent it from moving laterally in your network:

- Create a separate zone for guests and restrict guest access to that zone. To prevent lateral movement, don't allow guest access to other zones.
- Allow only sanctioned applications, use URL filtering to prevent access to risky URL categories, and apply the [best practice Security profiles](#).
- Apply a [No Decrypt decryption policy and profile](#) to prevent guests from accessing websites with unknown or expired CAs.

All employees, contractors, partners, and other users should use your normal corporate infrastructure and you should decrypt and inspect their traffic.

Size the Decryption Firewall Deployment

Decrypting encrypted traffic consumes firewall CPU resources and can affect throughput. In general, the tighter the security (the more SSL traffic you decrypt combined with the more stringent your protocol settings), the more firewall resources decryption consumes. Work with your Palo Alto Networks SE/CE to size your firewall deployment and avoid sizing mistakes. Factors that affect decryption resource consumption and therefore how much traffic the firewall can decrypt include:

- The amount of SSL traffic you want to decrypt. This varies from network to network. For example, some applications must be decrypted to prevent the injection of malware or exploits into the network or unauthorized data transfers, some applications can't be decrypted due to local laws and regulations or business reasons, and other applications are cleartext (unencrypted) and don't need to be decrypted. The more traffic you want to decrypt, the more resources you need.
- The TLS protocol version. Higher versions are more secure but consume more resources. Use the highest TLS protocol version you can to maximize security.
- The key size. The larger the key size, the better the security, but also the more resources the key processing consumes.
- The key exchange algorithm. Perfect Forward Secrecy (PFS) ephemeral key exchange algorithms such as Diffie-Hellman Ephemeral (DHE) Elliptic-Curve Diffie-Hellman Exchange (ECDHE) consume more processing resources than Rivest-Shamir-Adleman (RSA) algorithms. PFS key exchange algorithms provide greater security than RSA key exchange algorithms because the firewall has to generate a new cipher key for each session—but generating the new key consumes more firewall resources. However, if an attacker compromises a session key, PFS prevents the attacker from using it to decrypt any other sessions between the same client and server and RSA does not.
- The encryption algorithm. The key exchange algorithm determines whether the encryption algorithm is PFS or RSA.
- The certificate authentication method. RSA (not the RSA key exchange algorithm) consumes less resources than Elliptic Curve Digital Signature Algorithm (ECDSA) but ECDSA is more secure.



The combination of the key exchange algorithm and the certificate authentication method affect throughput performance as shown in [RSA and ECDSA benchmark tests](#). The performance cost of PFS trades off against the higher security that PFS achieves, but PFS may not be needed for all types of traffic. You can save firewall CPU cycles by using RSA for traffic that you want to decrypt and inspect for threats but that isn't sensitive.

- Average transaction sizes. For example, small average transaction sizes consume more processing power to decrypt. Measure the average transaction size of all traffic, then measure the average transaction size

of traffic on port 443 (the default port for HTTPS encrypted traffic) to understand the proportion of encrypted traffic going to the firewall in relation to your total traffic and the average transaction sizes. Eliminate anomalous outliers such as unusually large transactions to get a truer measurement of average transaction size.

- The firewall model and resources. Newer firewall models have more processing power than older models.

The combination of these factors determines how decryption consumes firewall processing resources. To best utilize the firewall's resources, understand the risks of the data you're protecting. If firewall resources are an issue, use stronger decryption for higher-priority traffic and use less processor-intensive decryption to decrypt and inspect lower-priority traffic until you can increase the available resources. For example, you could use RSA instead of ECDHE and ECDSA for traffic that isn't sensitive or high-priority to preserve firewall resources for using PFS-based decryption for higher priority, sensitive traffic. (You're still decrypting and inspecting the lower-priority traffic, but trading off consuming fewer computational resources with using algorithms that aren't as secure as PFS.) The key is to understand the risks of different traffic types and treat them accordingly.

Measure firewall performance so that you understand the currently available resources, which helps you understand whether you need more firewall resources to decrypt the traffic you want to decrypt. Measuring firewall performance also sets a baseline for performance comparisons after deploying decryption.

When you size the firewall deployment, base it not only on your current needs, but also on your future needs. Include headroom for the growth of decryption traffic because Gartner predicts that through 2019, more than 80 percent of enterprise web traffic will be encrypted, and more than 50 percent of new malware campaigns will use various forms of encryption. Work with your Palo Alto Networks representatives and take advantage of their experience in sizing firewalls to help you size your firewall decryption deployment.

Plan a Staged, Prioritized Deployment

Plan to roll out decryption in a controlled manner, piece by piece. Don't roll out your entire decryption deployment at one time. Test and ensure that decryption is working as planned and that users understand what you are doing and why. Rolling out decryption in this manner makes it easier to troubleshoot in case anything doesn't work as expected and helps users adjust to the changes.

Educating stakeholders, employees, and other users such as contractors and partners is critical because decryption settings may change their ability to access some websites. Users should understand how to respond to situations in which previously reachable websites become unreachable and what information to give technical support. Support should understand what is being rolled out when and how to help users who encounter issues. Before you roll out decryption to the general population:

- Identify early adopters to help champion decryption and who will be able to help other employees who have questions during the full rollout. Enlist the help of department managers and help them understand the benefits of decrypting traffic.
- Set up proof-of-concept (POC) trials in each department with early adopters and other employees who understand why decrypting traffic is important. Educate POC participants about the changes and how to contact technical support if they run into issues. In this way, decryption POCs become an opportunity to work with technical support to POC how to support decryption and to develop the most painless method for supporting the general rollout. The interaction between POC users and technical support also allows you to fine-tune policies and how to communicate with users.

POCs enable you to experiment with prioritizing what to decrypt first, so that when you phase in decryption in the general population, your POC experience helps you understand how to phase in decrypting different URL Categories. Measure the way decryption affects firewall CPU and memory utilization to help understand if the firewall sizing is correct or if you need to upgrade. POCs can also

reveal applications that break decryption technically (decrypting them blocks their traffic) and need to be added to the Decryption Exclusion list.

When you set up POCs, also set up a user group that can certify the operational readiness and procedures prior to the general rollout.

- Educate the user population before the general rollout, and plan to educate new users as they join the company. This is a critical phase of deploying decryption because the deployment may affect websites that users previously visited but are not safe, so those sites are no longer reachable. The POC experience helps identify the most important points to communicate.
- Phase in decryption. You can accomplish this several ways. You can decrypt the highest priority traffic first (for example, the URL Categories most likely to harbor malicious traffic, such as gaming) and then decrypt more as you gain experience. Alternatively, you can take a more conservative approach and decrypt the URL Categories that don't affect your business first (so if something goes wrong, no issues occur that affect business), for example, news feeds. In all cases, the best way to phase in decryption is to decrypt a few URL Categories, take user feedback into account, run reports to ensure that decryption is working as expected, and then gradually decrypt a few more URL Categories and verify, and so on. Plan to make [Decryption Exclusions](#) to exclude sites from decryption if you can't decrypt them for technical reasons or because you choose not to decrypt them.

If you [Enable Users to Opt Out of SSL Decryption](#) (users see a response page that allows them either to opt out of decryption and end the session without going to the site or to proceed to the site and agree to have the traffic decrypted), educate them about what it is, why they're seeing it, and what their options are.

- Create realistic deployment schedules that allow time to evaluate each stage of the rollout.



Place firewalls in positions where they can see all of the network traffic so that no encrypted traffic inadvertently gains access to your network because it bypasses the firewall.

Define Traffic to Decrypt

A Decryption policy rule allows you to define traffic that you want the firewall to decrypt and to define traffic that you choose to **exclude** from decryption because the traffic is personal or because of local regulations, for example.

Attach a Decryption profile to each Decryption policy rule to enable certificate checks, session mode checks, failure checks, and protocol and algorithm checks, depending on the profile. These checks prevent risky connections, such as sessions with untrusted certificate issuers, weak protocols, ciphers, and algorithms, and servers that have certificate issues.



Review [the Decryption deployment best practices checklist](#) to ensure that you understand the recommended best practices.

As a best practice, you should always block known dangerous [URL Filtering categories](#) such as malware, phishing, dynamic-dns, unknown, command-and-control, proxy-avoidance-and-anonymizers, copyright-infringement, extremism, newly-registered-domain, grayware, and parked. If you must allow any of these categories for business reasons, you must decrypt them and apply strict Security profiles to the traffic.

URL categories that you should always decrypt if you allow them include: online-storage-and-backup, web-based-email, web-hosting, personal-sites-and-blogs, and content-delivery-networks.



In Security policy, block Quick UDP Internet Connections (QUIC) protocol unless for business reasons, you want to allow encrypted browser traffic. Chrome and some other browsers establish sessions using QUIC instead of TLS, but QUIC uses proprietary encryption that the firewall can't decrypt, so potentially dangerous traffic may enter the network as encrypted traffic. Blocking QUIC forces the browser to fall back to TLS and enables the firewall to decrypt the traffic.

Create a Security policy rule to block QUIC on its UDP service ports (80 and 443) and create a separate rule to block the QUIC application. For the rule that blocks UDP ports 80 and 443, create a Service (Objects > Services) that includes UDP ports 80 and 443:



Use the Service to specify the UDP ports to block for QUIC. In the second rule, block the QUIC application:



- [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 SSL traffic that you *choose* to **exclude** from decryption. (If a server breaks SSL decryption technically due to certificate pinning or other

reasons, add the server to the Decryption [Exclusion](#) list.) Depending on your needs, create Decryption profiles to:

- Block sessions based on certificate status, including blocking sessions with expired certificates, untrusted issuers, unknown certificate status, certificate status check timeouts, and certificate extensions.
- Block sessions with unsupported versions and cipher suites, and that require using client authentication.
- Block sessions if the resources to perform decryption are not available or if a hardware security module is not available to sign certificates.
- Define the protocol versions and key exchange, encryption, and authentication algorithms allowed for SSL Forward Proxy and SSL Inbound Inspection traffic in the SSL Protocol Settings.

Don't weaken the main Decryption profile that you apply to most sites to accommodate weaker sites. Instead, create one or more separate Decryption profiles for sites that you need to support but that don't support strong ciphers and algorithms. You can also create different Decryption profiles for different URL Categories to fine tune security vs. performance for traffic that contains no sensitive material; however, you should always decrypt and inspect all the traffic you can.

After you create a decryption profile, attach it to a decryption policy rule; the firewall then enforces the decryption profile settings on traffic that matches 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. However, the best practice is to enable tighter decryption controls as described in [SSL Forward Proxy Decryption Profile](#), [SSL Inbound Inspection Decryption Profile](#), and [SSL Protocol Settings Decryption Profile](#).



Avoid supporting weak protocols or algorithms because they contain known vulnerabilities that attackers can exploit. If you must allow a weaker protocol or algorithm to support a key partner or contractor who uses legacy systems with weak protocols, create a separate Decryption profile for the exception and attach it to a Decryption policy rule that applies the profile only to the relevant traffic (for example, the source IP address of the partner). Don't allow the weak protocol for all traffic.

STEP 1 | Create a new decryption profile.

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) Enable an Ethernet Interface for the firewall to use to copy and forward decrypted traffic.

Separate from this task, follow the steps to [Configure Decryption Port Mirroring](#). Be aware of local privacy regulations that may prohibit mirroring or control the type of traffic that you can mirror. Decryption port mirroring requires a decryption port mirror license.

STEP 4 | (Optional) Block and control SSL tunneled and/or inbound traffic:



Although applying a Decryption profile to decrypted traffic is optional, it is a best practice to always apply a Decryption profile to the policy rules to protect your network against encrypted threats. You can't protect yourself against threats you can't see.

Select **SSL Decryption**:

- Select **SSL Forward Proxy** to configure the 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 configured to perform SSL Forward Proxy decryption.
- Select **SSL Inbound Inspection** to configure the settings to 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 performs SSL Inbound Inspection.
- Select **SSL Protocol Settings** to configure the settings that control minimum and maximum protocol versions and key exchange, encryption, and authentication algorithms to enforce on decrypted 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 choose to [Create a Policy-Based Decryption Exclusion](#).



Although applying a Decryption profile to traffic that you choose not to decrypt is optional, it is a best practice to always apply a Decryption profile to the policy rules to protect your network against sessions with expired certificates or untrusted issuers.

Select **No Decryption** to configure the [Decryption Profile for No Decryption](#) and check the **Block sessions with expired certificates** and **Block sessions with untrusted issuers** boxes to validate certificates for traffic that is excluded from decryption. Create policy-based exclusions only for traffic that you choose not to decrypt. If a server breaks decryption for technical reasons, don't create a policy-based exclusion, add the server to the SSL Decryption Exclusion list (**Device > Certificate Management > SSL Decryption Exclusion**).

These settings 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 decrypted SSH traffic.

Select **SSH Proxy** to configure the [SSH Proxy Decryption Profile](#) and configure settings to enforce supported protocol versions and to block sessions if system resources are not available to perform decryption.

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 when you [Create a Decryption Policy Rule](#).

The firewall applies the decryption profile to and enforces the profile's settings on the traffic that matches the decryption policy rule.

STEP 8 | **Commit** the configuration.

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 Mirroring](#).

STEP 1 | Add a new decryption policy rule.

Select **Policies > Decryption**, **Add** a new decryption policy rule, and give the policy rule a descriptive **Name**.

STEP 2 | 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 can be useful when you [Create a Policy-Based Decryption Exclusion](#). 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**).
 - Palo Alto Networks predefined [URL categories](#), which make it easy to decrypt entire categories of allowed traffic. This option is also useful when you create policy-based decryption exclusions because you can exclude sensitive sites by category instead of individually. For example, although you can create a custom URL category to group sites that you do not want to decrypt, you can also exclude financial or healthcare-related sites from decryption based on the predefined Palo Alto Networks URL categories. In addition, you can block risky URL Categories and create comfort pages to communicate the reason the sites are blocked or [Enable Users to Opt Out of SSL Decryption](#).

You can use the predefined high-risk and medium-risk URL categories to create a Decryption policy rule that decrypts all high-risk and medium-risk URL traffic. Place the rule at the bottom of the rulebase (all decryption exceptions must be above this rule so that you don't decrypt sensitive information) as a safety net to ensure that you decrypt and inspect all risky traffic. However, if high-risk or medium-risk sites to which you allow access contain personally identifiable information (PII) or other sensitive information that you don't want to decrypt, either block those sites to avoid allowing encrypted risky traffic while also avoiding privacy issues, or create a No Decryption rule to handle the sensitive traffic.

- Custom URL categories (see **Objects > Custom Objects > URL Category**). For example, you can create a custom URL Category to specify a group of sites you need to access for business purposes but that don't support the safest protocols and algorithms, and then apply a customized Decryption profile to allow the looser protocols and algorithms for just those sites (that way, you don't decrease security by downgrading the Decryption profile you use for most sites).

STEP 3 | Set the rule to either decrypt matching traffic or to exclude matching traffic from decryption.

Select **Options** and set the policy rule **Action**:

To decrypt matching traffic:

1. Set the **Action** to **Decrypt**.
2. Set the **Type** of decryption for the firewall to perform on matching traffic:
 - [SSL Forward 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.
- [SSH Proxy](#)

To exclude matching traffic from decryption:

Set the **Action** to **No Decrypt**.

STEP 4 | (Optional) Select a **Decryption Profile** to perform additional checks on traffic that matches the policy rule.



Although applying a Decryption profile to decrypted traffic is optional, it is a best practice to always apply a Decryption profile to the policy rules to protect your network against encrypted threats. You can't protect yourself against threats you can't see.

For example, attach a decryption profile to a policy rule to ensure that server certificates are valid and to block sessions using unsupported protocols or ciphers. To [Create a Decryption Profile](#), select **Objects > Decryption Profile**.

1. Create a decryption policy rule or open an existing rule to modify it.
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 the firewall applies to matching traffic depends 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 different decryption profiles with different types of decryption policy rules that apply to different types of traffic and users.

3. Click **OK**.

STEP 5 | Click **OK** to save the policy.

STEP 6 | Choose your next step to fully enable the firewall to decrypt traffic...

- [ConfigureSSL Forward Proxy](#)
- [ConfigureSSL Inbound Inspection](#)
- [ConfigureSSH Proxy](#)
- Create policy-based [Decryption Exclusions](#) for traffic you *choose* not to decrypt and add sites that break decryption for technical reasons such as pinned certificates or mutual authentication to the SSL Decryption Exclusion list.

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 (proxy) to the session between the client and the server. The firewall can use certificates signed by an enterprise certificate authority (CA) or self-signed certificates generated on the firewall as *Forward Trust certificates* to authenticate the SSL session with the client.

- **(Recommended Best Practice) Enterprise CA-signed Certificates**—An enterprise CA can issue a signing certificate that the firewall can use to sign the certificates for sites which require SSL decryption. When the firewall trusts the CA that signed the certificate of the destination server, the firewall can send a copy of the destination server certificate to the client, signed by the enterprise CA. This is a best practice because usually all network devices already trust the Enterprise CA (it is usually already installed in the devices' CA Trust storage), so you don't need to deploy the certificate on the endpoints, so the rollout process is smoother.
- **Self-signed Certificates**—The firewall can act as a CA and generate self-signed certificates that the firewall can use to sign the certificates for sites which require SSL decryption. The firewall can sign a copy of the server certificate to present to the client and establish the SSL session. This method requires that you need to install the self-signed certificates on all of your network devices so that those devices recognize the firewall's self-signed certificates. Because the certificates must be deployed to all devices, this method is better for small deployments and proof-of-concept (POC) trials than for large deployments.

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.



Regardless of whether you generate Forward Trust certificates from your Enterprise Root CA or use a self-signed certificate generated on the firewall, generate a separate subordinate Forward Trust CA certificate for each firewall. The flexibility of using separate subordinate CAs enables you to [revoke](#) one certificate when you decommission a device (or device pair) without affecting the rest of the deployment and reduces the impact in any situation in which you need to revoke a certificate. Separate Forward Trust CAs on each firewall also helps troubleshoot issues because the CA error message the user sees includes information about the firewall the traffic is traversing. If you use the same Forward Trust CA on every firewall, you lose the granularity of that information.

After setting up the Forward Trust and Forward Untrust certificates required for SSL Forward Proxy decryption, create a Decryption policy rule to define the traffic you want the firewall to decrypt and create a Decryption profile to apply SSL controls and checks to the traffic. The Decryption policy decrypts SSL tunneled traffic that matches the rule into clear text traffic. The firewall blocks and restricts traffic based on the Decryption profile attached to the Decryption policy and on the firewall Security policy. The firewall re-encrypts traffic as it exits the firewall.

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 a trusted CA has signed the server certificate. You can use an enterprise CA-signed certificate or a self-signed certificate as the forward trust certificate.

(**Recommended Best Practice**) Use an enterprise CA-signed certificate as the Forward Trust certificate. Create a uniquely named Forward Trust certificate on each firewall:

1. Generate a Certificate Signing Request (CSR) for the enterprise CA to sign and validate:

1. Select **Device > Certificate Management > Certificates** and click **Generate**.
2. Enter a **Certificate Name**. Use a unique name for each firewall.
3. In the **Signed By** drop-down, select **External Authority (CSR)**.
4. (**Optional**) If your enterprise CA requires it, add **Certificate Attributes** to further identify the firewall details, such as Country or Department.
5. Click **Generate** to save the CSR. The pending certificate is now displayed on the **Device Certificates** tab.

2. Export the CSR:

1. Select the pending certificate displayed on the **Device Certificates** tab.
2. 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.

3. 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 to import onto the firewall.
4. Import the enterprise CA-signed certificate onto the firewall:
 1. Select **Device > Certificate Management > Certificates** and click **Import**.
 2. Enter the pending **Certificate Name** exactly. The **Certificate Name** that you enter must exactly match the pending certificate name in order for the pending certificate to be validated.
 3. Select the signed **Certificate File** that you received from your enterprise CA.
 4. Click **OK**. The certificate is displayed as valid with the Key and CA check boxes selected.
5. Select the validated certificate 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.

Use a self-signed certificate as the Forward Trust certificate:

1. Create a [self-signed Root CA certificate](#).
2. Click the self-signed root CA certificate (**Device > Certificate Management > Certificates > Device Certificates**) to open **Certificate information** and then click the **Trusted Root CA** checkbox.
3. Click **OK**.
4. Generate new subordinate CA certificates for each firewall:
 1. Select **Device > Certificate Management > Certificates**.
 2. Click **Generate** at the bottom of the window.
 3. Enter a **Certificate Name**.
 4. Enter a **Common Name**, such as 192.168.2.1. This should be the IP or FQDN that will appear in the certificate. In this case, we are using the IP of the trust interface. Avoid using spaces in this field.
 5. In the **Signed By** field, select the self-signed Root CA certificate that you created.
 6. Click the **Certificate Authority** check box to enable the firewall to issue the certificate. Selecting this check box creates a certificate authority (CA) on the firewall that is imported to the client browsers, so clients trust the firewall as a CA.

7. **Generate** the certificate.
5. Click the new certificate to modify it and click the **Forward Trust Certificate** checkbox to configure the certificate as the Forward Trust Certificate.
6. Click **OK** to save the self-signed forward trust certificate.
7. Repeat this procedure to generate a unique subordinate CA certificate on each firewall.

STEP 3 | Distribute the forward trust certificate to client system certificate stores.

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 installed in the local trusted root CA list, you can skip this step. (The client systems trust the subordinate CA certificates you generate on the firewall because the Enterprise Trusted Root CA has signed them.)



If you do not install the forward trust certificate on client systems, users see certificate warnings for each SSL site they visit.

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 self-signed firewall Trusted Root CA certificate to the Trusted Root CA section. After GlobalProtect distributes the firewall's Trusted Root CA certificate to client systems, the client systems trust the firewall's subordinate CA certificates because the clients trust the firewall's Root CA certificate.
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 firewall Trusted Root CA certificate so that you can import it into client systems. Highlight the certificate and click **Export** at the bottom of the window. Choose PEM format.



Do not select the Export private key checkbox! The private key should remain on the firewall and should not be exported to client systems.

Import the firewall's Trusted Root CA certificate into the browser Trusted Root CA list on the client systems in order for the clients to trust it. When importing into the client browser, ensure that you add the certificate 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 option, such as an Active Directory Group Policy Object (GPO).

STEP 4 | Configure the Forward Untrust certificate (use the same Forward Untrust certificate for all firewalls).

1. Click **Generate** at the bottom of the certificates page.
2. Enter a **Certificate Name**, such as my-ssl-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.

-
- Click the new my-ssl-fwd-untrust certificate to modify it and enable the **Forward Untrust Certificate** option.



Do not export the Forward Untrust certificate to the Certificate Trust Lists of your network devices! Do not install the Forward Untrust certificate on client systems. This is critical because installing the Untrust certificate in the Trust List results in devices trusting websites that the firewall does not trust. In addition, users won't see certificate warnings for untrusted sites, so they won't know the sites are untrusted and may access those sites, which could expose your network to threats.

- Click **OK** to save.

STEP 5 | (Optional) **Configure the Key Size for SSL Forward Proxy Server 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.

STEP 6 | **Create a Decryption Policy Rule** to define traffic for the firewall to decrypt and **Create a Decryption Profile** to apply SSL controls to the traffic.



Although Decryption profiles are optional, it is a best practice to include a Decryption profile with each Decryption policy rule to prevent weak, vulnerable protocols and algorithms from allowing questionable traffic on your network.

- Select **Policies > Decryption**, Add or modify an existing rule, and define traffic to be decrypted.
- Select **Options** and:
 - Set the rule **Action** to **Decrypt** matching traffic.
 - Set the rule **Type** to **SSL Forward Proxy**.
 - (Optional but a best practice) Configure or select an existing **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).
- 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 Exclusions** to disable decryption for certain types of traffic.

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 load the server certificate onto the firewall). With an SSL Inbound Inspection Decryption policy enabled, the firewall decrypts all SSL traffic identified by the policy to clear text traffic and inspects it. The firewall blocks, restricts, or allows the traffic based on the Decryption profile attached to the policy and the Security policy that applies to the traffic, including and any configured Antivirus, Vulnerability Protection, Anti-Spyware, URL-Filtering, and File Blocking profiles. As a best practice, enable the firewall to [forward decrypted SSL traffic for WildFire analysis](#) and signature generation.

Configuring [SSL Inbound Inspection](#) includes installing the targeted server certificate on the firewall, creating an SSL Inbound Inspection Decryption policy, and applying a Decryption profile to the policy.

STEP 1 | Ensure that the appropriate interfaces are configured as either Tap, Virtual Wire, Layer 2, or Layer 3 interfaces.



You cannot use a Tap mode interface for SSL inbound inspection if the negotiated cyphers include PFS key-exchange algorithms (DHE and ECDHE).

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 the interface type.

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 and [Create a Decryption Profile](#) to apply SSL controls to the traffic.



Although Decryption profiles are optional, it is a best practice to include a Decryption profile with each Decryption policy rule to prevent weak, vulnerable protocols and algorithms from allowing questionable traffic on your network.

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 but a best practice**) Configure or select an existing **Decryption Profile** to block and control various aspects of the decrypted traffic (for example, create a Decryption profile to terminate sessions with unsupported algorithms and unsupported cipher suites).



When you configure the [SSL Protocol Settings Decryption Profile](#) for [SSL Inbound Inspection traffic](#), create separate profiles for servers with different security capabilities. For example, if one set of servers supports only RSA, the [SSL Protocol Settings](#) only need to support RSA. However, the [SSL Protocol Settings](#) for servers that support PFS should support PFS. Configure [SSL Protocol Settings](#) for the highest level of security that the server supports, but check performance to ensure that the firewall resources can handle the higher processing load that higher security protocols and algorithms require.

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.

STEP 6 | Choose your next step...

- [Enable Users to Opt Out of SSL Decryption](#).
- Configure [Decryption Exclusions](#) to disable decryption for certain types of traffic.

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, the firewall decrypts SSH traffic and blocks and or restricts the SSH traffic based on your decryption policy and decryption profile settings. Traffic is re-encrypted as it exits the firewall.

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 and [Create a Decryption Profile](#) to apply checks to the SSH traffic.



Although Decryption profiles are optional, it is a best practice to include a Decryption profile with each Decryption policy rule to prevent weak, vulnerable protocols and algorithms from allowing questionable traffic on your network.

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 but a best practice**) Configure or select an existing **Decryption Profile** to block and control various aspects of the decrypted traffic (for example, create a Decryption profile to terminate sessions with unsupported versions and unsupported algorithms).
3. Click **OK** to save.

STEP 3 | **Commit** the configuration.

STEP 4 | (**Optional**) Continue to [Decryption Exclusions](#) to disable decryption for certain types of traffic.

Configure Server Certificate Verification for Undecrypted Traffic

You create no-decryption policies for traffic that you *choose* not to decrypt because the traffic is personal, sensitive, or subject to local laws and regulations. For example, you may choose not to decrypt the traffic of certain executives or traffic between finance users and finance servers that contain personal information. (Don't exclude traffic that you can't decrypt because a site breaks decryption for technical reasons such as a pinned certificate or mutual authentication by policy. Instead, add the hostname to the [Decryption Exclusion List](#).)

However, just because you don't decrypt the traffic doesn't mean you should let any and all undecrypted traffic on your network. It is a best practice to apply a No Decryption profile to undecrypted traffic to block sessions with expired certificates and untrusted issuers.

STEP 1 | [Create a Decryption Policy Rule](#) to identify the undecrypted traffic and [Create a Decryption Profile](#) to block bad sessions.

1. Select **Policies > Decryption** and Add or modify an existing rule to identify the undecrypted traffic.
2. Select **Options** and:
 - Set the rule **Action** to **No Decrypt** so that the firewall doesn't decrypt traffic that matches the rule.
 - Ignore the rule **Type** because the traffic is not decrypted.
 - (**Optional but a best practice**) Configure or select an existing [Decryption profile for undecrypted traffic](#) to block sessions with expired certificates and untrusted certificate issuers.

STEP 2 | **Commit** the configuration.

STEP 3 | Choose your next step:

- [Enable Users to Opt Out of SSL Decryption](#).
- Configure [Decryption Exclusions](#) to disable decryption for certain types of traffic.

Decryption Exclusions

You can exclude two types of traffic from decryption:

- Traffic that breaks decryption for *technical reasons*, such as using a pinned certificate, an incomplete certificate chain, unsupported ciphers, or mutual authentication (decrypting blocks the traffic). Palo Alto Networks provides a predefined SSL Decryption Exclusion list (**Device > Certificate management > SSL Decryption Exclusion**) that excludes hosts with applications and services that are known to break decryption technically from SSL Decryption by default. If you encounter sites that break decryption technically and are not on the SSL Decryption Exclusion list, you can add them to list manually by server hostname. The firewall blocks sites whose applications and services break decryption technically unless you add them to the SSL Decryption Exclusion list.
- Traffic that you *choose* not to decrypt because of business, regulatory, personal, or other reasons, such as financial-services, health-and-medicine, or government traffic. You can choose to exclude traffic based on source, destination, URL category, and service.

You can use asterisks (*) as wildcards to create decryption exclusions for multiple hostnames associated with a domain. Asterisks behave the same way that carets (^) behave for URL category exceptions—each asterisk controls one variable subdomain (label) in the hostname. This enables you to create both very specific and very general exclusions. For example:

- mail.*.com matches mail.company.com but does not match mail.company.sso.com.
- *.company.com matches tools.company.com but does not match eng.tools.company.com.
- *.*.company.com matches eng.tools.company.com but does not match eng.company.com.
- *.*.*.company.com matches corp.exec.mail.company.com, but does not match corp.mail.company.com.
- mail.google.* matches mail.google.com, but does not match mail.google.uk.com.
- mail.google.*.* matches mail.google.co.uk, but does not match mail.google.com.

For example, to use wildcards to exclude video-stats.video.google.com from decryption but not to exclude video.google.com from decryption, exclude *.*.google.com.



*Regardless of the number of asterisk wildcards that precede a hostname (without a non-wildcard label preceding the hostname), the hostname matches the entry. For example, *.google.com, *.*.google.com, and *.*.*.google.com all match google.com. However, *.dev.*.google.com does not match google.com because one label (dev) is not a wildcard.*

To increase visibility into traffic and reduce the attack surface as much as possible, don't make decryption exceptions unless you must.

- [Palo Alto Networks Predefined Decryption Exclusions](#)
- [Exclude a Server from Decryption for Technical Reasons](#)
- [Create a Policy-Based Decryption Exclusion](#)

Palo Alto Networks Predefined Decryption Exclusions


The firewall provides a predefined SSL Decryption Exclusion list to exclude from decryption commonly used sites that break decryption because of technical reasons such as pinned certificates and mutual authentication. The predefined decryption exclusions are enabled by default and Palo Alto Networks delivers new and updated predefined decryption exclusions to the firewall as part of the Applications and Threats content update (or the Applications content update, if you do not have a Threat Prevention license). The firewall does not decrypt traffic that matches predefined exclusions and allows the encrypted traffic based on the Security policy that governs that traffic. However, the firewall can't inspect the encrypted traffic or enforce Security policy on it.



The SSL Decryption Exclusion list is **not** for sites that you choose not to decrypt for legal, regulatory, business, privacy, or other volitional reasons, it is only for sites that break decryption technically (decrypting these sites blocks their traffic). For traffic such as IP addresses, users, URL categories, services, and even entire zones that you choose not to decrypt, [Create a Policy-Based Decryption Exclusion](#).

Because the traffic of sites on the SSL Decryption Exclusion list remains encrypted, the firewall does not inspect or provide further security enforcement the traffic. You can disable a predefined exclusion. For example, you may choose to disable predefined exclusions to enforce a strict security policy that allows only applications and services that the firewall can inspect and on which the firewall can enforce Security policy. However, the firewall blocks sites whose applications and services break decryption technically if they are not enabled on the SSL Decryption Exclusion list.

You can view and manage all Palo Alto Networks predefined SSL decryption exclusions directly on the firewall (**Device > Certificate Management > SSL Decryption Exclusions**).



The **Hostname** displays the name of the host that houses the application or service that breaks decryption technically. You can also **Add** hosts to [Exclude a Server from Decryption for Technical Reasons](#) if it is not on the predefined list.

The **Description** displays the reason the firewall can't decrypt the site's traffic, for example, **pinned-cert** (a pinned certificate) or **client-cert-auth** (client authentication).

The firewall automatically removes enabled predefined SSL decryption exclusions from the list when they become obsolete (the firewall removes an application that decryption previously caused to break when the application becomes supported with decryption). **Show Obsoletes** checks if any disabled predefined exclusions remain on the list and are no longer needed. The firewall does not remove disabled predefined decryption exclusions from the list automatically, but you can select and **Delete** obsolete entries.

You can select a hostname's checkbox and then click **Disable** to remove predefined sites from the list. Use the SSL Decryption Exclusion list only for sites that break decryption for technical reasons, don't use it for sites that you choose not to decrypt.

Exclude a Server from Decryption for Technical Reasons

If decryption breaks an important application or service technically (decrypting the traffic blocks it), you can add the hostname of the site that hosts to the application or service to the Palo Alto Networks predefined SSL Decryption Exclusion list to create a custom decryption exception. The firewall doesn't decrypt, inspect, and enforce Security policy on traffic that the SSL Decryption Exclusion list allows because the traffic remains encrypted, so be sure that the sites you add to the list really are sites with applications or

services you need for business. For example, some business-critical internal custom applications may break decryption and you can add them to the list so that the firewall allows the encrypted custom application traffic.



*The SSL Decryption Exclusion list is **not** for sites that you choose not to decrypt for legal, regulatory, business, privacy, or other volitional reasons, it is only for sites that break decryption technically. For traffic (IP addresses, users, URL categories, services, and even entire zones) that you choose not to decrypt, [Create a Policy-Based Decryption Exclusion](#).*

Reasons that sites break decryption technically include pinned certificates, mutual authentication, incomplete certificate chains, and unsupported ciphers. For HTTP public key pinning (HPKP), most browsers that use HPKP permit Forward Proxy decryption as long as you install the enterprise CA certificate (or the certificate chain) on the client.



If the technical reason for excluding a site from decryption is an incomplete certificate chain, the next-generation firewall doesn't automatically fix the chain as a browser would. If you need to add a site to the SSL Decryption Exclusion list, manually review the site to ensure it's a legitimate business site, then download the missing sub-CA certificates and [load and deploy](#) them onto the firewall.

After you add a server to the SSL Decryption Exclusion list, the firewall compares the server hostname that you use to define the decryption exclusion against the common name (CN) in the certificate a server presents. If a single server hosts multiple websites using different certificates, the firewall compares the hostname against the server name indication (SNI) that the client presents to indicate the server to which it wants to connect.

STEP 1 | Select **Device > Certificate Management > SSL Decryption Exclusions**.

STEP 2 | **Add** a new decryption exclusion, or select an existing custom entry to modify it.

STEP 3 | Enter the **hostname** of the website or application you want to exclude from decryption.



The hostname is case-sensitive.

You can [use wildcards](#) to exclude multiple hostnames associated with a domain. The firewall excludes all sessions where the server presents a CN that matches the domain from decryption.

Make sure that the hostname field is unique for each custom entry. If a predefined exclusion matches a custom entry, the custom entry takes precedence.

STEP 4 | (**Optional**) Select **Shared** to share the exclusion across all virtual systems in a multiple virtual system firewall.

STEP 5 | **Exclude** the application from decryption. Alternatively, if you are modifying an existing decryption exclusion, you can clear this checkbox to start decrypting an entry that was previously excluded from decryption.

STEP 6 | Click **OK** to save the new exclusion entry.

Create a Policy-Based Decryption Exclusion

Policy-based decryption exclusions are for excluding traffic that you *choose* not to decrypt. You can create a policy-based decryption exclusion based on any combination of the traffic's source, destination, service, or URL Category. Examples of traffic you may choose not to decrypt include:

- Traffic that you should never decrypt because it contains personally identifiable information (PII) or other sensitive information, such as the [URL Filtering categories](#) financial-services, health-and-medicine, and government.
- Traffic that originates or is destined for executives or other users whose traffic shouldn't be decrypted.
- Some devices such as finance servers may need to be excepted from decryption.
- Depending on the business, some companies may value privacy and the user experience more than security for some applications.
- Laws or local regulations that prohibit decryption of some traffic.

An example of not decrypting traffic for regulatory and legal compliance is the European Union (EU) General Data Protection Regulation (GDPR). The EU GDPR will require strong protection of all personal data for all individuals. The GDPR affects all companies, including foreign companies, that collect or process the personal data of EU residents.

Different regulations and compliance rules may mean that you treat the same data differently in different countries or regions. Businesses usually can decrypt personal information in their corporate data centers because the business owns the information. The best practice is to decrypt as much traffic as possible so that you can see it and apply security protection to it.

You can use predefined URL Categories to except entire categories of websites from decryption, you can create custom URL Categories to define a customized list of URLs that you don't want to decrypt, or you can create an [External Dynamic List](#) (EDL) to define a customized list of URLs that you don't want to decrypt.

In environments such as Office 365 that have dynamically changing IP addresses or in environments where you make frequent changes to the list of URLs that you want to exclude from decryption, it's often preferable to use an EDL instead of a URL Category to specify the excluded URLs. Using an EDL is less disruptive in dynamic environments because editing an EDL changes the URL categories dynamically, without a **Commit**, while editing a custom URL Category requires a **Commit** to take effect.



Create an EDL or a custom URL Category that contains all the categories you choose not to decrypt so that one Decryption policy rule governs the encrypted traffic you choose to allow. Apply a No Decryption profile to the rule. The ability to add categories to an EDL or a custom URL Category makes it easy to exclude traffic from decryption and helps keep the rulebase clean.



Similar to Security policy rules, the firewall compares incoming traffic to Decryption policy rules in the policy rulebase's sequence. Place Decryption exclusion rules at the top of the rulebase to prevent inadvertently decrypting sensitive traffic or traffic that laws and regulations prevent you from decrypting.

If you create policy-based decryption exclusions, the best practice is to place the following exclusion rules at the top of the decryption rulebase, in the following order:

1. IP-address based exceptions for sensitive destination servers.
2. Source-user based exceptions for executives and other users or groups.
3. Custom URL or EDL based exceptions for destination URLs.

4. Sensitive predefined URL Category based exceptions for destination URLs of entire categories such as financial-services, health-and-medicine, and government.

Place rules that decrypt traffic after these rules in the decryption rulebase.


STEP 1 | Exclude traffic from decryption based on 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 **Add** or modify a decryption policy rule.
2. Define the traffic that you want to exclude from decryption.

In this example:

1. Give the rule a descriptive **Name**, such as No-Decrypt-Finance-Health.
2. Set the **Source** and **Destination** to **Any** to apply the No-Decrypt-Finance-Health rule to all SSL traffic destined for an external server.
3. 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 but a best practice) Create a [policy-based decryption exclusion](#) to validate certificates for sessions the firewall does not decrypt. Attach the No Decryption profile to the rule and set the profile to **Block sessions with expired certificates** and **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 rulebase.

The firewall enforces decryption rules against incoming traffic in the rulebase sequence and enforces the first rule that match the traffic.

Select the **No-Decrypt-Finance-Health** policy (**Decryption > Policies**), and click **Move Up** until it appears at the top of the list, or drag and drop the rule.

STEP 3 | Save the configuration.

Click **Commit**.

Enable Users to Opt Out of SSL Decryption

In privacy-sensitive situations, you may want to alert your users that the firewall is decrypting certain web traffic and allow them either to continue to the site with the understanding that their traffic is decrypted or to terminate the session and be block from going to the site. (There is no option to go to the site and also avoid decryption.)

The first time a user attempts to browse to an HTTPS site or application that matches the decryption policy, the firewall displays a response page notifying users 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 redisplay 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 redisplay 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. However, the best practice is to not allow users to opt out of decryption.



Custom response pages larger than the maximum supported size are not decrypted or displayed to users. In PAN-OS 8.1.2 and earlier PAN-OS 8.1 releases, custom response pages on a decrypted site cannot exceed 8,191 bytes; the maximum size is increased to 17,999 bytes in PAN-OS 8.1.3 and later releases.

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:

```

```

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.

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.



Temporarily Disable SSL Decryption

In some cases you may want to temporarily disable SSL decryption. For example, if you deployed SSL decryption too hastily and something doesn't work correctly but you're not sure what it is, and you have a lot of rules to examine, you can use the CLI to temporarily turn off decryption and give yourself time to analyze and solve the issue. After solving the issue, you can use the CLI to turn SSL decryption back on again. Because temporarily disabling and then re-enabling decryption using the CLI doesn't require a Commit operation, you can do it without disrupting network traffic.

The following CLI commands temporarily disable SSL decryption without a Commit and re-enable decryption without a Commit.



The command to disable SSL decryption doesn't persist in the configuration after a reboot. If you turn off decryption temporarily and then reboot the firewall, regardless of whether the issue has been fixed, decryption is turned on again.

- 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
```


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.

Keep in mind that the decryption, storage, inspection, and/or use of SSL traffic is regulated in certain countries and user consent may 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.

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 website](#) 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 and still want to set up decryption port mirroring, click **I understand and wish to proceed**.
5. Click **Activate**.

DEVICE LICENSES

Serial Number: 009C100103
Model: PAN-PA-5050-B
Device Name: PM Lab Firewall

Authorization Code: **Add**

Feature Name	Authorization Code	Expiration Date	Actions
Threat Prevention	14344239	01/06/2019	
PAN-DB URL Filtering	19544847	01/06/2019	
Virtual Systems	18729162	Perpetual	
Premium Support	17480971	12/29/2015	

AVAILABLE FEATURE LICENSES

☐ Decryption Port Mirror

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.

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**.

Verify Decryption

After you configure a best practice decryption profile and apply it to traffic, check the log files to verify that the firewall is decrypting the traffic that you intend to decrypt and that the firewall is not decrypting the traffic that you don't want to decrypt. In addition, [follow post-deployment decryption best practices](#) to maintain the deployment.


- **View Decrypted Traffic Sessions**—Filter the Traffic Logs (**Monitor > Logs > Traffic**) using the filter (**flags has proxy**).

This filter displays only logs in which the SSL proxy flag is on, meaning only decrypted traffic—every log entry has the value **yes** in the **Decrypted** column.



	Receive Time	Type	From Zone	To Zone	Source	Session ID	Destination	To Port	Decrypted	Application	Action	Rule
	07/12 14:08:38	deny	L3-Trust	L3-UnTrust	192.168.102.34	45484	54.148.39.150	443	yes	ssl	allow	Test
	07/12 14:08:38	deny	L3-Trust	L3-UnTrust	192.168.102.34	45482	54.148.39.150	443	yes	ssl	allow	Test
	07/12 14:08:38	deny	L3-Trust	L3-UnTrust	192.168.102.34	45481	54.148.39.150	443	yes	ssl	allow	Test
	07/12 14:08:38	deny	L3-Trust	L3-UnTrust	192.168.102.34	45479	54.148.39.150	443	yes	ssl	allow	Test
	07/12 14:08:38	deny	L3-Trust	L3-UnTrust	192.168.102.34	45480	54.148.39.150	443	yes	ssl	allow	Test
	07/12 14:08:38	deny	L3-Trust	L3-UnTrust	192.168.102.34	45483	54.148.39.150	443	yes	ssl	allow	Test
	07/12 14:08:38	deny	L3-Trust	L3-UnTrust	192.168.102.34	45478	54.148.39.150	443	yes	ssl	allow	Test
	07/12 13:36:37	deny	L3-Trust	L3-UnTrust	192.168.102.34	45456	54.186.39.123	443	yes	ssl	allow	Test
	07/12 13:36:37	deny	L3-Trust	L3-UnTrust	192.168.102.34	45452	54.186.39.123	443	yes	ssl	allow	Test
	07/12 13:36:37	deny	L3-Trust	L3-UnTrust	192.168.102.34	45454	54.186.39.123	443	yes	ssl	allow	Test
	07/12 13:36:37	deny	L3-Trust	L3-UnTrust	192.168.102.34	45451	54.186.39.123	443	yes	ssl	allow	Test

You can filter the traffic in a more granular fashion by adding more terms to the filter. For example, you can filter for decrypted traffic going only to the destination IP address 172.217.3.206 by adding the filter (**addr.dst in 172.217.3.206**):



	Receive Time	Type	From Zone	To Zone	Source	Session ID	Destination	To Port	Decrypted	Application	Action	Rule
	07/05 18:06:36	deny	L3-Trust	L3-UnTrust	192.168.102.34	41276	172.217.3.206	443	yes	ssl	allow	Test
	07/05 00:06:25	deny	L3-Trust	L3-UnTrust	192.168.102.34	41268	172.217.3.206	443	yes	ssl	allow	Test
	07/05 00:06:09	deny	L3-Trust	L3-UnTrust	192.168.102.34	41355	172.217.3.206	443	yes	ssl	allow	Test
	07/05 17:39:25	end	L3-Trust	L3-UnTrust	192.168.102.34	22995	172.217.3.206	443	yes	ssl	allow	Test
	07/05 17:39:25	end	L3-Trust	L3-UnTrust	192.168.102.34	22997	172.217.3.206	443	yes	ssl	allow	Test
	07/05 17:39:25	end	L3-Trust	L3-UnTrust	192.168.102.34	22998	172.217.3.206	443	yes	ssl	allow	Test
	07/05 13:07:05	end	L3-Trust	L3-UnTrust	192.168.102.34	4272	172.217.3.206	443	yes	ssl	allow	Test
	07/05 13:06:23	end	L3-Trust	L3-UnTrust	192.168.102.34	4266	172.217.3.206	443	yes	ssl	allow	Test
	07/05 13:06:22	end	L3-Trust	L3-UnTrust	192.168.102.34	4259	172.217.3.206	443	yes	ssl	allow	Test
	07/05 13:05:45	end	L3-Trust	L3-UnTrust	192.168.102.34	4221	172.217.3.206	443	yes	ssl	allow	Test

- **View SSL Traffic Sessions That Are Not Decrypted**—Filter the Traffic Logs (**Monitor > Logs > Traffic**) using the filter (**not flags has proxy**) and (**app eqssl**).

This filter displays only logs in which the SSL proxy flag is off (meaning only encrypted traffic) and the traffic is SSL traffic; every log entry has the value **no** in the **Decrypted** column and the value **ssl** in the **Application** column.

(not flags has proxy) and japp eq ssl

Receive Time	Type	From Zone	To Zone	Source	Session ID	Destination	To Port	Decrypted	Application	Action
07/25 16:50:02	deny	L3-Trust	L3-UnTrust	192.168.102.34	57908	216.58.217.206	443	no	ssl	allow
07/25 16:50:02	deny	L3-Trust	L3-UnTrust	192.168.102.34	57907	216.58.217.206	443	no	ssl	allow
07/25 16:49:47	deny	L3-Trust	L3-UnTrust	192.168.102.34	57906	216.58.217.206	443	no	ssl	allow
07/25 16:49:47	deny	L3-Trust	L3-UnTrust	192.168.102.34	57905	216.58.217.206	443	no	ssl	allow
07/25 15:49:56	deny	L3-Trust	L3-UnTrust	192.168.102.34	57884	216.58.217.206	443	no	ssl	allow
07/25 15:49:56	deny	L3-Trust	L3-UnTrust	192.168.102.34	57883	216.58.217.206	443	no	ssl	allow
07/25 15:49:45	deny	L3-Trust	L3-UnTrust	192.168.102.34	57882	216.58.217.206	443	no	ssl	allow
07/25 15:49:45	deny	L3-Trust	L3-UnTrust	192.168.102.34	57881	216.58.217.206	443	no	ssl	allow
07/25 14:49:59	deny	L3-Trust	L3-UnTrust	192.168.102.34	57845	172.217.4.174	443	no	ssl	allow
07/25 14:49:39	deny	L3-Trust	L3-UnTrust	192.168.102.34	57844	172.217.4.174	443	no	ssl	allow

Similar to the example for viewing decrypted traffic logs, you can add terms to filter the traffic that you don't decrypt in a more granular fashion.

- **View The Log for a Particular Session**—To view the decryption log for a particular session, filter on the Session ID.

For example, to see the log for a session with the ID 362370, filter using the term (**sessionid eq 362370**). You can find the ID number in the Session ID column in the log output, as shown in the previous screens. If the Session ID column isn't displayed, add the column to the output.

sessionid eq 362370


Receive Time	Type	From Zone	To Zone	Source	Session ID	Destination	To Port	Decrypted	Application	Action
06/22 19:22:39	end	L3-Trust	L3-UnTrust	192.168.102.34	362370	172.217.3.206	443	yes	ssl	allow

- **View All TLS and SSH Traffic**—Filter the Traffic Logs (**Monitor > Logs > Traffic**) to view both decrypted and undecrypted TLS and SSH traffic, use the filter (**s_encrypted neq 0**):

s_encrypted neq 0

RECEIVE TIME	DECRYPTED	NAT SOURCE PORT	TYPE	FROM PORT	FROM ZONE	TO ZONE	SOURCE	DESTINATION	TO PORT	APPLICATION	ACTION	RULE
07/09 02:38:02	yes	24908	end	13429	TRUST	UNTRUST	192.168.46.2	199.167.32.141	443	web-browsing	allow	ALLOW ALL
07/09 03:34:50	no	11794	end	58727	TRUST	UNTRUST	192.168.46.2	130.211.8.336	443	personal data security	allow	ALLOW ALL
07/09 03:35:23	yes	21390	end	57251	TRUST	UNTRUST	192.168.46.2	199.167.32.141	443	web-browsing	allow	ALLOW ALL
07/09 03:37:51	yes	22386	end	90036	TRUST	UNTRUST	192.168.46.2	199.167.32.141	443	web-browsing	allow	ALLOW ALL
07/09 04:52:52	yes	8193	end	41707	TRUST	UNTRUST	192.168.46.2	199.167.32.141	443	web-browsing	allow	ALLOW ALL
07/09 05:37:53	yes	80421	end	49889	TRUST	UNTRUST	192.168.46.2	199.167.32.141	443	web-browsing	allow	ALLOW ALL
07/09 04:21:18	yes	7517	end	81833	TRUST	UNTRUST	192.168.46.2	199.167.32.141	443	web-browsing	allow	ALLOW ALL
07/09 04:09:05	no	60349	end	14431	TRUST	UNTRUST	192.168.46.2	130.211.8.336	443	personal data security	allow	ALLOW ALL
07/09 04:08:04	yes	17482	end	47936	TRUST	UNTRUST	192.168.46.2	199.167.32.141	443	web-browsing	allow	ALLOW ALL
07/09 06:32:49	yes	25821	end	39754	TRUST	UNTRUST	192.168.46.2	199.167.32.141	443	web-browsing	allow	ALLOW ALL
07/09 06:36:19	yes	25191	end	13623	TRUST	UNTRUST	192.168.46.2	199.167.32.141	443	web-browsing	allow	ALLOW ALL
07/09 06:34:50	no	35229	end	38136	TRUST	UNTRUST	192.168.46.2	130.211.8.336	443	personal data security	allow	ALLOW ALL

- **Drill Down Into the Details**—To view more information about a particular log entry, click the magnifying glass to see a detailed log view. For example, for Session ID 362370 (shown in the previous bullet), the detailed log looks like this:



The box for the **Decrypted** flag provides a second way to verify if traffic was decrypted.

You can also take upstream and downstream [packet captures](#) of decrypted traffic to view how the firewall processes SSL traffic and takes actions on packets, or perform deep packet inspection.

Decryption Broker

Decryption broker allows you to offload SSL decryption to the Palo Alto Networks next-generation 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 and to simplify your network security deployment: decryption broker eliminates the need for a third-party SSL decryption solution and allows you to reduce the number of third-party devices performing traffic analysis and enforcement. For networks without a dedicated SSL decryption appliance, decryption broker reduces latency because the traffic flow is decrypted only once.

Decryption broker is supported for PA-7000 Series, PA-5200 Series, PA-3200 Series devices and VM-300, VM-500, and VM-700 models. It requires SSL Forward Proxy decryption to be enabled, where the firewall is established as a trusted third party (or man-in-the-middle) to session traffic.



A firewall interface cannot be both a decryption broker and a GRE tunnel endpoint.

- [How Decryption Broker Works](#)
- [Decryption Broker Concepts](#)
- [Layer 3 Security Chain Guidelines](#)
- [Configure Decryption Broker with One or More Layer 3 Security Chain](#)
- [Transparent Bridge Security Chain Guidelines](#)
- [Configure Decryption Broker with a Single Transparent Bridge Security Chain](#)
- [Configure Decryption Broker with Multiple Transparent Bridge Security Chains](#)


How Decryption Broker Works

A firewall configured to perform SSL Forward Proxy decryption can be enabled as a decryption broker. Decryption broker uses dedicated decryption forwarding interfaces to connect with a security chain, a set of third-party security appliances. The firewall and the security chain together function as private analysis network.

After decrypting and inspecting SSL traffic, the firewall sends only allowed, clear text traffic on to the security chain for additional analysis and enforcement. As the firewall capacity to decrypt SSL traffic exceeds security device processing speeds, you can enable it to distribute decrypted SSL sessions among multiple security chains, in order to avoid oversubscribing any one chain. The first device in the security chain receives the clear text traffic, enforces it, and forwards allowed traffic to the next inline security chain device. The last security chain device sends the remaining allowed traffic back to the firewall. The firewall re-encrypts the traffic and forwards it to its original destination.

Two types of security chain deployments are supported: Layer 3 security chains and Transparent Bridge security chains. You might choose the type of deployment you want to set up based on the devices that make up your security chain (like if you are using stateless or stateful devices). With both security chain deployments, you can choose for the firewall to direct traffic through the security chain either unidirectionally or bidirectionally based on your analysis needs (see [Decryption Broker: Security Chain Session Flow](#) to learn more about when to use a unidirectional or bidirectional flow).

The following figure shows how decryption broker works.



Decryption Broker Concepts

A firewall acting as a decryption broker uses dedicated decryption forwarding interfaces to send decrypted traffic to a security chain—a set of inline, third-party security appliances—for additional analysis. Two types of security chain networks are supported with a decryption broker (Layer 3 security chains and Transparent Bridge security chains), and you can also choose for the firewall to direct traffic through the security chain unidirectionally or bidirectionally. A single firewall can distribute decrypted sessions among up to 64 security chains, and can monitor security chains to ensure that they are effectively processing traffic.

Review the following topics to learn more about decryption broker support and features.

- [Decryption Broker: Forwarding Interfaces](#)
- [Decryption Broker: Layer 3 Security Chain](#)
- [Decryption Broker: Transparent Bridge Security Chain](#)
- [Decryption Broker: Security Chain Session Flow](#)
- [Decryption Broker: Multiple Security Chains](#)
- [Decryption Broker: Security Chain Health Checks](#)

Decryption Broker: Forwarding Interfaces

A firewall enabled as a decryption broker uses a pair of dedicated Layer 3 interfaces to forward decrypted traffic to a security chain for inspection. The decryption forwarding interfaces must be assigned to a brand new virtual router (one that has no configured routes or other interfaces used to pass dataplane traffic); this ensures that the clear text sessions that the firewall forwards to a security chain for additional analysis are totally segmented from dataplane traffic.

In a decryption broker deployment with a Layer 3 Security Chain, a pair of two decryption forwarding interfaces can support up to 64 security chains.


A pair of decryption forwarding interfaces supports a single Transparent Bridge security chains; however, you can configure multiple decryption forwarding interface pairs to support multiple transparent bridge security chains.

Decryption Broker: Layer 3 Security Chain


In a Layer 3 security chain network, security chain devices use Layer 3 interfaces to connect to the security chain network, and each interface must have an assigned IP address and subnet mask. Security chain devices must be configured with static routes to direct inbound and outbound traffic to the next device in the security chain and back to the firewall.

Depending on the security chain session flow you choose (unidirectional or bidirectional), decrypted inbound and outbound sessions pass through the security chain in the same or opposite directions.

The figure below shows a firewall that is enabled as a decryption broker directing allowed, clear text traffic through a Layer 3 security chain bidirectionally. The firewall is configured with static routes that direct inbound sessions to a trusted, internal zone where clients reside (for example, to employees), and with a default route that directs outbound sessions to an untrusted, external zone (the Internet). For outbound sessions, the firewall uses the Primary Interface dedicated to decryption forwarding to forward inbound sessions to the first security chain device. The security chain devices use static routes to direct traffic to the next inline device; each security chain device's next hop is the subsequent device's ingress port IP address. The last security chain device's next hop is the firewall's Secondary Interface dedicated to decryption forwarding. (The flow for inbound sessions is exactly the opposite).




Alternatively, the following figure shows the same firewall enabled as a decryption broker directing decrypted traffic through a Layer 3 security chain; however, in this example, the firewall directs all sessions to flow through the security unidirectionally. The firewall uses the Primary Interface dedicated to decryption forwarding to forward both inbound and outbound sessions to the first security chain device. The last security chain device forwards both inbound and outbound sessions back to the firewall.



In both Layer 3 security chain deployments (bidirectional and unidirectional), the firewall re-encrypts the traffic the security chain returns and continues to forward it to its destination. [Configure Decryption Broker with One or More Layer 3 Security Chain](#) to get started with either of these deployments.

Decryption Broker: Transparent Bridge Security Chain

In a transparent bridge security chain network, all security chain devices are configured with two interfaces connected to the security chain network. These two interfaces are configured to be in Transparent Bridge mode; they do not have assigned IP addresses, subnet masks, default gateways, or local routing tables. Security chain devices in Transparent Bridge mode are serially connected, one after the other. They receive traffic on one interface, and then analyze and enforce the traffic. The traffic egresses the other interface and is passed to the next inline security chain device. The first image below shows a Transparent Bridge security chain deployment with a bidirectional session flow, and the second image shows a Transparent Bridge security chain with a unidirectional session flow. [Configure Decryption Broker with a Single Transparent Bridge Security Chain](#) to get started with either of these deployments.



Decryption Broker: Security Chain Session Flow

You can choose for the firewall to direct decrypted inbound and outbound sessions through a security chain in the same direction (unidirectionally) or in opposite directions (bidirectionally). For example, if you have a stateless device like a packet recorder in a security chain, you could enable traffic to flow unidirectionally through the security chain so the inbound and outbound traffic traverse the device in the same direction. The packet recorder receives both inbound and outbound traffic on the same port and can then examine packet captures from both sides of the session in order to detect changes to packet header

values. Alternatively, if the security chain includes devices like Data Loss Prevention (DLP) solutions that statefully inspect traffic, enable traffic to flow bidirectionally through the security chain instead.

Decryption Broker: Multiple Security Chains

A firewall enabled as a decryption broker supports forwarding to multiple security chains (Layer 3, Transparent Bridge, or a mix of both) in order provide redundancy and to balance the analysis load, avoiding oversubscribing a security chain or a single security chain device. Because the firewall capacity to decrypt and forward traffic can exceed the capacity of security chain devices to process traffic, you can configure the firewall to distribute clear text sessions to multiple security chain networks for inspection. The firewall can distribute sessions among both types of security chain networks, so that security chains can share the inspection load; however, the methods to enable session distribution varies depending on whether you are using Layer 3 security chains or Transparent Bridge security chains. A decryption broker forwarding to multiple Layer 3 Security Chains can distribute sessions for inspection using one of four methods:


- IP modulo—The firewall assigns sessions based on the modulo hash of the source and destination IP addresses.
- IP hash—The firewall assigns sessions based on the IP hash of the source and destination IP addresses and port numbers.
- Round robin—The firewall allocates sessions evenly amongst the security chains.
- Lowest latency—The firewall allocates more sessions to the security chain with the lowest latency.

A decryption broker forwarding to multiple Transparent Bridge Security Chains must be configured to perform policy-based session distribution; traffic matched to a policy rule is forwarded only to the security chain associated with that rule. For example, specify a different source address range for each decryption policy to dedicate a single Transparent Bridge security chain to analyze and enforce traffic originating from specified IP address ranges.

When configuring multiple security chains, make sure that you're deploying enough security chains to provide excess capacity in the event of a security chain failure. If you enable the firewall to perform Security Chain Health Checks, and a security chain fails, the firewall continues to distribute decrypted sessions among the healthy security chains. If there are not enough healthy chains to cover the additional load, that single security chain failure could result in cascading failures as the remaining healthy security chains are oversubscribed.

The first image below shows a decryption broker deployment with multiple Layer 3 security chains. Note that a single pair of Decryption Forwarding Interfaces can forward decrypted traffic to multiple Layer 3 security chains (up to 64).

The second image below shows a decryption broker deployment with multiple Transparent Bridge security chains; a dedicated pair of decryption forwarding interfaces is required to forward to each separate Transparent Bridge security chain.



Decryption Broker: Security Chain Health Checks

A decryption broker can monitor the status of security chains to ensure that they are effectively processing decrypted traffic. Periodic health checks monitor:

- Security device connectivity (Path Monitoring)
- Security device processing speed and efficiency (HTTP Latency Monitoring)
- Security device HTTP inspection capabilities (HTTP Monitoring)

For each type of monitoring you enable, you must define the conditions you want to trigger a health check failure. When a security chain fails a health check, the firewall can either:

- Block existing SSL sessions assigned to the failed security chain. The firewall will only commence to forward new decrypted sessions to that security chain for analysis when the security chain passes a subsequent health check. The traffic flow must pass both the firewall security policy check and the security chain check to be allowed to the Internet.
- (Layer 3 Security Chain Only) Allow traffic to bypass the failed security chain. Keep in mind that traffic that bypasses a security chain still undergoes firewall decryption and security policy enforcement; however, it does not undergo security chain analysis. This option is only supported with Layer 3 security chains. Because session distribution for Transparent Bridge security chains is policy-based traffic cannot bypass a failed chain as the traffic matched to a policy rule is assigned to a specific chain for inspection.

You might choose if you want the firewall to block sessions or bypass a security chain if a security chain fails based on your organization's compliance and usability needs.

When configuring multiple security chains, it is a **best practice** to deploy enough security chains to provide excess capacity in the event of a security chain failure. If you enable the firewall to perform Security Chain Health Checks, and a security chain fails, the firewall continues to distribute decrypted sessions among the healthy security chains. If there are not enough healthy chains to cover the additional load, that single security chain failure could result in cascading failures as the remaining healthy security chains are oversubscribed.

Layer 3 Security Chain Guidelines

Follow these guidelines to set up Layer 3 security chain devices to support decryption broker:

- Configure security chain devices with Layer 3 interfaces to connect to the security chain network. These Layer 3 interfaces must have an assigned IP address and subnet mask.
- Do not include devices that modify IP or TCP headers in a security chain, or be sure to disable any features that perform these functions. If the security chain returns a session to the firewall with a modified IP or TCP header, the firewall drops the session as it can no longer match it to the original pre-decrypted session.
- Set the default gateways for security chain devices:
 - For all security chain devices except the last device in the chain, configure the default gateway to be the IP address of the next inline device.
 - For the last security chain device, configure the default gateway to be the firewall's Secondary Interface IP address. This ensures that the last device returns the traffic flow to the firewall. (When you configure a decryption forwarding profile, you'll assign one of the decryption forwarding interfaces to be the decryption broker Secondary Interface. See Objects > Decryption > Forwarding Profile > Secondary Interface, and use this interface's IP address).
 - If you configured the firewall to direct sessions through the security chain bidirectionally, you must also set the default gateway of the first security chain device to be the firewall's Primary Interface IP address (When you configure a decryption forwarding profile, you'll assign one of the decryption forwarding interfaces to be the decryption broker Primary Interface. See Objects > Decryption > Forwarding Profile > Primary Interface, and use this interface's IP address).
- Confirm that the firewall and security chain can effectively communicate: check that the router that directs traffic between the firewall and the security chain is configured correctly, and that security chain devices are configured with static routes to appropriately direct traffic.
- Security chain devices should not originate traffic to a network outside of the security chain. The firewall blocks traffic that it cannot match to the original pre-decrypted session. However, if a security chain device requires Internet access to receive updates, make sure that the device can access a separate network (for example, via the device's management port) to facilitate those updates.
- When configuring multiple security chains, it is a best practice to deploy enough security chains to provide excess capacity in the event of a security chain failure. If you enable the firewall to perform

Security Chain Health Checks, and a security chain fails, the firewall continues to distribute decrypted sessions among the healthy security chains. If there are not enough healthy chains to cover the additional load, that single security chain failure could result in cascading failures as the remaining healthy security chains are oversubscribed.

Configure Decryption Broker with One or More Layer 3 Security Chain

Perform the following steps to enable the firewall to act as a decryption broker that distributes traffic to a Layer 3 Security Chain for additional analysis and enforcement. Enabling the firewall as a decryption broker includes:

- Set up a Layer 3 security chain that adheres to the Layer 3 Security Chain Guidelines.
- Activate the free decryption broker license ([Decryption Licenses](#)). This includes going to the Palo Alto Networks [Customer Support Portal](#) to activate the license, and then installing the license on the firewall.
- Enable at least two firewall interfaces as decryption forwarding interfaces. A pair of decryption forwarding interfaces can support up to 64 security chains.
- Configure a Decryption Forwarding profile to enable the firewall to forward decrypted sessions to one or multiple security chains, to distribute those sessions amongst multiple security chains, and to monitor security chain health.

STEP 1 | Follow the Layer 3 Security Chain Guidelines to make sure that you've set up your security chain to support decryption broker.

STEP 2 | Activate the free Decryption Broker license (see [Decryption Licenses](#)).

STEP 3 | Confirm that the firewall is enabled to perform SSL Forward Proxy decryption.

Select **Policies > Decryption** to Add or modify a decryption policy rule. You can also attach a decryption profile to a decryption policy rule, to perform certificate checks and to validate SSL protocols. For example, a decryption profile allows you to block sessions based on certificate status, using unsupported protocols or cipher suits, or if the resources to perform decryption are not available.

STEP 4 | Enable a pair of Layer 3 interfaces to forward decrypted traffic.

1. View configured interfaces on the **Network > Interfaces > Ethernet** tab. The Interface Type column displays if an interface is configured as a Layer 3 interface. Select a Layer 3 interface and complete the following steps for both Layer 3 interfaces that you want to enable as a Decrypt Forward pair.
2. Select the Config tab and assign the interface to a Virtual Router that has no configured routes or interfaces used to pass dataplane traffic. The virtual router must be dedicated to the decryption forwarding interfaces to ensure the clear text sessions that the firewall forwards for additional analysis are totally segmented from dataplane traffic.
3. Continue to assign the interface to a Security Zone. (Assign both interfaces to the same security zone).
4. On the Advanced tab, select Decrypt Forward.
5. Click OK to save the interface settings.
6. Repeat these steps for an even number of interfaces, pairing two as you go.
7. Make sure that the interfaces enabled to forward decrypted traffic are not being used to pass any other type of traffic.

STEP 5 | Create a Decryption Forwarding profile to define settings for the firewall to forward decrypted traffic to a Layer 3 security chain.

-
1. Select **Objects > Decryption > Forwarding Profile**, Add a new Decryption Forwarding Profile, and give the profile a descriptive Name.
 2. On the General tab, set the Security Chain Type to Routed (layer 3) to configure the firewall to forward decrypted traffic to a security chain with Layer 3 devices.
 3. Set the Flow Direction for decrypted traffic the firewall forwards: Unidirectional or Bidirectional.
 4. Select the Primary Interface and Secondary Interface the firewall uses to communicate with the security chain.

Together, the primary and secondary interfaces form a pair of decryption forwarding interfaces. Only interfaces that you have enabled to be Decrypt Forward interfaces are displayed here. Your security chain type (Layer 3 or Transparent Bridge) and the traffic flow direction (unidirectional or bidirectional) determine which of the two interfaces forwards allowed, clear text traffic to the security chain, and which interface receives the traffic back from the security chain after it has undergone additional enforcement.

5. Click OK to save the decryption profile.

STEP 6 | Connect the firewall to a security chain.

1. Select the Security Chains tab and Add a security chain.
2. Name and Enable the security chain.
3. Enter details for the First Device and Last Device in the security chain.

Give the device a descriptive Name, and select the IPv4 address of the first device in the security chain. Or, you can define a new Address Object to easily reference the device.

4. Click OK to save the security chain, and continue to repeat these steps to add another security chain. Or, continue on if you're only adding a single chain.

STEP 7 | (Multiple Security Chains Only) Continue on the Security Chains tab and choose the Session Distribution Method for the firewall to use to distribute decrypted sessions amongst security chains.

Choose for session distribution to be based on IP Modulo, IP Hash, Round Robin, or Lowest Latency. The Lowest Latency distribution method requires you to also enable the firewall to perform HTTP Latency Monitoring and HTTP Monitoring on the security chain.

STEP 8 | Select the Health Monitor tab to enable the firewall to perform Security Chain Health Checks on security chains.

If a security chain fails a health check, the firewall can then either block traffic until the security chain passes a subsequent health check and is able to process it, or the firewall can allow traffic to bypass a failed security chain.

1. On Health Check Failure, choose for the firewall to either Bypass Security Chain or Block Session.
2. Define a Health Check Failed Condition as an event where any of the health monitor conditions are met (an OR Condition), or when all of the conditions are met (an AND Condition).
3. Enable Path Monitoring, HTTP Latency Monitoring, and/or HTTP Monitoring. For each type of monitoring you want to enable, define the periods of time and/or counts that you want to trigger a health check failure.

Latency and HTTP monitoring are required to effectively support Lowest Latency session distribution (**Objects > Decryption > Forwarding Profile > Security Chains Session Distribution Method**).

STEP 9 | Save the Forwarding profile.

STEP 10 | Attach the Forwarding Profile to a decryption policy rule.

The firewall decrypts and inspects traffic the rule matches, and then forwards the clear text traffic to the security chain for further inspection and enforcement.

1. Select **Policies > Decryption** and select a decryption policy rule.
2. Select Options.
3. Set the Action to Decrypt and Forward.
4. Select the Forwarding Profile you created.
5. Click OK to save the policy rule and Commit your changes.

STEP 11 | Monitor the decrypted traffic that the firewalls forwards for additional inspection.

1. Select **Monitor > Logs > Traffic** and use the following filter: `(flags has decrypt-forwarded)`.
2. Check the details for a traffic log entry and look for the Decrypt Forwarded flag.

Transparent Bridge Security Chain Guidelines

Follow these guidelines when configuring Transparent Bridge security chain devices to support decryption brokering:

- Each security chain device must be configured with two interfaces in Transparent Bridge mode; these two interfaces connect the device to the security chain network. The security chain devices does not use a local routing table, and the Transparent Bridge interfaces do not have assigned IP addresses, subnet masks, default gateways.
- Do not include devices that modify IP or TCP headers in a security chain, or be sure to disable any features that perform these functions. If the security chain returns a session to the firewall with a modified IP or TCP header, the firewall drops the sessions as it can no longer match it to the original client-to-server or server-to-client session.
- When configuring multiple security chains, it is a best practice to deploy enough security chains to provide excess capacity in the event of a security chain failure. If you enable the firewall to perform Security Chain Health Checks, and a security chain fails, the firewall continues to distribute decrypted sessions among the healthy security chains. If there are not enough healthy chains to cover the additional load, that single security chain failure could result in cascading failures as the remaining healthy security chains are oversubscribed.

Configure Decryption Broker with a Single Transparent Bridge Security Chain

Perform the following steps to enable the firewall to act as a decryption broker that distributes traffic to a Transparent Bridge Security Chain for additional analysis and enforcement. Enabling the firewall as a decryption broker includes:

- Set up a Transparent Bridge security chain that adheres to the Transparent Bridge Security Chain Guidelines.
- Activate the free decryption broker license ([Decryption Licenses](#)). This includes going to the Palo Alto Networks [Customer Support Portal](#) to activate the license, and then installing the license on the firewall.
- Enable a pair of Layer 3 firewall interfaces as decryption forwarding interfaces. Each pair of decryption forwarding interfaces supports one transparent bridge security chain; you'll need to create multiple decryption forwarding interface pairs to support multiple Transparent Bridge security chains.
- Configure a Decryption Forwarding profile to enable the firewall to forward decrypted sessions to a Transparent Bridge security chain and to monitor security chain performance.

Even if you plan to enable decryption broker with multiple Transparent Bridge security chains, you must perform the following steps first.

STEP 1 | Set up a Transparent Bridge security chain following the Transparent Bridge Security Chain Guidelines.

STEP 2 | Activate the free Decryption Broker license (see [Decryption Licenses](#)).

STEP 3 | Confirm that the firewall is enabled to perform SSL Forward Proxy decryption.

Select **Policies > Decryption** to **Add** or modify a decryption policy rule. You can also attach a decryption profile to a decryption policy rule, to perform certificate checks and validate SSL protocols. For example, a decryption profile allows you to block sessions based on certificate status, using unsupported protocols or cipher suits, or if the resources to perform decryption are not available.

STEP 4 | Enable a pair of Layer 3 interfaces to forward decrypted traffic.

1. View configured interfaces on the **Network > Interfaces > Ethernet** tab.

The Interface Type column displays if an interface is configured as a Layer 3 interface. Select a Layer 3 interface and complete the following steps for both Layer 3 interfaces that you want to enable as a Decrypt Forward pair.

2. Select the **Config** tab and assign the interface to a **Virtual Router** that has no configured routes or interfaces used to pass dataplane traffic.

The virtual router must be dedicated to the decryption forwarding interfaces to ensure the clear text sessions that the firewall forwards for additional analysis are totally segmented from dataplane traffic.

3. Continue to assign the interface to a **Security Zone**. (Assign both interfaces to the same security zone).
4. On the **Advanced** tab, select **Decrypt Forward**.
5. Click **OK** to save the interface settings.
6. Repeat these steps so that at least two interfaces are enabled to forward decrypted traffic.

A pair of two decryption forwarding interfaces supports a single Transparent Bridge Security Chain. If you want the firewall to distribute decrypted sessions amongst multiple Transparent Bridge security chains, continue to enable a pair of decryption forwarding interfaces for each security chain you want to support. Make sure that the interfaces enabled to forward decrypted traffic are not being used to pass any other type of traffic.

STEP 5 | Create a Decryption Forwarding profile to define settings for the firewall to forward decrypted traffic to a Transparent Bridge security chain.

1. Select **Objects > Decryption > Forwarding Profile**, **Add** a new Decryption Forwarding Profile, and give the profile a descriptive **Name**.
2. On the **General** tab, set the **Security Chain Type** to **Transparent Bridge** to configure the firewall to forward decrypted traffic to a security chain with Transparent Bridge devices.
3. Set the **Flow Direction** for decrypted traffic the firewall forwards: **Unidirectional** or **Bidirectional**.
4. Select the **Primary Interface** and **Secondary Interface** the firewall uses to forward traffic to the security chain.

Together, the primary and secondary interfaces form a pair of decryption forwarding interfaces. Only interfaces that you have enabled to be Decrypt Forward interfaces are displayed here.

STEP 6 | Select the **Health Monitor** tab to enable the firewall to perform health checks on a Transparent Bridge security chain.

1. Set **On Health Check Failure** to **Block Session** if you want to drop traffic until the health check succeeds or set it to **Bypass Security Chain** to forward traffic without going through the security chain.

Transparent Bridge security chain session distribution is policy-based, so traffic cannot fail over to a different security chain (as it can in Layer 3 mode) because the traffic matched to a policy rule is assigned to a specific chain for inspection.

2. Define a **Health Check Failed Condition** as an event where any of the health monitor conditions are met (an **OR Condition**), or when all of the conditions are met (an **AND Condition**).
3. Enable **Path Monitoring**, **HTTP Latency Monitoring**, and/or **HTTP Monitoring**. For each type of monitoring you want to enable, define the periods of time and/or counts that you want to trigger a health check failure.

Latency and HTTP monitoring are required to effectively support Lowest Latency session distribution (**Objects > Decryption > Forwarding Profile > Security Chains > Session Distribution Method**).

STEP 7 | Save the Forwarding Profile.

STEP 8 | Attach the Forwarding Profile to a decryption policy rule.

The firewall decrypts and inspects traffic the rule matches, and then forwards the clear text traffic to the security chain for further inspection and enforcement.

1. Select **Policies > Decryption** and select a decryption policy rule.
2. Use the policy rule tabs to define the traffic that you want to forward to the associated Transparent Bridge security chain.

For example, select **Source** and **Add a Source Address** range, or click **New Address** to create an address objects that identifies traffic originating from a given IP address range. The policy rule will enforce only traffic that originates from this source.

3. Select **Options**.
4. Set the **Action** to **Decrypt and Forward**.
5. Select a **Transparent Bridge Forwarding Profile**.
6. Click **OK** to save the policy rule and **Commit** your changes.

STEP 9 | (Optional) Continue to [Configure Decryption Broker with Multiple Transparent Bridge Security Chains](#).

STEP 10 | Monitor the decrypted traffic that the firewall has forwarded for additional inspection.

- Select **Monitor > Logs > Traffic** and add the filter: (`flags has decrypt-forwarded`).
- Check the details for a traffic log entry and look for the Decrypt Forwarded flag.

Configure Decryption Broker with Multiple Transparent Bridge Security Chains

You can configure the firewall to distribute sessions among multiple Multiple Security Chains, where the security chains are in Transparent Bridge mode. For each Transparent Bridge security chain you want to support, you must configure:

- A pair of decryption forwarding interfaces that forward traffic only to that single Transparent Bridge security chain.
- A Decryption Forwarding profile that specifies settings only for that single Transparent Bridge security chain.
- A Decryption policy rule that specifies only for certain decrypted traffic to be forwarded to that single Transparent Bridge security chain. This allows you to distribute sessions more evenly among multiple Transparent Bridge security chains (in order to avoid oversubscribing any one security chain) based on traffic origin.

-
- STEP 1** | First, follow the steps to [Configure Decryption Broker with a Single Transparent Bridge Security Chain](#). For each Transparent Bridge security chain you want to support, this includes:
- On the firewall, enable a pair of Layer 3 interfaces to support forwarding of decrypted traffic.
 - Create a Decryption Forwarding profile to define settings for the firewall to forward decrypted traffic to a Transparent Bridge security chain.
- STEP 2** | Attach each Transparent Bridge Decryption Forwarding profile to a separate decryption policy rule.
- In addition to applying the decryption forwarding settings to matching traffic, attaching Transparent Bridge Decryption Forwarding profiles to decryption policies rules allows you to distribute sessions amongst the Transparent Bridge Security chains. Specify a different source address range for each policy rule to dedicate a single Transparent Bridge security chain to analyze and enforce traffic originating from that range.
1. Select **Policies > Decryption** and select a decryption policy rule.
 2. Select **Source** and **Add** a Source Address range, or click **New Address** to create a new address objects that identifies traffic originating from a given IP address range. Only traffic originating from this IP address range is forwarded to the associated Transparent Bridge security chain for analysis.
 3. Select **Options**.
 4. Set the **Action** to **Decrypt and Forward**.
 5. Select a **Transparent Bridge Forwarding Profile** to attach to the policy rule.
 6. Click **OK** to save the policy rule and **Commit** your changes.
- STEP 3** | Continue to repeat these steps—associated one Transparent Bridge decryption forwarding profile with one decryption policy—for as many security chains as you want to support.

Activate Free Licenses for Decryption Features

No licenses are required to decrypt [SSH traffic](#) and SSL traffic ([SSL internet traffic](#) or [SSL traffic to an internal server](#)).

However, you must activate a free license in order to enable [Decryption Broker](#) and [Decryption Mirroring](#). The free license requirement ensures that these features can only be used after the approved personnel purposefully activates the associated license.

Follow these steps on the Palo Alto Networks [Customer Support Portal](#) to activate a decryption broker or decryption mirroring feature license.

STEP 1 | Log in to the [Customer Support Portal](#).

STEP 2 | Select **Assets > Devices** on the left-hand navigation pane.

STEP 3 | Find the device on which you want to enable decryption broker or decryption port mirroring and select **Actions** (the pencil icon).

STEP 4 | Under Activate Licenses, select **Activate Feature License**.

STEP 5 | Select the feature for which you want to activate a free license: **Decryption Port Mirror** or **SSL Decryption Broker**.

DEVICE LICENSES

Serial Number: 00000000000000000000

Model: PAN-PA-VM-300

Device Name: 00000000000000000000

Feature Name	Authorization Code	Expiration Date	Actions
VM-300 Bundle	00000000000000000000	Perpetual	
Threat Prevention		03/25/2023	
PAN-DB URL Filtering		03/25/2023	
GlobalProtect Gateway		03/25/2023	
Premium Support		03/25/2023	
AutoFocus Device License	00000000000000000000	02/11/2023	
WildFire License		03/25/2023	
PA-VM		Perpetual	

ACTIVATE LICENSES

- ☐ Activate Auth-Code
- ☐ Activate Trial License
- ☒ Activate Feature License
- ☐ Activate Upgrade License

AVAILABLE FEATURE LICENSES

- ☒ Decryption Port Mirror
- ☒ SSL Decryption Broker

STEP 6 | Agree and Submit.

STEP 7 | Install the decryption broker or decryption mirroring license on the firewall.

1. Select **Device > Licenses**.
2. Click **Retrieve license keys from the license server**.
3. Verify that the **SSL Decryption Broker** or the **Decryption Port Mirror** license is now active on the firewall.
4. Restart the firewall (**Device > Setup > Operations**). Decryption port mirroring and decryption broker are not available for configuration until the firewall reloads.

URL Filtering

Palo Alto Networks URL Filtering allows you to monitor and control the sites users can access, to prevent phishing attacks by controlling the sites to which users can submit valid corporate credentials, and to enforce safe search for search engines like Google and Bing.

- > [About URL Filtering](#)
- > [How URL Filtering Works](#)
- > [URL Filtering Vendors](#)
- > [URL Filtering Use Cases](#)
- > [URL Categories](#)
- > [Plan Your URL Filtering Deployment](#)
- > [URL Filtering Best Practices](#)
- > [Configure URL Filtering](#)
- > [Monitor Web Activity](#)
- > [Log Only the Page a User Visits](#)
- > [Create a Custom URL Category](#)
- > [URL Category Exceptions](#)
- > [Use an External Dynamic List in a URL Filtering Profile](#)
- > [Allow Password Access to Certain Sites](#)
- > [Safe Search Enforcement](#)
- > [URL Filtering Response Pages](#)
- > [Customize the URL Filtering Response Pages](#)
- > [HTTP Header Logging](#)
- > [Request to Change the Category for a URL](#)
- > [Troubleshoot URL Filtering](#)
- > [PAN-DB Private Cloud](#)

About URL Filtering


Palo Alto Networks URL Filtering protects against web-based threats by giving you a way to safely enable web access while controlling how your users interact with online content.

With URL Filtering enabled, all web traffic (HTTP and HTTPS) on any port is compared against the URL filtering database, which contains a listing of millions of websites that have been categorized. You can use these URL categories as a match criteria to enforce security policy. You can also use URL filtering to enforce safe search settings for your users and to [Prevent Credential Phishing](#) based on URL category.

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-600 appliances that function as PAN-DB servers within your network.

How URL Filtering Works

PAN-DB—the URL Filtering cloud database—classifies websites based on site content, features, and safety. A URL can have up to four URL categories, including [risk categories](#) (high, medium, and low) that indicate the likelihood that the site will expose you to threats. As PAN-DB categorizes sites, firewalls with URL Filtering enabled can leverage that knowledge in real-time to enforce security policy.



When a user accesses a URL that's not cached, the firewall checks PAN-DB for the site's category and saves it. As the firewall saves new entries, it removes URLs that users have not accessed recently so that it accurately reflects the traffic in your network.

When the firewall checks PAN-DB for a URL, it also looks for critical updates, such as URLs that previously qualified as benign but are now malicious. Every 30 minutes, the firewall checks PAN-DB for such updates.

If you believe PAN-DB has incorrectly categorized a site, you can [submit a URL category change request](#) in your browser through [Test A Site](#) or directly from the firewall logs.



Did you know?

Technically, the firewall caches URLs on both the management plane and the dataplane:

- *PAN-OS 9.0 and later releases do not download PAN-DB seed databases. Instead, upon activation of the URL filtering license, the firewall populates the cache as URL queries are made.*
- *The management plane holds more URLs and communicates directly with PAN-DB. When the firewall cannot find a URL's category in the cache and performs a lookup in PAN-DB, it caches the retrieved category information in the management plane. The management plane passes that information along to the dataplane, which also caches it and uses it to enforce policy.*
- *The dataplane holds fewer URLs and receives information from the management plane. After the firewall checks [URL category exception lists](#) and [custom URL categories](#) for a URL, the next place it looks is the dataplane. Only if the firewall cannot find the URL*

categorized in the dataplane does it check the management plane and, if the category information is not there, PAN-DB.

URL Filtering Vendors

Palo Alto Networks firewalls support two URL filtering vendors:

- **PAN-DB**—The Palo Alto Networks-developed URL filtering database. 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 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 cyberthreats. URL categories that identify confirmed malicious content—malware, phishing, and C2 are updated every five minutes—to ensure that you can manage access to these sites within minutes of categorization.

To view a list of PAN-DB URL filtering categories, refer to <https://urlfiltering.paloaltonetworks.com/CategoryList.aspx>.

- **BrightCloud**—A third-party URL database that is owned by Webroot, Inc. and is integrated into PAN-OS firewalls. For information on the BrightCloud URL database, visit <http://brightcloud.com>.

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 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.

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:
 1. On a host with Internet access, go to the [Palo Alto Networks Customer Support web site](https://www.paloaltonetworks.com/support/tabs/overview.html), www.paloaltonetworks.com/support/tabs/overview.html, and log in.
 2. In the Resources section, click **Dynamic Updates**.
 3. In the BrightCloud Database section, click **Download** and save the file to the host.
2. Upload the database to the firewall:
 1. Log in to the firewall, select **Device > Dynamic Updates** and click **Upload**.
 2. For the **Type**, select **URL Filtering**.
 3. 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:
 1. Select **Device > Dynamic Updates** and click **Install From File**.
 2. For the **Type**, select **URL Filtering**. The firewall automatically selects the file you just uploaded.
 3. 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
```

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.


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, the firewall ceases to perform PAN-DB URL Filtering; URL category enforcement, 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.




STEP 2 | Activate PAN-DB URL Filtering.



PAN-OS 9.0 and later releases do not download PAN-DB seed databases. Instead, upon activation of the URL filtering license, the firewall populates the cache as URL queries are made.

1. Click **Activate**. The value in the Active field changes to Yes.



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 [Safe Search Enforcement](#).

URL Filtering Use Cases

There are many ways to use URL Filtering beyond only blocking and allowing certain sites. For example, you can use multiple categories per URL to allow users to access a site, but block particular functions like submitting corporate credentials or downloading files. You can also use URL categories to enforce different [types of policy](#), such as Authentication, Decryption, QoS, and Security.

Read on for more about the different ways that you can use URL Filtering.

Control web access based on URL category

You can [create a URL Filtering profile](#) that specifies an action for a URL category and attach the profile to a policy rule. The firewall enforces policy against traffic based on the settings in the profile. For example, to block all gaming websites you could 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.

Multi-Category URL Filtering

Every URL can have up to four categories, including a [risk category](#) that indicates the likelihood a site will expose you to threats. More granular URL categorizations means that you can move beyond a basic “block-or-allow” approach to web access. Instead, you can control how your users interact with online content that, while necessary for business, is more likely to be used as part of a cyberattack.

For instance, you might consider certain URL categories risky to your organization, but are hesitant to block them outright as they also provide valuable resources or services (like cloud storage services or blogs). Now, you can allow users to visit sites that fall into these types of URL categories while you protect your network by decrypting and inspecting traffic and enforcing read-only access to the content.

For a URL category that you want to tightly control, set the URL Filtering profile action to alert as part of the steps to [Configure URL Filtering](#). Then continue to follow the [Malicious URL Categories](#): decrypt the URL category, block dangerous file downloads, and turn on credential phishing prevention.

Block or allow corporate credential submissions based on URL category

[Prevent Credential Phishing](#) by enabling the firewall to detect corporate credential submissions to sites, and then control those submissions based on URL category. Block users from submitting credentials to malicious and untrusted sites, warn users against entering corporate credentials on unknown sites or reusing corporate credentials on non-corporate sites, and explicitly allow users to submit credentials to corporate and sanctioned sites.

Enforce Safe Search Settings

Many search engines have a safe search setting that filters out adult images and videos from search results. You can enable the firewall to block search results if the end user is not using the strictest safe search settings, and you can transparently enable safe search for your users. The firewall supports safe search enforcement for the following search providers: Google, Yahoo, Bing, Yandex, and YouTube. See how to get started with [Safe Search Enforcement](#).

Enforce Password Access to Certain Sites

You can block access to a site for most users while allowing certain users to access the site. See how to [Allow Password Access to Certain Sites](#).

Block high-risk file downloads from certain URL categories

You can block high-risk file downloads from specific URL categories by creating a security policy with a [File Blocking profile](#) attached.

Enforce Security, Decryption, Authentication, and QoS policies based on URL category

You can enforce different types of firewall policies based on URL categories. For example, suppose you have enabled [Decryption](#), but you want to exclude certain personal information from being decrypted. In this case you could create a decryption policy rule that excludes websites that match the URL categories *financial-services* and *health-and-medicine* from decryption. Another example would be to use the URL category *streaming-media* in a QoS policy to apply bandwidth controls to websites that fall in to this category.

The following table describes the policies that accept URL categories as match criteria:

Policy Type	Description
Decryption  <i>Decrypting traffic based on URL categories is both a best practice for both URL Filtering and Decryption.</i>	You can also use URL categories to phase-in decryption, and to exclude URL categories that might contain sensitive or personal information from decryption (like <i>financial-services</i> and <i>health-and-medicine</i>). Plan to decrypt the riskiest traffic first (URL Categories most likely to harbor malicious traffic, such as gaming or high-risk) and then decrypt more as you gain experience. Alternatively, decrypt the URL Categories that don't affect your business first (if something goes wrong, it won't affect business), for example, news feeds. In both cases, decrypt a few URL Categories, listen to user feedback, run reports to ensure that decryption is working as expected, and then gradually decrypt a few more URL Categories, and so on. Plan to make decryption exclusions to exclude sites from decryption if you can't decrypt them for technical reasons or because you choose not to decrypt them.
Authentication	To ensure that users authenticate before being allowed access to a specific category, you can attach a URL category as a match criterion for Authentication policy rules.
QoS	Use URL categories to allocate throughput levels for specific website categories. For example, you may want to allow the <i>streaming-media</i> category, but limit throughput by adding the URL category to a QoS policy rule.
Security	In security policy rules, you can use URL categories in two ways: • Enforce policy based on URL categories by selecting them as match criteria. • Attach a URL Filtering profile that specifies the policy action for each category.

Policy Type	Description
	If for example, the IT-security group in your company needs access to the <i>hacking</i> category, but all other users are denied access to the category, you must create the following rules: • A Security policy rule that allows the IT-Security group to access content categorized as <i>hacking</i>. The Security policy rule references the <i>hacking</i> category in the Services/URL Category tab and IT-Security group in the Users tab. • Another Security policy rule that allows general web access for all users. To this rule you attach a URL filtering profile that blocks the <i>hacking</i> category. You must list the policy that allows access to <i>hacking</i> before the policy that blocks <i>hacking</i>. This is because the firewall evaluates security policy rules from the top down, so when a user who is part of the security group attempts to access a <i>hacking</i> site, the firewall evaluates the policy rule that allows access first and grants the user access. The firewall evaluates users from all other groups against the general web access rule that blocks access to the <i>hacking</i> sites.

URL Categories

PAN-DB classifies websites based on site content, features, and safety. A URL can have up to four categories, including risk categories (high, medium, and low), which indicate how likely it is that the site will expose you to threats.

Visit [Test A Site](#) to see how PAN-DB categorizes a URL, and to learn about all available URL categories. You can also use Test A Site to submit a URL Category change request, or you can submit the request directly in the firewall: select **Monitor > Logs** and open the details for a log entry. Under the URL Category, you'll see the option to submit a change request.

Read on to learn more about URL categories:

- [Security-Focused URL Categories](#)
- [Malicious URL Categories](#)
- [Policy Actions You Can Take Based on a URL Category](#)

Security-Focused URL Categories

Security-focused URL categories can help you to reduce your attack surface by providing targeted decryption and enforcement for sites that pose varying levels of risk, but are not confirmed malicious. Websites are classified with a security-related category only so long as they meet the criteria for that category; as site content changes, policy enforcement dynamically adapts. You cannot submit a change request for security-focused URL Categories.

Security-Focused URL Categories

High-Risk

High-risk sites include:

- Sites previously confirmed to be malware, phishing, or C2 sites that have displayed only benign activity for at least 30 days.
- Unknown domains are classified as high-risk until PAN-DB completes site analysis and categorization.
- Sites that are associated with confirmed malicious activity. For example, a page might be high-risk if there are malicious hosts on the same domain, even if the page itself does not contain malicious content.
- Bulletproof ISP-hosted sites.
- Sites hosted on IPs from ASNs that are known to allow malicious content.

Default and Recommended Policy Action: Alert

Medium-Risk

Medium-risk sites include:

- All cloud storage sites (with the URL category **online-storage-and-backup**).
- Sites previously confirmed to be malware, phishing, or C2 sites that have displayed only benign activity for at least 60 days.
- Unknown IP addresses are categorized as medium-risk until PAN-DB completes site analysis and categorization.

Security-Focused URL Categories

	Default and Recommended Policy Action: Alert
Low-Risk	Sites that are not medium or high risk are considered low risk. These sites have displayed benign activity for a minimum of 90 days. Default and Recommended Policy Action: Allow
Newly-Registered Domains	Identifies sites that have been registered within the last 32 days. New domains are frequently used as tools in malicious campaigns. Default Policy Action: Alert Recommended Policy Action: Block  <i>Newly-registered domains are often generated purposefully or by domain generation algorithms and used for malicious activity. It is a best practice to block this URL category.</i>

Malicious URL Categories

We strongly recommend that you block the URL categories that identify malicious or exploitive content. To get started, you can clone the default URL Filtering profile which blocks malware, phishing, and command-and-control URL categories by default. The default URL Filtering profile also blocks the abused-drugs, adult, gambling, hacking, questionable, and weapons URL categories. Whether to block these URL categories depends on your business requirements. For example, a university probably won't want to restrict student access to most of these sites because availability is important, but a business that values security first may block some or all of them.

- **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. If availability is critical to your business and you must allow the traffic, alert on unknown sites, apply the best practice Security profiles to the traffic, and investigate the alerts.



PAN-DB Real-Time Updates learns unknown sites after the first attempt to access an unknown site, so unknown URLs are identified quickly and become known URLs that the firewall can then handle based on the actual URL category.

- **newly-registered-domain**—Newly registered domains are often generated purposely or by domain generation algorithms and used for malicious activity.
- **command-and-control**—Command-and-control URLs and domains used by malware and/or compromised systems to surreptitiously communicate with an attacker's remote server to receive malicious commands or exfiltrate data.
- **copyright-infringement**—Domains with illegal content, such as content that allows illegal download of software or other intellectual property, which poses a potential liability risk. This category was introduced to enable adherence to child protection laws required in the education industry as well as

laws in countries that require internet providers to prevent users from sharing copyrighted material through their service.

- **extremism**—Websites promoting terrorism, racism, fascism, or other extremist views discriminating against people or groups of different ethnic backgrounds, religions or other beliefs. This category was introduced to enable adherence to child protection laws required in the education industry. In some regions, laws and regulations may prohibit allowing access to extremist sites, and allowing access may pose a liability risk.
- **proxy-avoidance-and-anonymizers**—URLs and services often used to bypass content filtering products.
- **grayware**—Websites and services that do not meet the definition of a virus but are malicious or questionable and may degrade device performance and cause security risks. Prior to Content release version 8206, the firewall placed grayware in either the malware or questionable URL category. If you are unsure about whether to block grayware, start by alerting on grayware, investigate the alerts, and then decide whether to block grayware or continue to alert on grayware.
- **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.

For categories that you decide to alert on, instead of block, you can very strictly control how users interact with site content. For example, give users access to the resources they need (like developer blogs for research purposes or cloud storage services), but take the following precautions to reduce exposure to web-based threats:

- ❑ Follow the Anti-Spyware, Vulnerability Protection, and File Blocking [best practices](#). A protective measure would be to block downloads of dangerous file types and blocking obfuscated JavaScript for sites that you are alerting on.
- ❑ [Target decryption](#) based on URL category. A good start would be to decrypt high-risk and medium-risk sites.
- ❑ [Display a response page](#) to users when they visit high-risk and medium-risk sites. Alert them that the site they are attempting to access is potentially malicious, and advise them on how to take precautions if they decide to continue to the site.
- ❑ [Stop credential theft](#) by blocking users from submitting their corporate credentials to sites including those that are high-risk and medium-risk.

Policy Actions You Can Take Based on URL Categories

On the firewall, you can use a URL Filtering profile to specify how you would like to enforce URL categories. By default, site access for all URL categories is 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 is not logged. Customize the URL Filtering profile by deciding what type of **Site Access** you want to enforce for each category. To [Prevent Credential Phishing](#), you can also allow or disallow **User Credential Submissions** based on URL category (for example, you can block user credential submissions to medium and high-risk sites). Users can still access these sites, but cannot enter submit their corporate credentials to them.

To start enforcing the actions you've defined in a URL Filtering, you'll need to attach the profile to a security policy rule. The firewall enforces the profile actions on traffic that matches the security policy rule (for details, see [Configure URL Filtering](#)).



Learn more about configuring a [best practice URL Filtering profile](#) to ensure protection against URLs that have been observed hosting malware or exploitative content.

Action	Description
Site Access	

Action	Description
alert	The website is allowed and a log entry is generated in the URL filtering log.  <i>Set alert as the Action for categories of traffic you don't block to log and provide visibility into the traffic.</i>
allow	The website is allowed and no log entry is generated.  <i>Don't set allow as the Action for categories of traffic you don't block because you lose visibility into traffic you don't log. Instead, set alert as the Action for categories of traffic you don't block to log and provide visibility into the traffic.</i>
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. Blocking site access for a URL category also sets User Credential Submissions for that URL category to block.
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.  <i>The Continue page doesn't display properly on client systems configured to use a proxy server.</i>
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 Allow Password Access to Certain Sites. In earlier release versions, URL Filtering category overrides had priority enforcement ahead of custom URL categories. As part of the upgrade to PAN-OS 9.0, URL category overrides are converted to custom URL categories, and no longer receive priority enforcement over other custom URL categories. Instead of the action you defined for the category override in previous release versions, the new custom URL category is enforced by the security policy rule with the strictest URL Filtering profile action. From most strict to least strict, possible URL Filtering profile actions are: block, override, continue, alert, and allow. This means that, if you had URL category overrides with the action allow, there's a possibility the overrides might be blocked after they are converted to custom URL category in PAN-OS 9.0.

Action	Description
	 <i>The Override page doesn't display properly on client systems configured to use a proxy server.</i>
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.

User Credential Permissions

 *These settings require you to first [Set Up Credential Phishing Prevention](#).*

alert	Allow users to submit corporate credentials to sites in this URL category, but generate a URL Filtering alert log each time this occurs.
allow (default)	Allow users to submit corporate credentials to websites in this URL category.
block	Block users from submitting corporate credentials to websites in this category. A default anti-phishing response page is displayed to users when they access sites to which corporate credential submissions are blocked. You can choose to create a custom block page to display.
continue	Display a response page to users that prompts them to select Continue to access to access the site. By default, the Anti Phishing Continue Page is shown to user when they access sites to which credential submissions are discouraged. You can also choose to create a custom response page to display—for example, if you want to warn users against phishing attempts or reusing their credentials on other websites.

Plan Your URL Filtering Deployment

To first deploy URL filtering in your network, we recommend that you start with a basic setup that'll give you visibility into web activity patterns while blocking confirmed malicious content:

- ❑ Start with a (mostly) passive URL Filtering profile that alerts on most categories. This gives you visibility into the sites your users are accessing, so you can decide what you want allow, limit, and block.
- ❑ Block URL categories that we know are bad: malware, C2, and phishing.

Because alerting on all web activity might create a large amount of log files, you might decide you only want to do this as you're initially deploying URL Filtering.



At that time, 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.

STEP 1 | At any time, you can use [Test A Site](#) to see how PAN-DB—the URL Filtering cloud database—categorizes a specific URL, and to learn about all possible URL categories.


You can also use Test A Site to submit a [change request](#), if you disagree with how a specific URL is categorized.

STEP 2 | Create a passive URL Filtering profile, that alerts on all categories so you have visibility into web traffic.

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 3 | Configure the action for all categories to **alert**, except for malware, command-and-control, and phishing, which should remain blocked.

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. **Block** access to known dangerous URL categories.



Block access to malware, phishing, dynamic-dns, unknown, command-and-control, extremism, copyright-infringement, proxy-avoidance-and-anonymizers, newly-registered-domain, grayware, and parked URL categories.

4. Click **OK** to save the profile.

STEP 4 | 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.

STEP 5 | Save the configuration.

Click **Commit**.

STEP 6 | View the URL filtering logs to see all of the website categories that your users are accessing. The categories you've set to block are also logged.

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 set to any action other than **allow**. URL Filtering reports give you a view of web activity in a 24-hour period. (**Monitor > Reports**).

STEP 7 | Here's what to do next:

- **Strictly control how users interact with high-risk and medium-risk sites.**

PAN-DB categorizes every URL with up to four categories, and every URL has a risk category (high, medium, and low). While high and medium-risk sites are not confirmed malicious, they are closely associated with malicious sites. For example, they might be on the same domain as malicious sites or maybe they hosted malicious content until only very recently. For everything that you do not explicitly allow or block, you can [use risk categories to write simple policy based on website safety](#).

You can take precautionary measures to limit your users' interaction high-risk sites especially, as there might be some cases where you want to give your users access to sites that might also present safety concerns (for example, you might want to allow your developers to use developer blogs for research, yet blogs are a category known to commonly host malware).

- **Pair URL Filtering with [User-ID](#).**

Use URL Filtering and User-ID together to control web access based on organization or department and to block corporate credential submissions to unsanctioned sites:

- URL Filtering [prevents credential theft](#) by detecting corporate credential submissions to sites based on the site category. Block users from submitting credentials to malicious and untrusted sites, warn users against entering corporate credentials on unknown sites or reusing corporate credentials on non-corporate sites, and explicitly allow users to submit credentials to corporate sites.
- Add or update a security policy rule with the passive URL Filtering profile so that it applies to a department user group, for example, Marketing or Engineering (**Policies > Security > User**). Monitor the department activity, and get feedback from department members to understand the web resources that are essential to the work they do.
- Consider [all the ways you can use URL Filtering](#) to reduce your attack surface and to control web usage.

For example, if you're a school, you can use URL Filtering to enforce strict safe search settings, where search engines filter out adult images and videos from search results. Or, if you have a security operations center, you might give threat analysts password access to compromised or dangerous sites for research, that you might not want to otherwise open up to entire organizations or teams.

- **Follow the [URL Filtering Best Practices](#).**

URL Filtering Best Practices

Palo Alto Networks URL Filtering protects you from web-based threats, and gives you a simple way to monitor and control web activity. These best practices show you how to reduce your exposure to web-based threats, without limiting user access to web resources that they need.

- ❑ As part of building a best practice internet gateway security policy, [Identify the applications you want to allow](#) and [create application allow rules](#).

Allowed applications include not only the applications you provision and administer for business and infrastructure purposes, but also the applications that your users need to get their jobs done and applications you might want to allow for personal use. [Use URL categories to build application filters](#) that allow general types you want to permit either for business or personal use.

- ❑ [Get visibility](#) into your users' web activity, so you can plan the most effective URL Filtering policy for your organization, and roll it out smoothly:
 - Use [Test A Site](#) to see how PAN-DB—the URL Filtering cloud database—categorizes a specific URL you're interested in, and to learn about all possible URL categories.
 - Start with a passive URL Filtering profile that alerts on most URL categories. Alerting on URL categories means that log entries are generated for each site your users access, but no policy action is enforced. With this visibility into the sites your users are accessing, you can decide what you want to allow, limit, and block. Because alerting on most URL categories results in a large volume of logs, we recommend only using this passive profile as a planning tool.
 - Monitor web activity to assess the sites your users are accessing and see how they align with your business needs.
- ❑ [Block URL categories](#) that classify malicious and exploitive web content.

While we know that these categories are dangerous, always keep in mind that the URL categories that you decide to block might depend on your business needs. Review the [review the full list of URL categories](#) and decide if there are other categories you'd like to block (for example, schools might decide to block gaming and adult sites).

- ❑ [Use URL categories to phase-in decryption, and to exclude sensitive or personal information from decryption](#).



Targeting decryption and phasing it in based on URL categories are also [Decryption best practices](#).

- ❑ [Exclude certain web traffic from decryption](#)—Use URL categories to make [decryption exceptions](#), especially if you want to exclude certain types of sites from decryption because of privacy concerns (for example, financial-services and health-and-medicine sites).
- ❑ [Phase in decryption](#)—Plan to decrypt the traffic that is classified as highfirst. Alternatively, decrypt the URL Categories that don't affect your business first (for example, news feeds).
- ❑ [Turn on credential phishing prevention, to stop corporate credential theft](#).

For all categories, set User Credential Submissions to block (**Objects > Security Profiles > URL Filtering > Categories > User Credential Submission**). Users will still be able to use their corporate credentials to access [allowed](#) applications and resources; however, they won't be able to reuse their corporate credentials on non-corporate sites or inadvertently enter them on a phishing site.

- ❑ [Decrypt, inspect, and strictly limit how users interact with high-risk and medium-risk URL categories, as well as any Malicious URL Categories](#) if you're unable to block them for business reasons.

The applications that you allow and the malicious URL categories that you block outright are just a portion of your overall web traffic. The rest of the content your users are accessing is a combination of benign (the low-risk URL category) and risky content (high-risk and medium-risk URL categories). The high-risk and medium-risk categories identify content that is not confirmed malicious but is closely

associated with malicious sites. For example, a high-risk URL might be on the same domain as a malicious site, or it might have hosted malicious content in the past.

However, many sites that pose a risk to your organization also provide valuable resources and services to your users (cloud storage services are a good example). While these resources and services are necessary for business, they are also more likely to be used as part of a cyberattack.

Control how users interact with potentially dangerous content, while still providing a good user experience:

- **Decrypt** high-risk and medium-risk sites.

The firewall evaluates decryption policy rules from top to bottom (**Policies > Decryption**). Place the rule to decrypt high-risk and medium-risk URL categories below the rules that specify your **decryption exceptions**. This way, the firewall enforces decryption exceptions first and then decrypts high-risk and medium-risk web content.



*Review your **decryption exceptions** and consider blocking the high-risk and medium-risk sites that you don't decrypt (like financial-services and health-and-medicine). This is a protective measure you can take while continuing to adhere to requirements to exclude certain content from decryption:*

- Your users can continue to access the low-risk, encrypted content.
- You're meeting requirements to not decrypt content that belongs to the URL category (like health-and-medicine).
- You're protecting your organization from the high-risk and medium-risk sites that you're not able to decrypt and inspect due to privacy concerns or other business reasons.

To do this, create a URL Filtering profile with the high-risk and medium-risk URL categories set to block (Objects > Security Profiles > URL Category > Categories). Next, set up a security policy rule that specifies that URL categories you are excluding from decryption, such as health-and-medicine and financial-services, as policy match criteria (Policies > Security > Service/URL Category). Select the Action tab in the security policy rule, and attach the URL Filtering profile to the rule. Then, set the rule Action to Allow—the policy rule allows only low-risk web traffic that matches the URL categories you specified (such as health-and-medicine and financial-services) and blocks the matching URL category traffic that is high-risk or medium-risk.

- In a URL Filtering profile, set the high-risk and medium-risk categories to **continue** to **display a response page** that warns users they're visiting a potentially-dangerous site. Advise them how to take precautions if they decide to continue to the site. If you don't want to prompt users with a response page, alert on the high-risk and medium-risk categories instead.
 - For high-risk and medium-risk URL categories, follow the Anti-Spyware, Vulnerability Protection, and File Blocking **best practices**. A protective measure would be to block downloads of dangerous file types and blocking obfuscated JavaScript.
- ❑ **If your educational institution has policies against allowing students or faculty to access sites such as adult or weapons sites, **enforce safe search** to make sure these sites are blocked.**

You can enable safe search transparently for your users; this means that if it's not already turned on, you automatically redirect users to filtered search results.

- ❑ **(PAN-OS 9.0.4 and later) Enable the firewall to hold an initial web request as it looks up a website's URL category with PAN-DB.**

When a user visits a website, a firewall with URL Filtering enabled checks its local cache of URL categories to categorize the site. If the firewall doesn't find the URL's category in the cache, it performs a lookup in PAN-DB. By default, the firewall allows the user's web request during this cloud lookup and enforces policy when the server responds.

But when you choose to hold web requests, the firewall blocks the request until it either finds the URL category or times out. If the lookup times out, the firewall considers the URL category not-resolved.

1. Enter `configure` to access Configuration Mode.
2. Enter `set deviceconfig setting ctd hold-client-request yes` to enable the feature.
3. Commit your changes.

Configure URL Filtering

After you [Determine URL Filtering Policy Requirements](#), you should have a basic understanding of the types of websites and website categories your users are accessing. Use this information to create custom URL filtering profiles and attach them to the Security policy rules that allow web access. In addition to managing web access with a URL Filtering profile, if you configure User-ID™, you can manage the sites to which users can submit corporate credentials.

STEP 1 | Create a URL Filtering profile.



If you didn't already, configure a [best practice URL Filtering profile](#) to ensure protection against URLs hosting malware or exploitive content.

Select **Objects > Security Profiles > URL Filtering** and **Add** or modify a URL Filtering profile.

STEP 2 | Define site access for each URL category.

Select **Categories** and set the Site Access for each URL category:

- **allow** traffic destined for that URL category; allowed traffic is not logged.
- Select **alert** to have visibility into sites that users are accessing. Traffic matching that category is allowed but a URL Filtering log is generated to record when a user accesses a site in that category.
- Select **block** to deny access to traffic that matches that category and to enable logging of the blocked traffic.
- Select **continue** to display a page to users with a warning and require them to click **Continue** to proceed to a site in that category.
- To only allow access if users provide a configured password, select **override**. For more details, see [Allow Password Access to Certain Sites](#).

STEP 3 | Configure the URL Filtering profile to detect corporate credential submissions to websites that are in allowed URL categories.



To ensure the best performance and a low false positive rate, the firewall automatically skips checking the credential submissions for any App-ID™ associated with sites that have never been observed hosting malware or phishing content—even if you enable checks in the corresponding category. The list of sites for which the firewall skips credential checking is automatically updated through Applications and Threats content updates.

1. Select **User Credential Detection**.
2. Select one of the [Methods to Check for Corporate Credential Submissions](#) to web pages from the **User Credential Detection** drop-down:
 - **Use IP User Mapping**—Checks for valid corporate username submissions and verifies that the username matches the user logged in to the source IP address of the session. To use this method, the firewall matches the submitted username against its IP address-to-username mapping table. To use this method, you can use any of the user mapping methods described in [Map IP Addresses to Users](#).
 - **Use Domain Credential Filter**—Checks for valid corporate usernames and password submissions and verifies that the username maps to the IP address of the logged-in user. See [Configure User Mapping Using the Windows User-ID Agent](#) for instructions on how to set up User-ID to enable this method.

- **Use Group Mapping**—Checks for valid username submissions based on the user-to-group mapping table populated when you configure the firewall to [Map Users to Groups](#).

With group mapping, you can apply credential detection to **any** part of the directory or to a specific group, such as groups like IT that have access to your most sensitive applications.



This method is prone to false positives in environments that do not have uniquely structured usernames, so you should only use this method to protect your high-value user accounts.

3. Set the **Valid Username Detected Log Severity** that the firewall uses to log detection of corporate credential submissions (default is medium).

STEP 4 | Allow or block users from submitting corporate credentials to sites based on URL category to [Prevent Credential Phishing](#).



To ensure the best performance and a low false positive rate, the firewall automatically skips checking the credential submissions for any App-ID associated with sites that have never been observed hosting malware or phishing content—even if you enable checks in the corresponding category. The list of sites for which the firewall skips credential checking is automatically updated through Applications and Threats content updates.

1. For each URL category to which you allow **Site Access**, select how you want to treat **User Credential Submissions**:
 - **alert**—Allow users to submit credentials to the website but generate a URL Filtering alert log each time a user submits credentials to sites in this URL category.
 - **allow** (default)—Allow users to submit credentials to the website.
 - **block**—Displays the [Anti Phishing Block Page](#) to block users from submitting credentials to the website.
 - **continue**—Present the [Anti Phishing Continue Page](#) to require users to click **Continue** to access the site.
2. [Configure the URL Filtering profile to detect corporate credential submissions to websites that are in allowed URL categories.](#)

STEP 5 | Define [URL Category Exception Lists](#) to specify websites that should always be blocked or allowed, regardless of URL category.

For example, to reduce URL Filtering logs, you may want add you corporate websites in the allow list so that no logs are generated for those sites or, if there is a website that is being overly used and is not work-related, you can add that site to the block list.

Traffic to websites in the block list is always blocked regardless of the action for the associated category and traffic to URLs in the allow list is always allowed.

For more information on the proper format and wildcard usage, see [URL Category Exception Lists](#).

1. Select **Overrides** and enter URLs or IP addresses in the **Block List** and select an action:
 - **block**—Block the URL.
 - **continue**—Prompt users to click **Continue** before they can proceed to the web page.
 - **override**—Prompt users 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.

STEP 6 | Enable [Safe Search Enforcement](#).

STEP 7 | Log only [Container Pages](#) for URL filtering events.

1. Select **URL Filtering Settings**. Enable **Log container page only** (default) so that the firewall logs only the main page that matches the category, not subsequent pages or categories that loaded within the container page.
2. To enable logging for all pages and categories, disable the **Log container page only** option.

STEP 8 | Enable [HTTP Header Logging](#) for one or more of the supported HTTP header fields.

Select **URL Filtering Settings** and select one or more of the following fields to log:

- **User-Agent**
- **Referer**
- **X-Forwarded-For**

STEP 9 | Save the URL Filtering profile and commit your changes.

1. Click **OK**.
2. Click **Commit**.



To test the URL filtering configuration, access a website in a category that is set to either block or continue and then observe whether the firewall performs the appropriate action.

STEP 10 | ([PAN-OS 9.0.4 and later PAN-OS 9.0 releases only](#)) Enable **hold-client-request** to block client requests while the firewall performs URL category lookups.

1. [Access the CLI](#).
2. Enter **configure** to access Configuration Mode.
3. Enter **set deviceconfig setting ctd hold-client-request yes** to enable the feature.
4. [Commit](#) your changes.



Enable this feature as a [URL Filtering Best Practice](#).

STEP 11 | Set the amount of time before a URL category lookup times out.

1. [Access the CLI](#).
2. Enter the **configure** command to access configuration mode.
3. Enter the **set deviceconfig setting ctd url-wait-timeout** command followed by a number of seconds (default is 5).
4. [Commit](#) your changes.

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, 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 **Network Activity** widgets allow 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.



View logs and configure log options:

- From the ACC, you can jump directly to the logs () or select **Monitor > Logs > URL Filtering**.

The log action for each entry depends on the Site Access setting you defined for the corresponding category:

- **Alert log**—In this example, the computer-and-internet-info category is set to alert.

Receive Time	Category	URL	From Zone	To Zone	Source	Source User	Destination	Application	Action
12/14 18:11:04	computer-and-internet-info	outlook.office365.com	any	any	192.168.1.100	any	192.168.1.100	outlook-web-online	alert

- **Block log**—In this example, the insufficient-content category is set to continue. If the category had been set to block instead, the log Action would be block-url.

Receive Time	Category	URL	From Zone	To Zone	Source	Source User	Destination	Application	Action
12/14 17:53:05	insufficient-content	music4u.markets.net	untrust	trust	10.47.20.1	anonymous	10.47.20.1	url	block-continue

- **Alert log on encrypted website**—In this example, the category is private-ip-addresses and the application is web-browsing. This log also indicates that the firewall decrypted this traffic.

Receive Time	Category	URL	Decrypted	From Zone	To Zone	Source	Destination	Application	Action
12/15 09:22:33	private-ip-addresses	updates.update-service2.us...	yes	trust	untrust	10.47.20.1	10.47.20.1	web-browsing	alert

- 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.


Receive Time	Category	URL	From Zone	To Zone	Source	Source User	Destination	Application	Action
12/15 09:58:47	insufficient-content				10.47.20.1		10.47.20.1	url	block-continue
12/15 09:58:47	insufficient-content				10.47.20.1		10.47.20.1	url	block-continue
12/15 09:58:48	insufficient-content				10.47.20.1		10.47.20.1	url	block-continue
12/15 09:58:48	insufficient-content	updates.update-service2.us...			10.47.20.1		10.47.20.1	url	block-continue
12/15 09:58:48	insufficient-content	updates.update-service2.us...			10.47.20.1		10.47.20.1	url	block-continue
12/15 09:58:48	insufficient-content	updates.update-service2.us...			10.47.20.1		10.47.20.1	url	block-continue

- 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.

PCAP	Receive Time	Type	Application	Action	Rule	Bytes	Severity	Category	Verdict	URL	File Name
	2016/12/15 11:55:37	url	outlook-web-online	allow	Allowed Personal Apps	15796		computer-and-internet-info			
	2016/12/15 11:55:48	url	outlook-web-online	alert	Allowed Personal Apps		informat...	computer-and-internet-info		outlook...	

- Generate 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 cover the 24-hour period of the date you select on the calendar. You can also export the report to PDF, CSV, or XML.



View the User Activity Report

This report provides a quick method of viewing user or group activity and also provides an option to view browse time activity.

STEP 1 | Configure a User Activity Report.

1. Select **Monitor > PDF Reports > User Activity Report**.
2. **Add** a report and enter a **Name** for it.
3. Select the report **Type**:
 - Select **User** to generate a report for one person.
 - 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.*

4. Enter the **Username/IP Address** for a user report or enter the group name for a user group report.
5. Select the time period. You can select an existing time period, or select **Custom**.
6. Select the **Include Detailed Browsing** check box, so browsing information is included in the report.

User Activity Report

Name: Doc Team

Type: ☐ User ☒ Group

Group Name: webgrubs

Time Period: Last 30 Days

☒ Include Detailed Browsing

Run Now OK Cancel


STEP 2 | Run the report.

1. Click **Run Now**.
2. When the firewall finishes generating report, click one of the links to download it:
 - Click **Download User Activity Report** to download a PDF version of the report.
 - Click **Download URL Logs** to download a CSV file of the corresponding log entries.



3. After downloading the report, click **Cancel**.
4. If you want to save the user activity report settings so you can run the same report again later, click **OK**; otherwise click **Cancel**.

STEP 3 | View the user activity report by opening the file that you downloaded. The PDF version of the report shows the user or group on which you based the report, the report time frame, and a table of contents:



STEP 4 | Click an item in the table of contents to view the report details. For example, click **Traffic Summary by URL Category** to view statistics for the selected user or group.

paloalto

Traffic Summary by URL Category


Category	Count	Bytes
computer-and-internet-ids	7.7%	279.3M
business-and-economy	1.3%	16.7M
private-ip-addresses	818	27.6M
google	347	1.3M
web-based-email	276	16.8M
ML_education	276	2.8M
search-engines	262	987.3k
web-advertisements	210	2.0M
internet-communications-and-chatting	178	1.8M
content-delivery-networks	147	5.0M
online-storage-and-backup	71	2.9M
internet-portals	47	251.2k
social-networking	40	592.7k
personal-education	28	129.8k
shopping	8	65.3k

Configure Custom URL Filtering Reports

To generate a detailed report that you can schedule to run regularly, configure a custom URL Filtering report. You can choose any combination of URL Filtering log fields on which to base the report.


STEP 1 | Add a new custom report.

1. Select **Monitor > Manage Custom Reports** and **Add** a report.
2. Give the report a unique **Name**, and optionally a **Description**.
3. Select the **Database** you want to use to generate the report. To generate a detailed URL Filtering report, select **URL** from the Detailed Logs section:




STEP 2 | Configure report options.

1. Select a predefined **Time Frame** or select **Custom**.
2. Select the log columns to include in the report from the Available Columns list add them (+) to the Selected Columns. For example, for a URL Filtering report you might select:
 - Action
 - App Category
 - Category
 - Destination Country
 - Source User
 - URL



3. If the firewall is enabled to [Prevent Credential Phishing](#), select the Attribute **Flags**, the Operator **has** and the Value **Credential Detected** to also include events in the report that record when a user submitted a valid corporate credential to a site.



4. (Optional) Select a **Sort By** option to set the attribute to use to aggregate the report details. If you do not select an attribute to sort by, the report will return the first N number of results without any aggregation. Select a **Group By** attribute to use as an anchor for grouping data. The following example shows a report with **Group By** set to **App Category** and **Sort By** set to a **Count** of **Top 5**.

Custom Report							
Report Setting		Weekly URL Filtering Summary (100%) x					
	App Category	Category	Action	Source User	Destination Country	URL	
6	general internet	online-storage-and-backup	alert		United States	g1d-keyvalueservice...	183
7		online-storage-and-backup	alert		United States	g29-keyvalueservice...	67
8		computer-and-internet info	alert		United States	g4p1.apple.com...	66
9		google	alert		United States	www.google.com/	57
10		online-storage-and-backup	alert		United States	g11-keyvalueservice...	55
11	collaboration	computer-and-internet info	alert		United States	outlook.office36...	762
12		computer-and-internet info	alert		United States	outlook.office36...	390
13		computer-and-internet info	alert		United States	outlook.office36...	251
14		computer-and-internet info	alert		United States	outlook.office36...	78
15		computer-and-internet info	alert		United States	outlook.office36...	67
16	media	streaming media	alert		United States	www.pandora.c...	214

STEP 3 | Run the report.

1. Click the **Run Now** icon to immediately generate the report, which opens in a new tab.
2. When you are done reviewing the report, go back to the **Report Setting** tab and either tune the settings and run the report again, or continue to the next step to schedule the report.
3. Select 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.

STEP 4 | Commit the configuration.

STEP 5 | View the custom report.

1. Select **Monitor > Reports**.
2. Expand the **Custom Reports** pane in the right column and select the report you want to view. The latest report displays automatically.
3. To view the report for a previous date, select the date from the calendar. You can also export the report to PDF, CSV, or XML.

Log Only the Page a User Visits

A container page is the main page that a user accesses when visiting a website, but additional pages might be loaded along with the main page. If the **Log Container page only** option is enabled in a URL filtering profile (**Objects > Security Profiles > URL Filtering**), 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 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.

Create a Custom URL Category

You can create a custom URL Filtering object to specify exceptions to URL category enforcement, and to create a custom URL category based on multiple URL categories:

- **Define exceptions to URL category enforcement**—Create a custom list of URLs that you want to use as match criteria in a security policy rule. This is a good way to specify exceptions to URL categories, where you'd like to enforce specific URLs differently than the URL category to which they belong.
- **Define a custom URL category based on multiple PAN-DB categories**—This allows you to target enforcement for websites that match a set of categories. The website or page must match *all* the categories defined as part of the custom category.

Follow these steps to create a custom URL category, and define how you'd like the firewall to enforce the custom URL category:

STEP 1 | Select **Objects > Custom Objects > URL Category**.

STEP 2 | **Add** or modify a custom URL Category, and give the category a descriptive **Name**.


STEP 3 | Set the category **Type** to either **Category Match** or **URL List**:

- **URL List**—Add URLs that you want to enforce differently than the URL category to which they belong. Use this list type to define exceptions for URL Category enforcement, or to define a list of URLs as belonging to a custom category. For details on how to populate this list, such as guidelines on how to use wildcards, see [URL Category Exceptions](#).
- **Category Match**—Provide targeted enforcement for websites that match a set of categories. The website or page must match *all* the categories defined as part of the custom category.

STEP 4 | Select **OK** to save the custom URL category.

STEP 5 | Select **Objects > Security Profiles > URL Filtering** and **Add** or modify a URL Filtering profile.

Your new custom category will be listed under the **Custom URL Categories** drop down:



STEP 6 | Decide how you want to enforce **Site Access** and **User Credential Submissions** for the custom URL Category. (To control the sites to which users can submit their corporate credentials, see [Prevent Credential Phishing](#)).

STEP 7 | Attach the URL Filtering Profile to a security policy rule, to enforce traffic that matches that rule.

Select **Policies > Security > Actions** and specify for the security policy rule to enforce traffic based on the URL Filtering profile you just updated. Make sure to **Commit** your changes.



You can also use custom URL categories as security policy match criteria. In this case, you do not need to define how the category should be enforced as part of a URL Filtering Profile. After setting up the custom category, go directly to the Security policy rule to which you want to add the custom URL category (Policies > Security). Select Service/URL Category to use the custom URL category as match criteria for the rule.

URL Category Exceptions

You can exclude specific websites from URL category enforcement, ensuring that these websites are blocked or allowed regardless of their associated URL category. For example, you could block a URL category but choose to allow certain sites that fall within that category. To create these kinds of exceptions to URL category enforcement:

- Add the IP addresses or URLs of the sites you want to explicitly block or allow by creating a [Create a Custom URL Category](#) list (**Objects > Custom Objects > URL Category**).
- Use an External Dynamic List in a URL Filtering profile. The benefit to using an External Dynamic List to specify the sites you want to enforce separately from their URL categories is that you can update the External Dynamic List without performing a configuration change or commit on the firewall.

The following guidelines describe how to populate URL Category block and allow lists, or a text file that you're using as the source of an external dynamic list for URLs:

- [Basic Guidelines For URL Category Exception Lists](#)
- [Wildcard Guidelines for URL Category Exception Lists](#)
- [URL Category Exception List—Wildcard Examples](#)

Basic Guidelines For URL Category Exception Lists

- Enter the IP addresses or URLs of websites that you want to enforce separately from the associated URL category.
- List entries must be an exact match and are case-insensitive.
- Enter a string that is an exact match to the website (and possibly, specific subdomain) for which you want to control access, or use wildcard characters to allow an entry to match to multiple website subdomains. For details on using wildcard characters, review [Wildcard Guidelines for URL Category Exception Lists](#).
- Omit `http` and `https` from URL entries.
- Each URL entry can be up to 255 characters in length.

Wildcard Guidelines for URL Category Exception Lists

You can use wildcards in URL Category exception lists to easily configure a single entry to match to multiple website subdomains and pages, without having to specify exact subdomains and pages.

Follow these guidelines when creating wildcard entries:

- The following characters are considered token separators: `.` `/` `&` `=` `;` `+`
Every string separated by one or two of these characters is a token. Use wildcard characters as token placeholders, indicating that a specific token can contain any value.
- In place of a token, you can use either an asterisk (*) or a caret (^) to indicate a wildcard value.
- Wildcard characters must be the only character within a token; however, an entry can contain multiple wildcards.

How to Use Asterisk (*) and Caret (^) Wildcards

You can use either * and ^ as wildcards in your custom categories and URL EDLs, but you cannot use both at the same time. This means that if you use * to represent a wildcard in one custom category or URL EDL, you must use * as the wildcard in every other custom category or URL EDL in your configuration—you can no longer use ^. For example, you cannot have one custom category or EDL that contains `^ .foo.com` and a separate custom category or URL EDL that contains `www.xyz.com/*`.

For details about how to use each wildcard, see below:

*	Use to indicate one or more variable subdomains. If you use *, the entry will match any additional subdomains, whether at the beginning or the end of the URL. Use a forward slash at the end of the entry if you do not want to match any additional subdomains beyond that point. Ex: • *.paloaltonetworks.com matches www.paloaltonetworks.com and www.paloaltonetworks.com.uk. • *.paloaltonetworks.com/ matches www.paloaltonetworks.com but not www.paloaltonetworks.com.uk.
^	Use to indicate one variable subdomain. Ex: mail.^ .com matches to mail.company.com but not mail.company.sso.com.



Do not create an entry with consecutive asterisk (*) wildcards or more than nine consecutive caret (^) wildcards—entries like these can affect firewall performance.

For example, do not add an entry like **mail.*.*.com**; instead, depending on the range of websites you want to control access to, enter **mail.*.com** or **mail.^.^ .com**. An entry like **mail.*.com** matches to a greater number of sites than **mail.^.^ .com**; **mail.*.com** matches to sites with any number of subdomains and **mail.^.^ .com** matches to sites with exactly two subdomains.

URL Category Exception List—Wildcard Examples

The following table lists examples of URL exception lists entries using wildcards, and examples of the sites that these entries match to.

URL Exception List Entry	Matching Sites
Example Set 1	
*.company.com	eng.tools.company.com support.tools.company.com tools.company.com docs.company.com
^.company.com	tools.company.com docs.company.com

URL Exception List Entry	Matching Sites
^.^.company.com	eng.tools.company.com support.tools.company.com
Example Set 2	
mail.google.*	mail.google.com mail.google.co.uk
mail.google.^^	mail.google.com
mail.google.^^.^^	mail.google.co.uk

Use an External Dynamic List in a URL Filtering Profile

To protect your network from newly-discovered threats and malware, you can use [External Dynamic Lists](#) in URL Filtering profiles. External dynamic lists give you the ability to update the list without a configuration change or commit on the firewall. 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 the list is updated on the web server, the firewall retrieves the changes and applies policy to the modified list without requiring a commit on the firewall.

The firewall dynamically imports the list at the configured interval and enforces policy for the URLs (IP addresses or domains are ignored) in the list. For URL formatting guidelines, see [URL Category Exceptions](#).

For more information, see [External Dynamic List](#).

STEP 1 | Configure the Firewall to Access an External Dynamic List.

- Ensure that the list does not include IP addresses or domain names; the firewall skips non-URL entries.
- Verify the formatting of the list (see).
- Select **URL List** from the Type drop-down.

STEP 2 | 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.
 1. Select **Policies > Security**.
 2. Select the **Actions** tab and, in the Profile Setting section, select the new profile in the **URL Filtering** drop-down.
 3. Click **OK** and **Commit**.

STEP 3 | Test that the policy action is enforced.

1. [View External Dynamic List Entries](#) for the URL list, and attempt to access a URL from the list.
2. Verify that the action you defined is enforced in the browser.
3. To monitor the activity on the firewall:
 1. Select **ACC** and add a URL Domain as a global filter to view the Network Activity and Blocked Activity for the URL you accessed.
 2. Select **Monitor > Logs > URL Filtering** to access the detailed log view.

STEP 4 | 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 firewall model.



To check whether you have reached the limit for an external dynamic list type, select **Objects > External Dynamic Lists** and click **List Capacities**.

Use the following CLI command on a firewall to review the details for a list.

```
request system external-list show type url name <list_name>
```

For example:

```
request system external-list show type url name My_URL_List
vsys5/My_URL_List:
Next update at: Tue Jan 3 14:00:00 2017
Source: http://example.com/My_URL_List.txt
Referenced: Yes
Valid: Yes
Auth-Valid: Yes

Total valid entries: 3
Total invalid entries: 0
Valid urls:
www.URL1.com
www.URL2.com
www.URL3.com
```

Allow Password Access to Certain Sites

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:

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 302 to prompt for the password, which applies on a per-vsyt level.



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**.

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:
 1. Select **Network > Interface Mgmt** and click **Add**.
 2. 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**.

Safe Search Enforcement

Many search engines have a safe search setting that filters out adult images and videos in search query return traffic. You can enable the firewall to block search results if the end user is not using the strictest safe search settings, and you can also transparently enable safe search for your users. The firewall supports safe search enforcement for the following search providers: Google, Yahoo, Bing, Yandex, and YouTube. Consider that safe search is a best-effort setting and service providers do not guarantee that it works with every website, and search providers classify sites as safe or unsafe (not Palo Alto Networks).

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 then blocks any matching search query return traffic that is not using the strictest safe search settings. There are two methods to enforce safe search:

- [Block Search Results when Strict Safe Search is not Enabled](#)—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.
- [Transparently Enable Safe Search for Users](#)—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 is supported for Google, Yahoo, and Bing searches.

As safe search settings differ by search provider, get started by reviewing the different safe search implementations. There are then two ways you can enforce safe search: you can block search results when safe search is disabled, or you can transparently enable safe search for your users:

- [Safe Search Settings for Search Providers](#)
- [Block Search Results when Strict Safe Search is not Enabled](#)
- [Transparently Enable Safe Search for Users](#)

Safe Search Settings for Search Providers

Safe search settings differ for each search provider—review the following settings to learn more.

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 <code>FF=</code> 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

Search Provider	Safe Search Setting Description
	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.  <i>PAN-OS supports safe search enforcement for YouTube through HTTP header insertion. HTTP header insertion is not currently supported for HTTP/2. To enforce safe search for YouTube, App-ID and HTTP/2 Inspection downgrade HTTP/2 connections to HTTP/1.1 using the Strip ALPN feature in the appropriate decryption profile.</i> <i>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.</i>
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 <code>vm=</code> and passed to the server each time the user performs a Yahoo search. Appending <code>vm=r</code> to a Yahoo search query URL also enables the strictest safe search settings.  <i>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.</i>
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 <code>adt1=</code> and passed to the server each time the user performs a Bing search. Appending <code>adt=strict</code> 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.

Block Search Results when Strict Safe Search is not Enabled

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 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 [Transparently Enable Safe Search for Users](#).

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:
 1. On the **Categories** tab, set the **search-engines** category to **block**.
 2. 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 site access for each URL category](#).
 - [Define Block and Allow Lists to specify websites that should always be blocked or allowed, regardless of URL category](#).
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:
 1. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.
 2. Enter a **Name** for the category, such as `SearchEngineDecryption`.
 3. **Add** the following to the Sites list:

www.bing.*

www.google.*

search.yahoo.*

4. 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 4 | (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:
 1. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.
 2. Enter a **Name** for the category, such as EnableBingSafeSearch.
 3. **Add** the following to the Sites list:
`www.bing.com/images/*`
`www.bing.com/videos/*`
 4. Click **OK** to save the custom URL category object.
2. Create another URL filtering profile to block the custom category you just created:
 1. Select **Objects > Security Profiles > URL Filtering**.
 2. **Add** a new profile and give it a descriptive **Name**.
 3. Locate the custom category in the Category list and set it to **block**.
 4. Click **OK** to save the URL filtering profile.
3. Add a security policy rule to block Bing SSL traffic:
 1. Select **Policies > Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.
 2. On the **Actions** tab, attach the URL filtering profile you just created to block the custom Bing category.
 3. On the **Service/URL Category** tab **Add a New Service** and give it a descriptive **Name**, such as bingssl.
 4. Select **TCP** as the **Protocol** and set the **Destination Port** to 443.
 5. Click **OK** to save the rule.
 6. 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**.

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:

Search Blocked

User: 192.168.2.10

Your search results have been blocked because your search settings are not in accordance with company policy. In order to continue, please update your search settings so that Safe Search is set to the strictest setting, and try your search again.

For more information, please refer to: <http://www.bing.com/account/general>

Please contact your system administrator if you believe this message is in error.

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.

Transparently Enable Safe Search for Users

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.

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**.

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:
 1. On the **Categories** tab, set the **search-engines** category to **block**.
 2. 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 site access for each URL category.](#)
 - [Define Block and Allow Lists to specify websites that should always be blocked or allowed, regardless of URL category.](#)
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.

STEP 4 | (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:
 1. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.
 2. Enter a **Name** for the category, such as `EnableBingSafeSearch`.
 3. **Add** the following to the Sites list:
`www.bing.com/images/*`
`www.bing.com/videos/*`
 4. Click **OK** to save the custom URL category object.
2. Create another URL filtering profile to block the custom category you just created:
 1. Select **Objects > Security Profiles > URL Filtering**.
 2. **Add** a new profile and give it a descriptive **Name**.
 3. Locate the custom category you just created in the Category list and set it to **block**.
 4. Click **OK** to save the URL filtering profile.
3. **Add** a security policy rule to block Bing SSL traffic:
 1. Select **Policies > Security** and **Add** a policy rule that allows traffic from your trust zone to the Internet.
 2. On the **Actions** tab, attach the URL filtering profile you just created to block the custom Bing category.
 3. On the **Service/URL Category** tab **Add a New Service** and give it a descriptive **Name**, such as `bingssl`.
 4. Select **TCP** as the **Protocol**, set the **Destination Port** to **443**.
 5. Click **OK** to save the rule.
 6. 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 following text and then save the file.

```
<html>
<head>
  <title>Search Blocked</title>
  <meta http-equiv="Content-Type" content="text/html; charset=utf-8">
  <meta http-equiv="pragma" content="no-cache">
  <meta name="viewport" content="initial-scale=1.0">
</style>
  #content {
    border:3px solid#aaa;
    background-color:#fff;
    margin:1.5em;
    padding:1.5em;
    font-family:Tahoma,Helvetica,Arial,sans-serif;
    font-size:1em;
  }
```

```

    h1 {
    font-size:1.3em;
    font-weight:bold;
    color:#196390;
    }
    b {
    font-weight:normal;
    color:#196390;
    }
</style>
</head>
<body bgcolor="#e7e8e9">
  <div id="content">
    <h1>Search Blocked</h1>
    <p>
      <b>User:</b>
      <user/>
    </p>
    <p>Your search results have been blocked because your search
settings are not in accordance with company policy. In order to
continue, please update your search settings so that Safe Search is
set to the strictest setting. If you are currently logged into your
account, please also lock Safe Search and try your search again.</p>
    <p>
      For more information, please refer to:
      <a href="<ssurl/>">
        <ssurl/>
      </a>
    </p>
    <p id="java_off"> Please enable JavaScript in your browser.<br></
p>
    <p><b>Please contact your system administrator if you believe this
message is in error.</b></p>
  </div>
</body>
<script>
  // Grab the URL that's in the browser.
  var s_u = location.href;
  //bing
  // Matches the forward slashes in the beginning, anything, then
".bing." then anything followed by a non greedy slash. Hopefully the
first forward slash.
  var b_a = /^.*\/\/(.+\.bing\..+?)\/\/.exec(s_u);
  if (b_a) {
    s_u = s_u + "&adlt=strict";
    window.location.replace(s_u);
    document.getElementById("java_off").innerHTML = 'You are being
redirected to a safer search!';
  }
  //google
  // Matches the forward slashes in the beginning, anything, then
".google." then anything followed by a non greedy slash. Hopefully the
first forward slash.
  var g_a = /^.*\/\/(.+\.google\..+?)\/\/.exec(s_u);
  if (g_a) {
    s_u = s_u.replace(/&safe=off/ig, "");
    s_u = s_u + "&safe=active";
    window.location.replace(s_u);
    document.getElementById("java_off").innerHTML = 'You are being
redirected to a safer search!';
  }
  //yahoo

```

```
// Matches the forward slashes in the beginning, anything, then
".yahoo."" then anything followed by a non greedy slash. Hopefully the
first forward slash.
var y_a = /^.*\/\/(.+\.yahoo\..+?)\/\//.exec(s_u);
if (y_a) {
    s_u = s_u.replace(/&vm=p/ig, "");
    s_u = s_u + "&vm=r";
    window.location.replace(s_u);
    document.getElementById("java_off").innerHTML = 'You are being
redirected to a safer search!';
}
document.getElementById("java_off").innerHTML = ' ';
</script>
</html>
```

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.

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:
 1. Select **Objects > Custom Objects > URL Category** and **Add** a custom category.
 2. Enter a **Name** for the category, such as `SearchEngineDecryption`.
 3. **Add** the following to the Sites list:


```
www.bing.*
www.google.*
search.yahoo.*
```
 4. 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**.

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 [Container Pages](#) is enabled:

- **URL Filtering and Category Match Block Page**




- **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, ([Allow Password Access to Certain Sites](#)), 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 rule with a URL Filtering profile that has the Safe Search Enforcement option enabled (see [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.




- **Anti Phishing Block Page**

This page displays to users when they attempt to enter corporate credentials (usernames or passwords) on a web page in a category for which credential submissions are blocked. The user can continue to access the site but remains unable to submit valid corporate credentials to any associated web forms. To control the sites to which users can submit corporate credentials, the firewall must be configured with User-ID and enabled to [Prevent Credential Phishing](#) based on URL category.



- **Anti Phishing Continue Page**

This page warns users against submitting credentials (usernames and passwords) to a web site. Warning users against submitting credentials can help to discourage them from reusing corporate credentials and to educate them about possible phishing attempts. They must select Continue to proceed to credentials on the site. To control the sites to which users can submit corporate credentials, the firewall must be configured with User-ID and enabled to [Prevent Credential Phishing](#) based on URL category.



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.

Table 2: URL Filtering Response Page Variables

Variable	Usage
<user/>	The firewall replaces the variable with the username (if available via User-ID) or IP address of the user when displaying the response page.
<url/>	The firewall replaces the variable with the requested URL when displaying the response page.
<category/>	The firewall replaces the variable with the URL filtering category of the blocked request.
<pan_form/>	HTML code for displaying the Continue button on the URL Filtering Continue and Override page.

You can also add code that triggers the firewall to display different messages depending on what URL category the user is attempting to access. For example, the following code snippet from a response page specifies to display Message 1 if the URL category is games, Message 2 if the category is travel, or Message 3 if the category is kids:

```
var cat = "<category/>";
switch(cat)
{
  case 'games':
    document.getElementById("warningText").innerHTML = "Message 1";
    break;
  case 'travel':
    document.getElementById("warningText").innerHTML = "Message 2";
    break;
  case 'kids':
    document.getElementById("warningText").innerHTML = "Message 3";
    break;
}
```

Only a single HTML page can be loaded into each virtual system for each type of block page. However, other resources such as images, sounds, and cascading style sheets (CSS files) can be loaded from other servers at the time the response page is displayed in the browser. All references must include a fully qualified URL.

Table 3: Response Page References

Reference Type	Example HTML Code
Image	<pre></pre>
Sound	<pre><embed src="http://simplythebest.net/sounds/WAV/WAV_files/movie_WAV_files/do_not_go.wav" volume="100" hidden="true" autostart="true"></pre>

Reference Type	Example HTML Code
Style Sheet	<pre><link href="http://example.com/style.css" rel="stylesheet" type="text/css" /></pre>
Hyperlink	<pre>View Corporate Policy</pre>

Customize the URL Filtering Response Pages

The firewall provides predefined [URL Filtering Response Pages](#) that display by default when a user:

- A user attempts to browse to a site in a category with restricted access.
- A user submits valid corporate credentials to a site for which credential detection is enabled ([Prevent Credential Phishing](#) based on URL category).
- [Log Only the Page a User Visits](#) blocks a search attempt.

However, you can create your own custom response pages with your corporate branding, acceptable use policies, and links to your internal resources.



Custom response pages larger than the maximum supported size are not decrypted or displayed to users. In PAN-OS 8.1.2 and earlier PAN-OS 8.1 releases, custom response pages on a decrypted site cannot exceed 8,191 bytes; the maximum size is increased to 17,999 bytes in PAN-OS 8.1.3 and later releases.

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 [Table 2: 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 [Table 3: 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.

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.



You can also use HTTP headers to manage access to SaaS applications. You don't need a URL Filtering license to do this, but you must use a URL Filtering profile to turn this feature on.

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.

Request to Change the Category for a URL

If you think that a URL is not categorized accurately, you can request for us to categorize it differently. Submit a change request directly in the firewall, or use [Test A Site](#). A change request triggers PAN-DB—the URL Filtering cloud—to do an immediate analysis of the URL for which you're suggesting a category change. If PAN-DB validates that the new category suggestion is accurate, the change request is approved. If PAN-DB does not find the new category suggestion to be accurate, the change request is then reviewed by human editors from the Palo Alto Networks threat research and data science teams.

After you've submitted a change request, you'll receive an email from us confirming that we've received your request. When we've completed our investigation, you'll receive a second email confirming the results.

You cannot request to change the risk category a URL receives (**high risk**, **medium risk**, or **low risk**), or to URLs categorized as **insufficient content** or **newly-registered domains**.

- [Make a Change Request Online](#)
- [Make a Bulk Change Request](#)
- [Make a Change Request From the Firewall](#)

Make a Change Request Online

Visit Palo Alto Networks URL Filtering [Test A Site](#) to make a change request online.

STEP 1 | Go to [Test A Site](#).

You do not need to log in to submit a change request, though you will need to provide your email as part of the change request form. If you decide not to log in, you'll need to take a CAPTCHA test to confirm that you're a human being (log in to avoid the CAPTCHA test).

STEP 2 | Enter a URL to check it's categories:




STEP 3 | Review the URL categories, and if you don't think that they're accurate, select **Request Change**.



STEP 4 | Continue to populate and submit the change request form.

Include at least one (and up to two) new category suggestions, and leave an (optional) comment to tell us more about your suggestion.



Make a Bulk Change Request

You can also use [Test A Site](#) to make a bulk change request, where you want to submit change requests for multiple URLs at a single time.

STEP 1 | Go to [Test A Site](#).

You don't need to log in to make a change request; however, you'll need to provide your email as part of completing the change request form. If you decide not to log in, you'll need to take a CAPTCHA test to confirm that you're a human being (log in to avoid the CAPTCHA test).

STEP 2 | Choose the option to submit a bulk change request:



STEP 3 | Complete and submit the bulk change request form.

Change Multiple Sites

File Format # Multiple Category to Single Category

Description: The multiple category submission should be used if your change requests are for two or more categories. For example, if your request is to have three sites changed to the "Games" category and two also changed to the "Tracking" category, then you'll need to use this upload method.

- The uploaded file must be in CSV format
- It must not exceed 1000 entries
- It cannot be larger than 2MB in size
- It should have one change request per line with the format: `URL, [suggested category], [old (and new) comment]`
- If there are comments in your URL, or any (any) comment, please quote them with double quotation marks.

CSV File Example:

```
www.palmonlinevents.com,business-and-economy,"this is my comment"
www.palmonlinevents.com,business-and-economy,"this is my comment"
www.palmonlinevents.com,business-and-economy,"this is my comment"
```

Here's a downloadable list of possible suggested categories:

URLs List upload: No file chosen


Comment:

Your Email:

☐ Receive Email Notification?

Make a Change Request from the Firewall

You can also submit a URL category change request directly from the firewall. In the URL Filtering logs, the details for each log entry include an option to **Request Categorization Change**(Monitor > Logs > URL Filtering).



From here you can complete the request form, and submit it.

Request Categorization Change

URL:

Log Category:

Current Category:

Suggested Category:

Email: ☐ educational-institutions

Confirm Email: ☐ entertainment-and-arts

Comments: ☐ extremism

☐ financial-services

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](#)

Problems Activating PAN-DB

Use the following workflow to 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 | Check the [PAN-DB cloud connection status](#).

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 :                  serverlist.urlcloud.paloaltonetworks.com
Cloud connection :                      connected
Cloud mode :                           public
URL database version - device :         20200624.20296
URL database version - cloud :          20200624.20296 ( last update time
2020/06/24 12:39:19 )
URL database status :                   good
URL protocol version - device :         pan/2.0.0
URL protocol version - cloud :          pan/2.0.0
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 : 0000.00.00.000
URL protocol version - device : pan/0.0.2
```

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 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 serverlist.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 serverlist.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:

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.

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>
```

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>
```

PAN-DB Private Cloud

The PAN-DB private cloud is an on-premise solution for organizations that restrict the usage of cloud services. With this on-premise solution, you can deploy one or more M-600 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-600 appliance(s) to have direct internet access or keep it completely offline. Because the M-600 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-600 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-600 appliance, the appliance automatically generates a new set of certificates to authenticate the firewalls.

- [M-600 Appliance for PAN-DB Private Cloud](#)
- [Set Up the PAN-DB Private Cloud](#)

M-600 Appliance for PAN-DB Private Cloud

To deploy a PAN-DB private cloud, you need one or more M-600 appliances. The [M-600 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-600 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-200 appliance cannot be deployed as a PAN-DB private cloud.

The M-600 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-600 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.

Table 4: 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 PAN-DB public cloud updates the URL categories malware and phishing every five minutes. The firewall checks for critical updates whenever it queries the cloud servers for URL lookups.	Content updates and full URL database updates are available once a day during the work week.
URL Categorization Requests	Submit URL categorization change requests using the following options: • 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-600 appliance(s) in the PAN-DB private cloud. If there is no match for the URL, the PAN-DB private cloud sends a category <i>unknown</i> response to the firewall; the request is not sent to the public cloud unless you have configured the M-600 appliance to access the PAN-DB public cloud. If the M-600 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.

Set Up the PAN-DB Private Cloud

To deploy one or more M-600 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 Authentication with Custom Certificates on the PAN-DB Private Cloud](#)

Configure the PAN-DB Private Cloud

STEP 1 | Rack mount the M-600 appliance.

Refer to the [M-600 Hardware Reference Guide](#) for instructions.

STEP 2 | Register the M-600 appliance.

For instructions on registering the M-600 appliance, see [Register the Firewall](#).

STEP 3 | Perform Initial Configuration of the M-600 Appliance.



The M-600 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-600 appliance in one of the following ways:
 - Attach a serial cable from a computer to the Console port on the M-600 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-600 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 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 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

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-600
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-600
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
```

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:

1. To check whether a new version is published use:

```
request pan-url-db upgrade check
```

2. To check the version that is currently installed on your server use:

```
request pan-url-db upgrade info
```

3. To download and install the latest version:

- **request pan-url-db upgrade download latest**
 - **request pan-url-db upgrade install <version latest | file>**
4. To schedule the M-600 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>**

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:
 1. **configure**
 2. **set mgt-config users <username> permissions role-based <superreader | superuser> yes**
 3. **set mgt-config users <username> password**
 4. **Enter password:xxxxxx**
 5. **Confirm password:xxxxxx**
 6. **commit**

- To set up an administrative user with RADIUS authentication:

1. 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>
```

2. 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>
```

3. Attach the authentication-profile to the user.

```
set mgt-config users <username> authentication-  
profile <auth_profile_name>
```

4. 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.

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.

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
```

Configure Authentication with Custom Certificates on the PAN-DB Private Cloud

By default, a PAN-DB server uses predefined certificates for mutual authentication to establish the SSL connections used for management access and inter-device communication. However, you can configure authentication using custom certificates instead. Custom certificates allow you to establish a unique chain of trust to ensure mutual authentication between your PAN-DB server and firewalls. In the case of a PAN-DB private cloud, the firewall acts as the client and the PAN-DB server acts as the server.

STEP 1 | Obtain key pairs and certificate authority (CA) certificates for the PAN-DB server and firewall.

STEP 2 | Import the CA certificate to validate the certificate on the firewall.

1. Log in to the CLI on the PAN-DB server and enter configuration mode.

```
admin@M-600> configure
```

2. Use TFTP or SCP to import the CA certificate.

```
admin@M-600# {tftp | scp} import certificate from <value> file <value>
remote-port <1-65535> source-ip <ip/netmask> certificate-name <value>
passphrase <value> format {pkcs12 | pem}
```

STEP 3 | Use TFTP or SCP to import the key pair that contains the server certificate and private key for the PAN-DB M-600 appliance.

```
admin@M-600# {tftp | scp} import keypair from <value> file <value>
remote-port <1-65535> source-ip <ip/netmask> certificate-name <value>
passphrase <value> format {pkcs12 | pem}
```

STEP 4 | Configure a certificate profile that includes the root CA and intermediate CA. This certificate profile defines the device authentication between the PAN-DB server and the firewall.

1. In the CLI of the PAN-DB server, enter configuration mode.

```
admin@M-600> configure
```

2. Name the certificate profile.

```
admin@M-600# set shared certificate-profile <name>
```

3. (Optional) Set the user domain.

```
admin@M-600# set shared certificate-profile <name> domain <value>
```

4. Configure the CA.



Default-ocsp-url and ocsp-verify-cert are optional parameters.

```
admin@M-600# set shared certificate-profile <name> CA <name>
```

```
admin@M-600# set shared certificate-profile <name> CA <name> [default-ocsp-url <value>]
```

```
admin@M-600# set shared certificate-profile <name> CA <name> [ocsp-verify-cert <value>]
```

STEP 5 | Configure an SSL/TLS profile for the PAN-DB M-600 appliance. This profile defines the certificate and protocol range that PAN-DB and client devices use for SSL/TLS services.

1. Identify the SSL/TLS profile.

```
admin@M-600# set shared ssl-tls-service-profile <name>
```

2. Select the certificate.

```
admin@M-600# set shared ssl-tls-service-profile <name> certificate <value>
```

3. Define the SSL/TLS range.



PAN-OS 8.0 and later releases support TLS 1.2 and later TLS versions only. You must set the max version to TLS 1.2 or max.

```
admin@M-600# set shared ssl-tls-service-profile <name> protocol-settings min-version {tls1-0 | tls1-1 | tls1-2
```

```
admin@M-600# set shared ssl-tls-service-profile <name> protocol-settings max-version {tls1-0 | tls1-1 | tls1-2 | max
```

STEP 6 | Configure secure server communication on PAN-DB.

1. Set the SSL/TLS profile. This SSL/TLS service profile applies to all SSL connections between PAN-DB and firewalls.

```
admin@M-600# set deviceconfig setting management secure-conn-server ssl-  
tls-service-profile <ssltls-profile>
```

2. Set the certificate profile.

```
admin@M-600# set deviceconfig setting management secure-conn-server  
certificate-profile <certificate-profile>
```

3. Set the disconnect wait time in number of minutes that PAN-DB should wait before breaking and reestablishing the connection with its firewall (range is 0 to 44,640).

```
admin@M-600# set deviceconfig setting management secure-conn-server  
disconnect-wait-time <0-44640>
```

STEP 7 | Import the CA certificate to validate the certificate for the PAN-DB M-600 appliance.

1. Log in to the firewall web interface.
2. [Import the CA certificate](#).

STEP 8 | Configure a local or a SCEP certificate for the firewall.

1. If you are a local certificate, then [import the key pair for the firewall](#).
2. If you are a SCEP certificate for the firewall, [configure a SCEP profile](#).

STEP 9 | Configure the certificate profile for the firewall. You can configure this on each firewall individually or you can push the configuration from Panorama to the firewalls as part of a template.

1. Select **Device > Certificate Management > Certificate Profile** for firewalls or **Panorama > Certificate Management > Certificate Profile** for Panorama.
2. [Configure a Certificate Profile](#).

STEP 10 | Deploy custom certificates on each firewall. You can either deploy certificates centrally from Panorama or configure them manually on each firewall.

1. Log in to the firewall web interface.
2. Select **Device > Setup > Management** for a firewall or **Panorama > Setup > Management** for Panorama and **Edit** the Secure Communication
3. Select the **Certificate Type**, **Certificate**, and **Certificate Profile** from the respective drop-downs.
4. In the Customize Communication settings, select **PAN-DB Communication**.
5. Click **OK**.
6. **Commit** your changes.

After committing your changes, the firewalls do not terminate their current sessions with the PAN-DB server until after the **Disconnect Wait Time**. The disconnect wait time begins counting down after you enforce the use of custom certificates in the next step.

STEP 11 | After deploying custom certificates on all firewalls, enforce custom certificate authentication.

1. Log in to the CLI on the PAN-DB server and enter configuration mode.

```
admin@M-600> configure
```

2. Enforce the use of custom certificates.

```
admin@M-600# set deviceconfig setting management secure-conn-server  
disable-pre-defined-cert yes
```

After committing this change, the disconnect wait time begins counting down (if you configured setting on PAN-DB). When the wait time ends, PAN-DB and its firewall connect using only the configured certificates.

STEP 12 | You have two choices when adding new firewalls or Panorama to your PAN-DB private cloud deployment.

- If you did not enable **Custom Certificates Only** then you can add a new firewall to the PAN-DB private cloud and then deploy the custom certificate as described above.
- If you enabled **Custom Certificates Only** on the PAN-DB private cloud, then you can must deploy the custom certificates on the firewalls before connecting them to the PAN-DB private cloud.

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 model. Select two or more product models and click **Compare Now** to view QoS feature support for each model (for example, you can check if your firewall model 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-5200 Series, and PA-3200 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.




Figure 7: QoS Traffic Flow

The QoS configuration options allow you to control the traffic flow and define it at different points in the flow. The figure [QoS Traffic Flow](#) indicates where the configurable options define the traffic flow. 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.

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.

Because QoS is enforced on traffic as it egresses the firewall, your QoS policy rule is applied to traffic after the firewall has enforced all other security policy rules, including Network Address Translation (NAT) rules. If you want to apply QoS treatment to traffic based on source, make sure to specify the post-NAT source address in a QoS policy rule (do not use the pre-NAT source address).

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 to eight QoS classes, as well as the total bandwidth allotted 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, see Step [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 [Add a QoS profile rule](#)). 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, see Step [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, see Step [Enable QoS on a physical interface](#).

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.



- The egress interface for Alice's download traffic is Ethernet 1/2. To prioritize or limit her download traffic, Alice enables QoS on Ethernet 1/2.
- The egress interface for Alice's upload traffic is Ethernet 1/1. To prioritize or limit her upload traffic, Alice enables QoS on Ethernet 1/1.

Because QoS is enforced on traffic as it egresses the firewall, your QoS policy rule is applied to traffic after the firewall has enforced all other security policy rules, including Network Address Translation (NAT) rules. If you want to apply QoS treatment to traffic based on source, you must specify the post-NAT source address in a QoS policy rule (do not use the pre-NAT source address).

Learn more about 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.

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:



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.



Because QoS is enforced on traffic as it egresses the firewall, your QoS policy rule is applied to traffic after the firewall has enforced all other security policy rules, including Network Address Translation (NAT) rules. If you want to apply QoS treatment to traffic based on source, you must specify the post-NAT source address in a QoS policy rule (do not use the pre-NAT source address).

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**.

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:
 1. **Add** a class to the QoS Profile.
 2. Select the **Priority** for the class: real-time, high, medium, or low.

3. 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.

Class	Priority	Egress Max	Egress Guaranteed
class2	medium	50	2
class4	high	3000	0
class1	medium	3000	0
class3	medium	3000	0
class5	medium	3000	0
class6	medium	3000	0
class7	medium	3000	0
class8	medium	3000	0

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 firewall model 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.



1. (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.
2. Click **OK**.


STEP 6 | Commit your changes.

Click **Commit**.

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.

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.

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**.

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.

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:
 1. Click **Add** to add a class to the QoS Profile.
 2. Select the **Priority** for the class.
 3. Enter an **Egress Max** for a class to set the overall bandwidth limit for that individual class.
 4. Enter an **Egress Guaranteed** for the class to set the guaranteed bandwidth for that individual class.
6. Click **OK** to save the QoS profile.

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.


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:
 1. 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).



2. 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-vsyz 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.

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** 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.

-
- 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.

- 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.
- 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.

- 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](#).
- 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.
- Select a **Priority** for the class of traffic, such as **high**.
- 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).

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.

- Select **Policies > Security** and **Add** or modify a security policy.
- Select **Actions** and in the **QoS Marking** drop-down, choose **Follow-Client-to-Server-Flow**.
- 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 | Commit the configuration.

Commit your changes.

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.

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:



Class	Priority	Egress Max	Egress Guaranteed
class1	medium	1000	50


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.

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 prior step). 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](#).):



Source Zone	Source Address	Source User
Any	Any	companynetwork/CEO

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):

Name	Zone	Address	User	Zone	Address	Application	Service	Class
1. Policy	any	any	any	any	any	google-maps	any	1
2. Policy	any	any	any	any	any	any	any	2
3. Policy	any	any	any	any	any	any	any	3
4. Policy (CEO traffic)	any	any	any	any	any	any	any	1

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:

QoS Interface

Physical Interface: Clear Text Traffic Tunnelled Traffic

Interface Name: ethernet1/2

Egress Max (Mbps): 1000

☒ Turn on QoS feature on this interface

Default Profile:

Clear Text: CEO_traffic

Tunnel Interface: None


OK Cancel

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.

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:

Name	Guaranteed Egress (Mbps)	Maximum Egress (Mbps)	Profile	Enabled	Statistics
ethernet1/2		1,000.000		☒	Statistics
Tunnelled Traffic	0.000	0.000			
Clear Text Traffic	0.000	0.000	CEO_traffic		
ethernet1/28		0.000		☒	Statistics
Tunnelled Traffic	0.000	0.000			
Clear Text Traffic	0.000	0.000	Lmt Facebook apps		
Facebook Apps (ethernet1/19 - any)			Lmt Facebook apps		

STEP 5 | 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.

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.




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**:

		Source			Destination			
Name	Zone	Address	User	Zone	Address	Application	Class	
1 Video	any	any	any	any	any	google-video http-video youtube	1	
2 HTTPS	any	any	compared...	any	any	web-browsing	2	
4 Voice-Video	any	any	any	any	any	voip-video-ss	2	

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.

The screenshot shows the 'QoS Interface' configuration window. The 'Physical Interface' tab is selected. The 'Interface Name' field is set to 'ethernet1/2' and is circled in red. Below it, 'Egress Max (Mbps)' is set to '1000'. The checkbox 'Turn on QoS feature on this interface' is checked. Under 'Default Profile', 'Clear Text' is set to 'ensure voice-video traffic' and 'Tunnel Interface' is set to 'None'. 'OK' and 'Cancel' buttons are at the bottom right.

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).

The screenshot shows the 'QoS Interface' configuration window for a second interface. The 'Physical Interface' tab is selected. The 'Interface Name' field is set to 'ethernet1/1' and is circled in red. Below it, 'Egress Max (Mbps)' is set to '1000'. The checkbox 'Turn on QoS feature on this interface' is checked. Under 'Default Profile', 'Clear Text' is set to 'ensure voice-video traffic' and 'Tunnel Interface' is set to 'None'. 'OK' and 'Cancel' buttons are at the bottom right.

STEP 4 | The admin selects **Network > QoS** to confirm that QoS is enabled for both incoming and outgoing voice and video traffic:

Name	Guaranteed Egress (Mbps)	Maximum Egress (Mbps)	Profile	Enabled	
 Ethernet1/1		1,000,000		☒	Statistics
 Tunnelled Traffic	0.000	0.000			
 Clear Text Traffic	250.000	0.000	<u>ensure voice-video traffic</u>		
 Ethernet1/2		1,000,000		☒	Statistics
 Tunnelled Traffic	0.000	0.000			
 Clear Text Traffic	250.000	0.000	<u>ensure voice-video traffic</u>		

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.

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 8: VPN Deployments

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 9: 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. If you configure any proxy IDs, the proxy ID is counted toward any IPSec tunnel capacity.

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.

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.

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

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.

Table 5: 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

ESP	AH
• 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
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.
 - 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 172.16.0.0/16 and its destination address to 192.16.0.0/16. Thus, the narrowing down accommodates the addresses of gateway A and the traffic selectors of the two gateways are in agreement.
- If gateway B (configured with source IP address 0.0.0.0) is the Initiator instead of the Responder, gateway A will respond with its more specific IP addresses, and gateway B will narrow down its addresses to reach agreement.

Hash and URL Certificate Exchange

IKEv2 supports Hash and URL Certificate Exchange, which is used during an IKEv2 negotiation of an SA. You store the certificate on an HTTP server, which is specified by a URL. The peer fetches the certificate from the server based on receiving the URL to the server. The hash is used to check whether the content of the certificate is valid or not. Thus, the two peers exchange certificates with the HTTP CA rather than with each other.

The hash part of Hash and URL reduces the message size and thus Hash and URL is a way to reduce the likelihood of packet fragmentation during IKE negotiation. The peer receives the certificate and hash that it expects, and thus IKE Phase 1 has validated the peer. Reducing fragmentation occurrences helps protect against DoS attacks.

You can enable the Hash and URL certificate exchange when configuring an IKE gateway by selecting **HTTP Certificate Exchange** and entering the **Certificate URL**. The peer must also use Hash and URL certificate exchange in order for the exchange to be successful. If the peer cannot use Hash and URL, X.509 certificates are exchanged similarly to how they are exchanged in IKEv1.

If you enable the Hash and URL certificate exchange, you must export your certificate to the certificate server if it is not already there. When you export the certificate, the file format should be **Binary Encoded Certificate (DER)**. See [Export a Certificate for a Peer to Access Using Hash and URL](#).

SA Key Lifetime and Re-Authentication Interval

In IKEv2, two IKE crypto profile values, **Key Lifetime** and **IKEv2 Authentication Multiple**, control the establishment of IKEv2 IKE SAs. The key lifetime is the length of time that a negotiated IKE SA key is effective. Before the key lifetime expires, the SA must be re-keyed; otherwise, upon expiration, the SA must begin a new IKEv2 IKE SA re-key. The default value is 8 hours.

The re-authentication interval is derived by multiplying the **Key Lifetime** by the **IKEv2 Authentication Multiple**. The authentication multiple defaults to 0, which disables the re-authentication feature.

The range of the authentication multiple is 0-50. So, if you were to configure an authentication multiple of 20, for example, the system would perform re-authentication every 20 re-keys, which is every 160 hours. That means the gateway could perform Child SA creation for 160 hours before the gateway must re-authenticate with IKE to recreate the IKE SA from scratch.

In IKEv2, the Initiator and Responder gateways have their own key lifetime value, and the gateway with the shorter key lifetime is the one that will request that the SA be re-keyed.

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.



If your VPN traffic is passing through (not originating or terminating on) a PA-7000 Series or PA-5200 Series firewall, configure bi-directional Security policy rules to allow the ESP or AH traffic in both directions.

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 pre-shared 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.

STEP 1 | Define the IKE Gateway.

1. Select **Network > Network Profiles > IKE Gateways**, Add a gateway, and enter the gateway **Name** (**General** tab).
2. Set the **Version** to **IKEv1 only mode**, **IKEv2 only mode**, or **IKEv2 preferred mode**. The IKE gateway begins its negotiation with its peer in the mode you specify 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** you select also determines which options are available for you to configure on the **Advanced Options** tab.

STEP 2 | Establish the local endpoint of the tunnel (gateway).

1. Select the **Address Type**: **IPv4** or **IPv6**.
2. Select the physical, outgoing **Interface** on the firewall where the local gateway resides.

3. From the **Local IP Address** list, select the IP address that the VPN connection will use as the endpoint; 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).

For **Peer IP Address Type**, select one of the following and enter the corresponding information for the peer:

- **IP**—Enter a **Peer Address** that is either an IPv4 or IPv6 address or enter an address object that is an IPv4 or IPv6 address.
- **FQDN**—Enter a **Peer Address** that is an FQDN string or an address object that uses an FQDN string. If the FQDN or FQDN address object resolves to more than one IP address, the firewall selects the preferred address from the set of addresses that match the Address Type (IPv4 or IPv6) of the IKE gateway as follows:
 - If no IKE security association (SA) is negotiated, the preferred address is the IP address with the smallest value.
 - If the IKE gateway uses an address that is in the set of returned addresses, the firewall selects that address (whether or not it is the smallest address in the set).
 - If the IKE gateway uses an address that isn't in the set of returned addresses, the firewall selects a new address, and it is the smallest address in the set.
- **Dynamic**—Select **Dynamic** if the peer IP address or FQDN value is unknown so that the peer will initiate the negotiation.



Using an FQDN or FQDN address object reduces issues in environments where the peer is subject to dynamic IP address changes (and would otherwise require you to reconfigure this IKE gateway peer address).

STEP 4 | Specify how to authenticate the peer.

Select the **Authentication** method: **Pre-Shared Key** or **Certificate**. If you choose a pre-shared key, proceed to the next step. If you choose a certificate, skip ahead to Step 6, Configure certificate-based authentication.

STEP 5 | Configure a pre-shared key.

1. Enter a **Pre-shared Key**, which is the security key 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)**, and **User FQDN (email address)**. Local identification defines the format and identification of the local gateway. If you do not specify a value, the local IP address is used as the local identification value.
3. For **Peer Identification**, choose from the following types and enter a value that you determine: **FQDN (hostname)**, **IP address**, **KEYID (binary format ID string in HEX)**, and **User FQDN (email address)**. Peer identification defines the format and identification of the peer gateway. If you do not specify a value, the peer IP address is used as the peer identification value.
4. Proceed to Step 7 (Configure advanced options for the gateway).

STEP 6 | Configure certificate-based authentication.

Perform the remaining steps in this procedure if you selected **Certificate** as the method of authenticating the peer gateway at the opposite end of the tunnel.

1. Select a **Local Certificate**—one that is already on the firewall, **Import** a certificate, or **Generate** a new certificate.
 - If you need to **Import** a certificate, then first [Import a Certificate for IKEv2 Gateway Authentication](#) and then return to this task.
 - If you want to **Generate** a new certificate, then first [generate a certificate on the firewall](#) and then return to this task.
2. (Optional) Enable (select) the **HTTP Certificate Exchange** 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—**Distinguished Name (Subject)**, **FQDN (hostname)**, **IP address**, or **User FQDN (email address)**—and then enter the value. Local identification defines the format and identification of the local gateway.
4. Select the **Peer Identification** type—**Distinguished Name (Subject)**, **FQDN (hostname)**, **IP address**, or **User FQDN (email address)**—and then enter the value. Peer identification defines the format and identification of the peer gateway.
5. Specify the type of **Peer ID Check**:
 - **Exact**—Ensures that the local setting and peer IKE ID payload match exactly.
 - **Wildcard**—Allows the peer identification to match as long as every character before the wildcard (*) matches. The characters after the wildcard need not match.
6. (Optional) **Permit peer identification and certificate payload identification mismatch** 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**. A certificate profile contains information about how to authenticate the peer gateway.
8. (Optional) **Enable strict validation of peer's extended key use** to strictly control how the key can be used.

STEP 7 | Configure advanced options for the gateway.

1. (Optional) **Enable Passive Mode** in the Common Options (**Advanced Options**) to specify that the firewall only respond to IKE connection requests and never initiate them.
2. If you have a device performing NAT between the gateways, **Enable NAT Traversal** to use UDP encapsulation on IKE and UDP protocols, which enables them to pass through intermediate NAT devices.
3. If you configured **IKEv1 only mode** in Step 1, then, on the IKEv1 tab:
 - Select the **Exchange Mode**: **auto**, **aggressive**, or **main**. When you set a firewall to use **auto** exchange mode, it can accept both **main** mode and **aggressive** mode negotiation requests; however, when possible, it will initiate exchanges in **main** mode.



If you do not set the exchange mode to auto, then 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 the **IKE Crypto Profile** list. If needed, you can [Define IKE Crypto Profiles](#).
 - (Only when using certificate-based authentication and when 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 to 100 seconds). For **Retry**, define the time to delay (range is 2 to 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.
4. If you configured **IKEv2 only mode** or **IKEv2 preferred mode** in Step 1, then on the IKEv2 tab:

- Select an **IKE Crypto Profile**, 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](#).
- (Optional) Enable **Strict Cookie Validation** [Cookie Activation Threshold and Strict Cookie Validation](#).
- (Optional) 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 as many as 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 | Click **OK** and **Commit** your changes.

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**.

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**.

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 | Continue to the next task.

Step [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.

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 | Commit your changes.

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](#).

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 0-65,535; default is 500).
2. Click **OK**.

STEP 2 | Commit your changes.

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.

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**).

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.

- [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.

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 (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 menus.

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**.



As a best practice, choose the strongest authentication and encryption algorithms the peer can support. For the authentication algorithm, use SHA-256 or higher (SHA-384 or higher preferred for long-lived transactions). Do not use SHA-1 or MD5. For the encryption algorithm, use AES; DES and 3DES are weak and vulnerable.

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; default is 0) that is multiplied by the **Key Lifetime** to determine the authentication count. The default value of 0 disables the re-authentication feature.

STEP 4 | Commit 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.

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.



As a best practice, select ESP (Encapsulating Security payload) over AH (Authentication Header) because ESP offers both confidentiality and authentication for the connection whereas AH offers only authentication.

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**.



As a best practice, choose the strongest authentication and encryption algorithms the peer can support. For the authentication algorithm, use SHA-256 or higher (SHA-384 or higher preferred for long-lived transactions). Do not use SHA-1, MD5, or none. For the encryption algorithm, use AES; DES and 3DES are weak and vulnerable.

- Authentication—**sha512**, **sha384**, **sha256**, **sha1**, **md5**.

STEP 2 | Select the DH Group to use for the IPsec SA negotiations in IKE phase 2.

From **DH Group**, select the key strength you want to use: **group1**, **group2**, **group5**, **group14**, **group19**, or **group20**. For highest security, choose the group with the highest number.

If you don't want to renew the key that the firewall creates during IKE phase 1, select **no-pfs** (no perfect forward secrecy); the firewall reuses the current key for the IPSec security association (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 **Lifeseize** or volume of data after which the keys must be renegotiated.

STEP 4 | Commit your IPSec profile.

Click **OK** and click **Commit**.

STEP 5 | Attach the IPSec Profile to an IPSec tunnel configuration.

See [Set up key exchange](#).

Set Up an IPSec Tunnel

The IPSec tunnel configuration allows you to authenticate and/or encrypt the data (IP packet) as it traverses the tunnel.

If you are setting up the 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 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 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.

STEP 1 | Select **Network > IPSec Tunnels** and then **Add** a new tunnel configuration.

STEP 2 | On the **General** tab, enter a **Name** for the tunnel.

STEP 3 | Select the **Tunnel interface** on which 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** list to define the zone as follows:

Use your trust zone as the termination point for the tunnel—Select the zone. 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.

Or:

Create a separate zone for VPN tunnel termination (Recommended)—Select **New Zone**, define a **Name** for the new zone (for example vpn-corp), and click **OK**.

1. For **Virtual Router**, select **default**.
2. (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.
3. Click **OK**.

STEP 4 | (Optional) Enable IPv6 on the tunnel interface.

1. Select the IPv6 tab on **Network > Interfaces > Tunnel > IPv6**.
2. Select **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.

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.
 1. 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.
 2. Select **Anycast** to include routing through the nearest node.

STEP 5 | Set up key exchange.

On the **General** tab, 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 **Local 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 and enter a **Key** and then **Confirm Key**.
5. For ESP, select the **Authentication** method 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 **Remote SPI** for the remote peer.
7. Enter the **Remote Address**, the IP address of the remote peer.

STEP 6 | Protect against a replay attack.

A replay attack occurs when a packet is maliciously intercepted and retransmitted by the interceptor.

On the **General** tab, select **Show Advanced Options** and 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 | (Optional) Select **Add GRE Encapsulation** to enable GRE over IPSec.

Add GRE encapsulation in cases where the remote endpoint requires traffic to be encapsulated within a GRE tunnel before IPSec encrypts the traffic. For example, some implementations require multicast traffic to be encapsulated before IPSec encrypts it. Add GRE Encapsulation when the GRE packet encapsulated in IPSec has the same source IP address and destination IP address as the encapsulating IPSec tunnel.

STEP 9 | 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. Select **Tunnel Monitor**.
2. Specify a **Destination IP** address on the other side of the tunnel to determine if the tunnel is working properly.
3. Select a **Profile** to determine the action upon tunnel failure. To create a new profile, see [Define a Tunnel Monitoring Profile](#).

STEP 10 | Create a Proxy ID to identify the VPN peers.

This step is required only if the VPN peer uses policy-based VPN.

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**:
 - **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 11 | Commit 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.

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** to take 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 (sec)** and **Threshold** to trigger the specified action.

- **Threshold** specifies the number of heartbeats to wait before taking the specified action (range is 2-100; default is 5).
- **Interval (sec)** specifies the time (in seconds) between heartbeats (range is 2-10; default is 3).

STEP 5 | Attach the monitoring profile to the IPsec Tunnel configuration. See [Enable Tunnel Monitoring](#).

View the Status of the Tunnels

The status of the tunnel informs you about whether or not valid IKE phase-1 and phase-2 SAs have been established, and whether the tunnel interface is up and available for passing traffic.

Because the tunnel interface is a logical interface, it cannot indicate a physical link status. Therefore, you must enable tunnel monitoring so that the tunnel interface can verify connectivity to an IP address and determine if the path is still usable. If the IP address is unreachable, the firewall will either wait for the tunnel to recover or failover. When a failover occurs, the existing tunnel is torn down and routing changes are triggered to set up a new tunnel and redirect traffic.

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 to make troubleshooting easier.

- 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 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

You can [Refresh or Restart an IKE Gateway or IPSec Tunnel](#). 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 <code>show</code> command in the CLI (after an initial <code>show</code> 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 <code>clear, test, show</code> command sequence in the CLI.

Phase	Refresh	Restart
IPSec Tunnel (IKE Phase 2)	Updates the onscreen statistics for the selected IPSec tunnel. Equivalent to issuing a second <code>show</code> command in the CLI (after an initial <code>show</code> 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

Keep in mind that the result of restarting an IKE gateway depends on whether it is IKEv1 or IKEv2. See [Refresh and Restart Behaviors](#) for an IKE gateway (IKEv1 and IKEv2) and for an 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.

- 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

Perform this task to 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 whether the Security Association displays. If it doesn't, 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 2 is set up:

```
show vpn ipsec-sa tunnel <tunnel_name>
```

In the output, check whether the Security Association displays. If it doesn't, 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
-----
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 6: 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.

If error is this:	Try this:
<code>pfs group mismatched:my: 2peer: 0</code> 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..


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.



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**.
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**, 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

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** to define the zone as follows:
 - To use your trust zone as the termination point for the tunnel, select the zone.
 - **(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.10
- **Security Zone**—vpn_tun
- **Virtual Router**—default
- **IPv4**—172.19.9.2/24

The configuration for VPN Peer B is:

- **Interface**—tunnel.11
- **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.10

The configuration for VPN Peer B is:

- **Destination**—172.19.9.0/24
- **Interface**—tunnel.11

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.

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.10
- **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.11
 - **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 | Commit any pending configuration changes.

Click **Commit**.

STEP 9 | Test VPN Connectivity.

See also [View the Status of the Tunnels](#).

Site-to-Site VPN with OSPF

In this example, each site uses OSPF for dynamic routing of traffic. The tunnel IP address on each VPN peer is statically assigned and serves as the next hop for routing traffic between the two sites.



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** list.
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** list, 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 | 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 **Security Zone** to define the zone as follows:
 - To use your trust zone as the termination point for the tunnel, select the zone.
 - (**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.

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 example 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.

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.

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**


```

admin@FW-A> show routing protocol ospf neighbor

Options: 0x80:reserved, O:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
N/P:NSSA option, MC:multicast, E:AS external LSA capability, T:TOS capability
*****
virtual router:          vr1
neighbor address:        2.1.1.140
local address binding:    0.0.0.0
type:                    dynamic
status:                  full
neighbor router ID:       192.168.100.140
area id:                 0.0.0.0
neighbor priority:        1
lifetime remain:         39
messages pending:         0
LSA request pending:     0
options:                 0x42: O E
hello suppressed:        no

admin@FW-B> show routing protocol ospf neighbor

Options: 0x80:reserved, O:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
N/P:NSSA option, MC:multicast, E:AS external LSA capability, T:TOS capability
*****
virtual router:          vr1
neighbor address:        2.1.1.141
local address binding:    0.0.0.0
type:                    dynamic
status:                  full
neighbor router ID:       192.168.100.141
area id:                 0.0.0.0
neighbor priority:        1
lifetime remain:         39
messages pending:         0
LSA request pending:     0
options:                 0x42: O E
hello suppressed:        no

```

- **show routing route type ospf**

```

admin@FW-A> show routing route type ospf

flags: A:active, ?:loose, C:connect, H:host, S:static, ~:internal, R:rip, O:ospf, B:bgp,
O1:ospf intra-area, Oo:ospf inter-area, Ol:ospf ext-type-1, O2:ospf ext-type-2

VIRTUAL ROUTER: vr1 (id 1)
*****
destination      nexthop      metric flags    age  interface    next-AS
2.1.1.0/24        0.0.0.0      10  O1           6760 tunnel.4#
172.16.101.0/24   0.0.0.0      10  O1           6854 ethernet1/1
192.168.1.0/24    2.1.1.140    20  A Oo         6754 tunnel.4#
total routes shown: 3

admin@FW-B> show routing route type ospf

flags: A:active, C:connect, H:host, S:static, R:rip, O:ospf,
O1:ospf intra-area, Oo:ospf inter-area, Ol:ospf ext-type-1, O2:ospf ext-type-2

VIRTUAL ROUTER: vr1 (id 1)
*****
destination      nexthop      metric flags    age  interface    next-AS
2.1.1.0/24        0.0.0.0      10  O1           20033 tunnel.4#
172.16.101.0/24   2.1.1.141    20  A Oo         6896 tunnel.4#
192.168.1.0/24    0.0.0.0      10  O1           8058 ethernet1/15
total routes shown: 3

```

STEP 9 | Test VPN Connectivity.


See [Set Up Tunnel Monitoring](#) and [View the Status of the Tunnels](#).

Site-to-Site VPN with Static and Dynamic Routing

In this example, one site uses static routes and the other site uses OSPF. When the routing protocol is not the same between the locations, the tunnel interface on each firewall must be configured with a static

IP address. Then, to allow the exchange of routing information, the firewall that participates in both the static and dynamic routing process must be configured with a Redistribution profile. Configuring the redistribution profile enables the virtual router to redistribute and filter routes between protocols—static routes, connected routes, and hosts— from the static autonomous system to the OSPF autonomous system. Without this redistribution profile, each protocol functions on its own and does not exchange any route information with other protocols running on the same virtual router.

In this example, the satellite office has static routes and all traffic destined to the 192.168.x.x network is routed to tunnel.41. The virtual router on VPN Peer B participates in both the static and the dynamic routing process and is configured with a redistribution profile in order to propagate (export) the static routes to the OSPF autonomous system.



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**.
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**, 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.

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.

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** to define the zone as follows:
 - To use your trust zone as the termination point for the tunnel, select the zone.
 - (**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

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.
 1. Select **Network > Virtual Routers**, and select the router you used above.
 2. Select **Redistribution Profiles**, and click **Add**.
 3. 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.
 4. Set **Source Type** as **static**, and click **OK**. The static route you defined in Step 6 will be used for the redistribution.
2. Inject the static routes in to the OSPF system.
 1. Select **OSPF > Export Rules** (for IPv4) or **OSPFv3 > Export Rules** (for IPv6).
 2. Click **Add**, and select the redistribution profile that you just created.
 3. 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**.
 4. 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.
 5. Click **OK**.

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.

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**

```
admin@FW-A> show routing protocol ospf neighbor

Options: 0x80:reserved, O:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
         N/P:NSSA option, MC:multicast, E:AS external LSA capability, T:TOS capability
=====
virtual router:          vr1
neighbor address:        2.1.1.140
local address binding:   0.0.0.0
type:                    dynamic
status:                  full
neighbor router ID:      192.168.100.140
area id:                 0.0.0.0
neighbor priority:       1
lifetime remain:         39
messages pending:        0
LSA request pending:     0
options:                 0x42: O E
hello suppressed:       no

admin@FW-B> show routing protocol ospf neighbor

Options: 0x80:reserved, O:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
         N/P:NSSA option, MC:multicast, E:AS external LSA capability, T:TOS capability
=====
virtual router:          vr1
neighbor address:        2.1.1.141
local address binding:   0.0.0.0
type:                    dynamic
status:                  full
neighbor router ID:      192.168.100.141
area id:                 0.0.0.0
neighbor priority:       1
lifetime remain:         39
messages pending:        0
LSA request pending:     0
options:                 0x42: O E
hello suppressed:       no
```

- **show routing route**

The following is an example of the output on each VPN peer.

```
VPN PeerA
destination      next hop      metric  flags  age      interface      next-AS
192.168.1.0/24   2.1.1.141    20      A S    -        tunnel.41      -
192.168.2.0/24   2.1.1.141    20      A S    -        tunnel.41      -
172.16.101.0/24  0.0.0.0      1       A H    -        ethernet1/1    -
2.1.1.140/24     2.1.1.141    20      A S    -        tunnel.41      -

VPN PeerB
destination      next hop      metric  flags  age      interface      next-AS
192.168.1.0/24   0.0.0.0      10      A Oo   -        ethernet1/1    -
192.168.2.0/24   0.0.0.0      10      A Oo   -        ethernet1/15   -
172.16.101.0/24  2.1.1.140    20      A H    -        tunnel.40      -
2.1.1.141/24     2.1.1.140    10      A C    -        tunnel.40      -
```

STEP 11 | Test VPN Connectivity.

See [Set Up Tunnel Monitoring](#) and [View the Status of the Tunnels](#).

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.



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 supports both IPv6 and IPv4 addressing for 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. GlobalProtect supports both IPv6 and IPv4 addressing for the tunnel interface.

For more information about portals, gateways, and satellites see [LSVPN Overview](#).

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.

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. Assign an IP address to the interface:
 - For an IPv4 address, select **IPv4** and **Add** the IP address and network mask to assign to the interface, for example 203.0.11.100/24.
 - For an IPv6 address, select **IPv6**, **Enable IPv6 on the interface**, and **Add** the IP address and network mask to assign to the interface, for example 2001:1890:12f2:11::10.1.8.160/80.
6. To save the interface configuration, click **OK**.

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**) To assign an IP address to the tunnel interface:
 - For an IPv4 address, select **IPv4** and **Add** the IP address and network mask to assign to the interface, for example 203.0.11.100/24.
 - For an IPv6 address, select **IPv6**, **Enable IPv6 on the interface**, and **Add** the IP address and network mask to assign to the interface, for example 2001:1890:12f2:11::10.1.8.160/80.
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 | Commit your changes.

Click **Commit**.

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). When using self-signed root CA certificates, 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:

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:
 1. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Generate**.
 2. Enter a **Certificate Name**.
 3. Enter the FQDN (**recommended**) or IP address of the interface where you plan to configure the gateway in the **Common Name** field.
 4. In the **Signed By** field, select the **LSVPN_CA** certificate you just created.
 5. 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**.
 6. **Generate** the certificate.
2. [Configure an SSL/TLS Service Profile](#) for the portal and each gateway:
 1. Select **Device > Certificate Management > SSL/TLS Service Profile** and click **Add**.
 2. Enter a **Name** to identify the profile and select the server **Certificate** you just created for the portal or gateway.
 3. Define the range of TLS versions (**Min Version** to **Max Version**) allowed for communicating with satellites and click **OK**.

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.
 1. Select **Device > Certificate Management > Certificates > Device Certificates**.
 2. Select the root CA certificate used to issue certificates for the LSVPN components and click **Export**.
 3. 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.
 1. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.
 2. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.
 3. **Browse** to the **Certificate File** you downloaded from the CA.
 4. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.
 5. Select the certificate you just imported on the **Device Certificates** tab to open it.
 6. Select **Trusted Root CA** and then click **OK**.
 7. **Commit** the changes.

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. (Recommended) Enable use of CRL and/or OCSP to enable certificate status verification.
5. Click **OK** to save the profile.

STEP 6 | Commit your changes.

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.

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.

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**.

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.

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 it 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.

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 [Add the user account to the local database](#).

Configure a server profile for the authentication service type:

-
- [Add a RADIUS server profile.](#)



You can use RADIUS to integrate with a [Multi-Factor Authentication](#) service.

- [Add a TACACS+ server profile.](#)
- [Add a SAML IdP server profile.](#)
- [Add a Kerberos server profile.](#)
- [Add an LDAP server profile.](#) If you use LDAP to connect to Active Directory (AD), create a separate LDAP server profile for every AD domain.

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**.

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.

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 each GlobalProtect gateway to participate in the LSVPN as follows:

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. Specify the **IP Address Type** and **IP address** for gateway access:
 - The IP address type can be **IPv4** (only), **IPv6** (only), or **IPv4 and IPv6**. Use **IPv4 and IPv6** if your network supports dual stack configurations, where IPv4 and IPv6 run at the same time.
 - The IP address must be compatible with the IP address type. For example, **172.16.1/0** for IPv4 addresses or **21DA:D3:0:2F3B** for IPv6 addresses. For dual stack configurations, enter both an IPv4 and IPv6 address.
3. Click **OK** to save changes.

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. You can specify an **IPv4** address, and **IPv6** address, or both. 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](#).

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.



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:

- [GlobalProtect Portal for LSVPN Prerequisite Tasks](#)
- [Configure the Portal](#)
- [Define the Satellite Configurations](#)

GlobalProtect Portal for LSVPN 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 [GlobalProtect Portal for LSVPN Prerequisite Tasks](#), configure the GlobalProtect portal as follows:

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. Specify the **IP Address Type** and **IP address** for satellite access to the portal:
 - The IP address type can be **IPv4** (for IPv4 traffic only), **IPv6** (for IPv6 traffic only, or **IPv4 and IPv6**. Use **IPv4 and IPv6** if your network supports dual stack configurations, where IPv4 and IPv6 run at the same time.
 - The IP address must be compatible with the IP address type. For example, **172.16.1/0** for IPv4 addresses or **21DA:D3:0:2F3B** for IPv6 addresses. For dual stack configurations, enter both an IPv4 and IPv6 address.
3. Click **OK** to save changes.

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.

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](#).

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 to 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**.

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.

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.

STEP 1 | Configure a Layer 3 Interface (see [Configure Layer 3 Interfaces](#)).

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**) To assign an IP address to the tunnel interface:
 - For an IPv4 address, select **IPv4** and **Add** the IP address and network mask to assign to the interface, for example 203.0.11.100/24.
 - For an IPv6 address, select **IPv6**, **Enable IPv6 on the interface**, and **Add** the IP address and network mask to assign to the interface, for example 2001:1890:12f2:11::10.1.8.160/80.
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:
 1. Select **Device > Certificate Management > Certificates > Device Certificates**.
 2. Select the CA certificate, and click **Export**.
 3. 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.
 1. Select **Device > Certificate Management > Certificates > Device Certificates** and click **Import**.
 2. Enter a **Certificate Name** that identifies the certificate as your client CA certificate.
 3. **Browse** to the **Certificate File** you downloaded from the CA.
 4. Select **Base64 Encoded Certificate (PEM)** as the **File Format** and then click **OK**.

-
5. Select the certificate you just imported on the **Device Certificates** tab to open it.
 6. Select **Trusted Root CA** and then click **OK**.

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 **IP Address** to use on the selected interface. You can select an **IPv4** address, an **IPv6** address, or both. Specify if you want **IPv6 preferred for portal registration**.

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 GlobalProtect Gateways for LSVPN](#).

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**.

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).

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](#)

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:

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](#).

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 **lsvpnservice**.

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:

- [Add a RADIUS server profile.](#)



You can use RADIUS to integrate with a [Multi-Factor Authentication](#) service.

- [Add a TACACS+ server profile.](#)
 - [Add a SAML IdP server profile.](#)
 - [Add a Kerberos server profile.](#)
 - [Add an LDAP server profile.](#) If you use LDAP to connect to Active Directory (AD), create a separate LDAP server profile for every AD domain.
2. [Configure an authentication profile.](#) In this example, the profile **lsvpn-sat** is used to authenticate satellites.

STEP 7 | [Configure GlobalProtect Gateways 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**—lsvpnservice
- **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.](#)

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

STEP 9 | Define the Satellite Configurations.

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.

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 [Create Interfaces and Zones for the LSVPN](#).

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.

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.


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-5200 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-5200s in an active/passive HA pair acting as the backup LSVPN gateway. The portal and gateways serve 500 PA-220s 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:

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

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 the 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 [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

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**.
 1. Name the redistribution profile **ToAllSat** and set the **Priority** to 1.
 2. Set Redistribute to **Redist**.
 3. **Add ethernet1/23** from the Interface drop-down.
 4. Click **OK**.
4. Select **BGP** on the Virtual Router to configure BGP.
 1. On **BGP > General**, select **Enable**.
 2. Enter the gateway IP address as the **Router ID (172.16.22.1)** and **1000** as the **AS Number**.
 3. In the Options section, select **Install Route**.
 4. On **BGP > Peer Group**, click **Add** a peer group with all the satellites that will connect to the gateway.
 5. 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.
 1. Name the redistribution profile **ToLSVPNGW** and set the **Priority** to 1.
 2. **Add** an **Interface ethernet1/2.1**.
 3. Click **OK**.
6. Select **BGP > General**, **Enable** BGP and configure the protocol as follows:
 1. Enter the gateway IP address as the **Router ID (172.16.22.1)** and **1000** as the **AS Number**.
 2. In the Options section, select **Install Route**.
 3. On **BGP > Peer Group**, **Add** a peer group containing all the satellites that will connect to the gateway.
 4. On **BGP > Redist Rules**, **Add** the **ToLSVPNGW** redistribution profile you created previously.
7. Click **OK**.

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.

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. Once your network interfaces have been configured, you can Export Configuration Table Data as a PDF or CSV for internal review or audits.

The following topics describe networking concepts and how to integrate Palo Alto Networks next-generation firewalls into your network.

- > [Configure Interfaces](#)
- > [Virtual Routers](#)
- > [Service Routes](#)
- > [Static Routes](#)
- > [RIP](#)
- > [OSPF](#)
- > [BGP](#)
- > [IP Multicast](#)
- > [Route Redistribution](#)
- > [GRE Tunnels](#)
- > [DHCP](#)
- > [DNS](#)
- > [Dynamic DNS Overview](#)
- > [Configure Dynamic DNS for Firewall Interfaces](#)
- > [NAT](#)
- > [NPTv6](#)
- > [NAT64](#)
- > [ECMP](#)
- > [LLDP](#)
- > [BFD](#)
- > [Session Settings and Timeouts](#)
- > [Tunnel Content Inspection](#)

Configure Interfaces

A Palo Alto Networks next-generation firewall can operate in multiple deployments at once because the deployments occur at the interface level. For example, you can configure some interfaces for Layer 3 interfaces to integrate the firewall into your dynamic routing environment, while configuring other interfaces to integrate into your Layer 2 switching network. The following topics describe each type of interface deployment and how to configure the corresponding interface types:

- [Tap Interfaces](#)
- [Virtual Wire Interfaces](#)
- [Layer 2 Interfaces](#)
- [Layer 3 Interfaces](#)
- [Configure an Aggregate Interface Group](#)
- [Use Interface Management Profiles to Restrict Access](#)

Tap Interfaces

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

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

By deploying the firewall in tap mode, you can get visibility into what applications are running on your network without having to make any changes to your network design. In addition, when in tap mode, the firewall can also identify threats on your network. Keep in mind, however, because the traffic is not running through the firewall when in tap mode it cannot take any action on the traffic, such as blocking traffic with threats or applying QoS traffic control.

To configure a tap interface and begin monitoring the applications and threats on your network:

STEP 1 | Decide which port you want to use as your tap interface and connect it to a switch configured with SPAN/RSPAN or port mirroring.

You will send your network traffic from the SPAN destination port through the firewall so you can have visibility into the applications and threats on your network.

STEP 2 | From the firewall web interface, configure the interface you want to use as your network tap.

1. Select **Network > Interfaces** and select the interface that corresponds to the port you just cabled.
2. Select **Tap** as the **Interface Type**.
3. On the **Config** tab, expand the **Security Zone** and select **New Zone**.
4. In the Zone dialog, enter a **Name** for new zone, for example TapZone, and then click **OK**.

STEP 3 | (Optional) Create any forwarding profiles you want to use.

- [Configure a Log Forwarding profile](#)
- [Configure Syslog Monitoring](#)

STEP 4 | Create [Security Profiles](#) to scan your network traffic for threats:

1. Select **Objects > Security Profiles**.

-
2. For each security profile type, **Add** a new profile and set the action to **alert**.

Because the firewall is not inline with the traffic you cannot use any block or reset actions. By setting the action to alert, you will be able to see any threats the firewall detects in the logs and ACC.

STEP 5 | Create a security policy rule to allow the traffic through the tap interface.

When creating a security policy rule for tap mode, both the source zone and destination zone must be the same.

1. Select **Policies > Security** and click **Add**.
2. In the **Source** tab, set the **Source Zone** to the TapZone you just created.
3. In the **Destination** tab, set the **Destination Zone** to the TapZone also.
4. Set the all of the rule match criteria (**Applications, User, Service, Address**) to **any**.
5. In the **Actions** tab, set the **Action Setting** to **Allow**.
6. Set **Profile Type** to **Profiles** and select each of the security profiles you created to alert you of threats.
7. Verify that **Log at Session End** is enabled.
8. Click **OK**.
9. Place the rule at the top of your rulebase.

STEP 6 | Commit the configuration.

STEP 7 | Monitor the firewall logs (**Monitor > Logs**) and the **ACC** for insight into the applications and threats on your network.


Virtual Wire Interfaces

In a virtual wire deployment, you install a firewall transparently on a network segment by binding two firewall ports (interfaces) together. The virtual wire logically connects the two interfaces; hence, the virtual wire is internal to the firewall.

Use a virtual wire deployment only when you want to seamlessly integrate a firewall into a topology and the two connected interfaces on the firewall need not do any switching or routing. For these two interfaces, the firewall is considered a *bump in the wire*.

A virtual wire deployment simplifies firewall installation and configuration because you can insert the firewall into an existing topology without assigning MAC or IP addresses to the interfaces, redesigning the network, or reconfiguring surrounding network devices. The virtual wire supports blocking or allowing traffic based on virtual LAN (VLAN) tags, in addition to supporting security policy rules, App-ID, Content-ID, User-ID, decryption, LLDP, active/passive and active/active HA, QoS, zone protection (with some exceptions), non-IP protocol protection, DoS protection, packet buffer protection, tunnel content inspection, and NAT.

Virtual Wire Deployment (No routing or switching performed by virtual wire interfaces)



Each virtual wire interface is directly connected to a Layer 2 or Layer 3 networking device or host. The virtual wire interfaces have no Layer 2 or Layer 3 addresses. When one of the virtual wire interfaces receives a frame or packet, it ignores any Layer 2 or Layer 3 addresses for switching or routing purposes, but applies your security or NAT policy rules before passing an allowed frame or packet over the virtual wire to the second interface and on to the network device connected to it.

You wouldn't use a virtual wire deployment for interfaces that need to support switching, VPN tunnels, or routing because they require a Layer 2 or Layer 3 address. A virtual wire interface doesn't use an interface management profile, which controls services such as HTTP and ping and therefore requires the interface have an IP address.

All firewalls shipped from the factory have two Ethernet ports (ports 1 and 2) preconfigured as virtual wire interfaces, and these interfaces allow all untagged traffic.



If you're using security group tags (SGTs) in a Cisco Trustsec network, it's a best practice to deploy inline firewalls in either Layer 2 or virtual wire mode. Firewalls in Layer 2 or virtual wire mode can inspect and provide threat prevention for the tagged traffic.



If you don't intend to use the preconfigured virtual wire, you must delete that configuration to prevent it from interfering with other settings you configure on the firewall. See [Set Up Network Access for External Services](#).

- [Layer 2 and Layer 3 Packets over a Virtual Wire](#)
- [Port Speeds of Virtual Wire Interfaces](#)
- [LLDP over a Virtual Wire](#)
- [Aggregated Interfaces for a Virtual Wire](#)
- [Virtual Wire Support of High Availability](#)
- [Zone Protection for a Virtual Wire Interface](#)
- [VLAN-Tagged Traffic](#)
- [Virtual Wire Subinterfaces](#)
- [Configure Virtual Wires](#)

Layer 2 and Layer 3 Packets over a Virtual Wire

A virtual wire interface will allow Layer 2 and Layer 3 packets from connected devices to pass transparently as long as the policies applied to the zone or interface allow the traffic. The virtual wire interfaces themselves don't participate in routing or switching.

For example, the firewall doesn't decrement the TTL in a traceroute packet going over the virtual link because the link is transparent and doesn't count as a hop. Packets such as Operations, Administration and Maintenance (OAM) protocol data units (PDUs), for example, don't terminate at the firewall. Thus, the virtual wire allows the firewall to maintain a transparent presence acting as a pass-through link, while still providing security, NAT, and QoS services.

In order for bridge protocol data units (BPDUs) and other Layer 2 control packets (which are typically untagged) to pass through a virtual wire, the interfaces must be attached to a virtual wire object that allows untagged traffic, and that is the default. If the virtual wire object **Tag Allowed** field is empty, the virtual wire allows untagged traffic. (Security policy rules don't apply to Layer 2 packets.)

In order for routing (Layer 3) control packets to pass through a virtual wire, you must apply a security policy rule that allows the traffic to pass through. For example, apply a security policy rule that allows an application such as BGP or OSPF.

If you want to be able to apply security policy rules to a zone for IPv6 traffic arriving at a virtual wire interface on the firewall, enable IPv6 firewalling. Otherwise, IPv6 traffic is forwarded transparently across the wire.

If you enable multicast firewalling for a virtual wire object and apply it to a virtual wire interface, the firewall inspects multicast traffic and forwards it or not, based on security policy rules. If you don't enable multicast firewalling, the firewall simply forwards multicast traffic transparently.

Fragmentation on a virtual wire occurs the same as in other interface deployment modes.

Port Speeds of Virtual Wire Interfaces

Different firewall models provide various numbers of copper and fiber optic ports, which operate at different speeds. A virtual wire can bind two Ethernet ports of the same type (both copper or both fiber optic), or bind a copper port with a fiber optic port. By default, the **Link Speed** of copper ports on the firewall is set to **auto**, which means the firewall automatically negotiates their speed and transmission mode (**Link Duplex**). When you [Configure Virtual Wires](#), you can also select a specific **Link Speed** and **Link Duplex** but the values for these settings must be the same for both ports in any single virtual wire.

LLDP over a Virtual Wire

Virtual wire interfaces can use **LLDP** to discover neighboring devices and their capabilities, and LLDP allows neighboring devices to detect the presence of the firewall in the network. LLDP makes troubleshooting easier especially on a virtual wire, where the firewall would typically go undetected by a ping or traceroute passing through the virtual wire. LLDP provides a way for other devices to detect the firewall in the network. Without LLDP, it is practically impossible for network management systems to detect the presence of a firewall through the virtual link.

Aggregated Interfaces for a Virtual Wire

You can [Configure an Aggregate Interface Group](#) of virtual wire interfaces, but virtual wires don't use LACP. If you configure LACP on devices that connect the firewall to other networks, the virtual wire will pass LACP packets transparently without performing LACP functions.



In order for aggregate interface groups to function properly, ensure all links belonging to the same LACP group on the same side of the virtual wire are assigned to the same zone.

Virtual Wire Support of High Availability

If you configure the firewall to perform path monitoring for [High Availability](#) using a virtual wire path group, the firewall attempts to resolve ARP for the configured destination IP address by sending ARP packets out both of the virtual wire interfaces. The destination IP address that you are monitoring must be on the same subnetwork as one of the devices surrounding the virtual wire.

Virtual wire interfaces support both active/passive and active/active HA. For an active/active HA deployment with a virtual wire, the scanned packets must be returned to the receiving firewall to preserve the forwarding path. Therefore, if a firewall receives a packet that belongs to the session that the peer HA firewall owns, it sends the packet across the HA3 link to the peer.

For PAN-OS 7.1 and later releases, you can configure the passive firewall in an HA pair to allow peer devices on either side of the firewall to pre-negotiate LLDP and LACP over a virtual wire before an HA failover occurs. Such a configuration for [LACP and LLDP Pre-Negotiation for Active/Passive HA](#) speeds up HA failovers.

Zone Protection for a Virtual Wire Interface

You can apply zone protection to a virtual wire interface, but because virtual wire interfaces don't perform routing, you can't apply [Packet-Based Attack Protection](#) to packets coming with a spoofed IP address, nor can you suppress ICMP TTL Expired error packets or ICMP Frag Needed packets.

By default, a virtual wire interface forwards all non-IP traffic it receives. However, you can apply a zone protection profile with [Protocol Protection](#) to block or allow certain non-IP protocol packets between security zones on a virtual wire.

VLAN-Tagged Traffic

Virtual wire interfaces by default allow all untagged traffic. You can, however, use a virtual wire to connect two interfaces and configure either interface to block or allow traffic based on the virtual LAN (VLAN) tags. 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.

Virtual Wire Subinterfaces

Virtual wire deployments can use virtual wire subinterfaces to separate traffic into zones. Virtual wire subinterfaces provide flexibility in enforcing distinct policies when you need to manage traffic from multiple customer networks. The subinterfaces allow 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 [Virtual Wire Deployment with Subinterfaces \(VLAN Tags only\)](#) shows an ISP using virtual wire subinterfaces with VLAN tags to separate traffic for two different customers.
- **VLAN tags in conjunction with IP classifiers (address, range, or subnet)**—The following example shows an ISP with two separate virtual systems on a firewall that manages traffic from two different customers. On each virtual system, the example illustrates how virtual wire subinterfaces with VLAN tags and IP classifiers are used to classify traffic into separate zones and apply relevant policy for customers from each network.

Virtual Wire Subinterface Workflow

- Configure two Ethernet interfaces as type virtual wire, and assign these interfaces to a virtual wire.

Virtual Wire Subinterface Workflow

- 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 subinterface(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 10: Virtual Wire Deployment with Subinterfaces (VLAN Tags only)

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.

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 11: 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)	Zone3	100	None
		e1/2.1 (egress)	Zone4	100	
	2	e1/1.2 (ingress)	Zone5	100	IP subnet 192.1.0.0/16
		e1/2.2 (egress)	Zone6	100	
	2	e1/1.3 (ingress)	Zone7	100	IP subnet 192.2.0.0/16
		e1/2.3 (egress)	Zone8	100	
B	3	e1/1.4 (ingress)	Zone9	200	None
		e1/2.4 (egress)	Zone10	200	

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

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



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

Configure Virtual Wires

The following task shows how to configure two [Virtual Wire Interfaces](#) (Ethernet 1/3 and Ethernet 1/4 in this example) to create a virtual wire. The two interfaces must have the same **Link Speed** and transmission mode (**Link Duplex**). For example, a full-duplex 1000Mbps copper port matches a full-duplex 1Gbps fiber optic port.

STEP 1 | Create the first virtual wire interface.

1. Select **Network > Interfaces > Ethernet** and select an interface you have cabled (**ethernet1/3** in this example).
2. Set the **Interface Type** to **Virtual Wire**.

STEP 2 | Attach the interface to a virtual wire object.

1. While still on the same Ethernet interface, on the **Config** tab, select **Virtual Wire** and click **New Virtual Wire**.
2. Enter a **Name** for the virtual wire.
3. For **Interface1**, select the interface you just configured (**ethernet1/3**). (Only interfaces configured as virtual wire interfaces appear in the list.)
4. For **Tag Allowed**, enter **0** to indicate untagged traffic (such as BPDUs and other Layer 2 control traffic) is allowed. The absence of a tag implies tag 0. Enter additional allowed tag integers or ranges of tags, separated by commas (default is 0; range is 0 to 4,094).
5. Select **Multicast Firewalling** if you want to be able to apply security policy rules to multicast traffic going across the virtual wire. Otherwise, multicast traffic is transparently forwarded across the virtual wire.
6. Select **Link State Pass Through** so the firewall can function transparently. When the firewall detects a link down state for a link of the virtual wire, it brings down the other interface in the virtual wire pair. Thus, devices on both sides of the firewall see a consistent link state, as if there were no firewall between them. If you don't select this option, link status is not propagated across the virtual wire.
7. Click **OK** to save the virtual wire object.

STEP 3 | Determine the link speed of the virtual wire interface.

1. While still on the same Ethernet interface, select **Advanced** and note or change the **Link Speed**. The port type determines the speed settings available in the list. By default, copper ports are set to **auto** negotiate link speed. Both virtual wire interfaces must have the same link speed.
2. Click **OK** to save the Ethernet interface.

STEP 4 | Configure the second virtual wire interface (**ethernet1/4** in this example) by repeating the preceding steps.

When you select the **Virtual Wire** object you created, the firewall automatically adds the second virtual wire interface as **Interface2**.

STEP 5 | Create a separate security zone for each virtual wire interface.

1. Select **Network > Zones** and **Add** a zone.
2. Enter the **Name** of the zone (such as **internet**).
3. For **Location**, select the virtual system where the zone applies.
4. For **Type**, select **Virtual Wire**.
5. **Add** the **Interface** that belongs to the zone.
6. Click **OK**.

STEP 6 | (Optional) Create security policy rules to allow Layer 3 traffic to pass through.

To allow Layer 3 traffic across the virtual wire, [Create a Security Policy Rule](#) to allow traffic from the user zone to the internet zone, and another to allow traffic from the internet zone to the user zone, selecting the applications you want to allow, such as BGP or OSPF.

STEP 7 | (Optional) Enable IPv6 firewalling.

If you want to be able to apply security policy rules to IPv6 traffic arriving at the virtual wire interface, enable IPv6 firewalling. Otherwise, IPv6 traffic is forwarded transparently.

1. Select **Device > Setup > Session** and edit Session Settings.
2. Select **Enable IPv6 Firewalling**.
3. Click **OK**.

STEP 8 | **Commit** your changes.

STEP 9 | (Optional) Configure an LLDP profile and apply it to the virtual wire interfaces (see [Configure LLDP](#)).

STEP 10 | (Optional) Apply non-IP protocol control to the virtual wire zones (see [Configure Protocol Protection](#)). Otherwise, all non-IP traffic is forwarded over the virtual wire.

Layer 2 Interfaces

In a Layer 2 deployment, the firewall provides switching between two or more networks. Devices are connected to a Layer 2 segment; the firewall forwards the frames to the proper port, which is associated with the MAC address identified in the frame. [Configure a Layer 2 Interface](#) when switching is required.



If you're using security group tags (SGTs) in a Cisco Trustsec network, it's a best practice to deploy inline firewalls in either Layer 2 or virtual wire mode. Firewalls in Layer 2 or virtual wire mode can inspect and provide threat prevention for the tagged traffic.

The following topics describe the different types of Layer 2 interfaces you can configure for each type of deployment you need, including details on using virtual LANs (VLANs) for traffic and policy separation among groups. Another topic describes how the firewall rewrites the inbound port VLAN ID number in a Cisco per-VLAN spanning tree (PVST+) or Rapid PVST+ bridge protocol data unit (BPDU).

- [Layer 2 Interfaces with No VLANs](#)


- [Layer 2 Interfaces with VLANs](#)
- [Configure a Layer 2 Interface](#)
- [Configure a Layer 2 Interface, Subinterface, and VLAN](#)
- [Manage Per-VLAN Spanning Tree \(PVST+\) BPDU Rewrite](#)

Layer 2 Interfaces with No VLANs

[Configure a Layer 2 Interface](#) on the firewall so it can act as a switch in your layer 2 network (not at the edge of the network). The Layer 2 hosts are probably geographically close to each other and belong to a single broadcast domain. The firewall provides security between the Layer 2 hosts when you assign the interfaces to security zones and apply security rules to the zones.

The hosts communicate with the firewall and each other at Layer 2 of the OSI model by exchanging frames. A frame contains an Ethernet header that includes a source and destination Media Access Control (MAC) address, which is a physical hardware address. MAC addresses are 48-bit hexadecimal numbers formatted as six octets separated by a colon or hyphen (for example, 00-85-7E-46-F1-B2).

The following figure has a firewall with three Layer 2 interfaces that each connect to a Layer 2 host in a one-to-one mapping.



The firewall begins with an empty MAC table. When the host with source address 0A-76-F2-60-EA-83 sends a frame to the firewall, the firewall doesn't have destination address 0B-68-2D-05-12-76 in its MAC table, so it doesn't know which interface to forward the frame to; it broadcasts the frame to all of its Layer 2 interfaces. The firewall puts source address 0A-76-F2-60-EA-83 and associated Eth1/1 into its MAC table.

The host at 0C-71-D4-E6-13-44 receives the broadcast, but the destination MAC address is not its own MAC address, so it drops the frame.


The receiving interface Ethernet 1/2 forwards the frame to its host. When host 0B-68-2D-05-12-76 responds, it uses the destination address 0A-76-F2-60-EA-83, and the firewall adds to its MAC table Ethernet 1/2 as the interface to reach 0B-68-2D-05-12-76.

Layer 2 Interfaces with VLANs

When your organization wants to divide a LAN into separate virtual LANs (VLANs) to keep traffic and policies for different departments separate, you can logically group Layer 2 hosts into VLANs and thus divide a Layer 2 network segment into broadcast domains. For example, you can create VLANs for the Finance and Engineering departments. To do so, [Configure a Layer 2 Interface, Subinterface, and VLAN](#).

The firewall acts as a switch to forward a frame with an Ethernet header containing a VLAN ID, and the destination interface must have a subinterface with that VLAN ID in order to receive that frame and forward it to the host. You configure a Layer 2 interface on the firewall and configure one or more logical subinterfaces for the interface, each with a VLAN tag (ID).

In the following figure, the firewall has four Layer 2 interfaces that connect to Layer 2 hosts belonging to different departments within an organization. Ethernet interface 1/3 is configured with subinterface .1 (tagged with VLAN 10) and subinterface .2 (tagged with VLAN 20), thus there are two broadcast domains on that segment. Hosts in VLAN 10 belong to Finance; hosts in VLAN 20 belong to Engineering.



In this example, the host at MAC address 0A-76-F2-60-EA-83 sends a frame with VLAN ID 10 to the firewall, which the firewall broadcasts to its other L2 interfaces. Ethernet interface 1/3 accepts the frame because it's connected to the host with destination 0C-71-D4-E6-13-44 and its subinterface .1 is assigned VLAN 10. Ethernet interface 1/3 forwards the frame to the Finance host.

Configure a Layer 2 Interface

Configure [Layer 2 Interfaces with No VLANs](#) when you want Layer 2 switching and you don't need to separate traffic among VLANs.

STEP 1 | Configure a Layer 2 interface.

1. Select **Network > Interfaces > Ethernet** and select an interface. The **Interface Name** is fixed, such as ethernet1/1.
2. For **Interface Type**, select **Layer2**.
3. Select the **Config** tab and assign the interface to a **Security Zone** or create a **New Zone**.
4. Configure additional Layer 2 interfaces on the firewall that connect to other Layer 2 hosts.

STEP 2 | Commit.

Click **OK** and **Commit**.

Configure a Layer 2 Interface, Subinterface, and VLAN

Configure [Layer 2 Interfaces with VLANs](#) when you want Layer 2 switching and traffic separation among VLANs. You can optionally control non-IP protocols between security zones on a Layer 2 interface or between interfaces within a single zone on a Layer 2 VLAN.

STEP 1 | Configure a Layer 2 interface and subinterface and assign a VLAN ID.

1. Select **Network > Interfaces > Ethernet** and select an interface. The **Interface Name** is fixed, such as ethernet1/1.
2. For **Interface Type**, select **Layer2**.
3. Select the **Config** tab.
4. For **VLAN**, leave the setting **None**.
5. Assign the interface to a **Security Zone** or create a **New Zone**.
6. Click **OK**.
7. With the Ethernet interface highlighted, click **Add Subinterface**.
8. The **Interface Name** remains fixed. After the period, enter the subinterface number, in the range 1-9,999.
9. Enter a VLAN **Tag ID** in the range 1-4,094.
10. Assign the subinterface to a **Security Zone**.
11. Click **OK**.

STEP 2 | Commit.

Click **Commit**.

STEP 3 | (Optional) Apply a Zone Protection profile with protocol protection to control non-IP protocol packets between Layer 2 zones (or between interfaces within a Layer 2 zone).

[Configure Protocol Protection](#).

Manage Per-VLAN Spanning Tree (PVST+) BPDU Rewrite

When an interface on the firewall is configured for a [Layer 2 interface](#), 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 the BPDU out. This default behavior beginning in PAN-OS 7.1 allows the firewall to correctly tag Cisco proprietary PVST+ and Rapid PVST+ frames between Cisco switches in VLANs on either side of the firewall so that spanning tree loop detection using Cisco PVST+ and Rapid PVST+ can function properly. The firewall is not participating in the Spanning Tree Protocol (STP) election process and there is no behavior change for other types of spanning tree.



The Cisco switch must have the loopguard disabled for the PVST+ or Rapid PVST+ BPDU rewrite to function properly on the firewall.

This feature is supported on Layer 2 Ethernet and Aggregated Ethernet (AE) interfaces only. The firewall supports a PVID range of 1 to 4,094 with a native VLAN ID of 1 to be compatible with the Cisco native VLAN implementation.


To support the PVST+ BPDU rewrite feature, PAN-OS supports the concept of a PVST+ native VLAN. Frames sent to and received from a native VLAN are untagged with a PVID equal to the native VLAN. All switches and firewalls in the same Layer 2 deployment must have the same native VLAN for PVST+ to function properly. Although the Cisco native VLAN defaults to vlan1, the VLAN ID could be a number other than 1.

For example, the firewall is configured with a VLAN object (named VLAN_BRIDGE), which describes the interfaces and subinterfaces that belong to a switch or broadcast domain. In this example, the VLAN includes three subinterfaces: ethernet1/21.100 tagged with 100, ethernet1/22.1000 tagged with 1000, and ethernet1/23.1500 tagged with 1500.

The subinterfaces belonging to VLAN_BRIDGE look like this:

Interface	Interface Type	Link State	Tag	VLAN / Virtual-Wire	Security Zone
ethernet1/21	Layer2		Untagged	none	none
ethernet1/21.100	Layer2		100	VLAN_BRIDGE	Zone_Trust
ethernet1/22	Layer2		Untagged	none	none
ethernet1/22.1000	Layer2		1000	VLAN_BRIDGE	Zone_Untrust
ethernet1/23	Layer2		Untagged	none	none
ethernet1/23.1500	Layer2		1500	VLAN_BRIDGE	Zone_Management

The sequence in which the firewall automatically rewrites the PVST+ BPDU is shown in the following graphic and explanation:



1. The Cisco switch port belonging to VLAN 100 sends a PVST+ BPDU—with the PVID and 802.1Q VLAN tag set to 100—to the firewall.
2. The firewall interfaces and subinterfaces are configured as a Layer 2 interface type. The ingress subinterface on the firewall is tagged with VLAN 100, which matches the PVID and VLAN tag of the incoming BPDU, so the firewall accepts the BPDU. The firewall floods the PVST+ BPDU to all other interfaces belonging to the same VLAN object (in this example, ethernet1/22.1000 and ethernet1/23.1500). If the VLAN tags did not match, the firewall would instead drop the BPDU.
3. When the firewall floods the BPDU out through other interfaces (belonging to the same VLAN object), the firewall rewrites the PVID and any 802.1Q VLAN tags to match the VLAN tag of the egress interface. In this example, the firewall rewrites the BPDU PVID from 100 to 1000 for one subinterface and from 100 to 1500 for the second subinterface as the BPDU traverses the Layer 2 bridge on the firewall.
4. Each Cisco switch receives the correct PVID and VLAN tag on the incoming BPDU and processes the PVST+ packet to detect possible loops in the network.

The following CLI operational commands allow you to manage PVST+ and Rapid PVST+ BPDUs.

- Globally disable or re-enable the PVST+ and Rapid PVST+ BPDU rewrite of the PVID (default is enabled).

```
set session rewrite-pvst-pvid <yes|no>
```

- Set the native VLAN ID for the firewall (range is 1 to 4,094; default is 1).



If the native VLAN ID on your switch is a value other than 1, you must set the native VLAN ID on the firewall to that same number; otherwise, the firewall will drop packets with that VLAN ID. This applies to trunked and non-trunked interfaces.

```
set session pvst-native-vlan-id <vid>
```

- Drop all STP BPDU packets.

```
set session drop-stp-packet <yes|no>
```

Examples of why you might want to drop all STP BPDU packets:

- If there is only one switch on each side of the firewall and no other connections between the switches that can cause a loop, then STP is not required and can be disabled on the switch or blocked by the firewall.
 - If there is a misbehaving STP switch inappropriately flooding BPDUs, you can stop the STP packets at the firewall to stop the BPDU flood.
- Verify whether PVST+ BPDU rewrite is enabled, view the PVST native VLAN ID, and determine whether the firewall is dropping all STP BPDU packets.

```
show vlan all
```

```
pvst+ tag rewrite: disabled
pvst native vlan id:      5
drop stp:                 disabled
total vlans shown:       1
name      interface      virtual interface
bridge    ethernet1/1
           ethernet1/2
           ethernet1/1.1
           ethernet1/2.1
```

- Troubleshoot PVST+ BPDU errors.

```
show counter global
```

Look at the `flow_pvid_inconsistent` counter, which counts the number of times the 802.1Q Tag and PVID fields inside a PVST+ BPDU packet don't match.

Layer 3 Interfaces

In a Layer 3 deployment, the firewall routes traffic between multiple ports. Before you can [Configure Layer 3 Interfaces](#), you must configure the [Virtual Routers](#) that you want the firewall to use to route the traffic for each Layer 3 interface.



If you're using security group tags (SGTs) in a Cisco Trustsec network, it's a best practice to deploy inline firewalls in either Layer 2 or virtual wire mode. However, if you need to use a Layer 3 firewall in a Cisco Trustsec network, you should deploy the Layer 3 firewall between two SGT exchange protocol (SXP) peers, and configure the firewall to allow traffic between the SXP peers.

The following topics describe how to configure Layer 3 interfaces, and how to use Neighbor Discovery Protocol (NDP) to provision IPv6 hosts and view the IPv6 addresses of devices on the link local network to quickly locate devices.

- [Configure Layer 3 Interfaces](#)
- [Manage IPv6 Hosts Using NDP](#)

Configure Layer 3 Interfaces

The following procedure is required to configure [Layer 3 Interfaces](#) (Ethernet, VLAN, loopback, and tunnel interfaces) with IPv4 or IPv6 addresses so that the firewall can perform routing on these interfaces. If a tunnel is used for routing or if tunnel monitoring is turned on, the tunnel needs an IP address. Before performing the following task, define one or more [Virtual Routers](#).

You would typically use the following procedure to configure an external interface that connects to the internet and an interface for your internal network. You can configure both IPv4 and IPv6 addresses on a single interface.



PAN-OS firewall models support a maximum of 16,000 IP addresses assigned to physical or virtual Layer 3 interfaces; this maximum includes both IPv4 and IPv6 addresses.

If you're using IPv6 routes, you can configure the firewall to provide [IPv6 Router Advertisements for DNS Configuration](#). The firewall provisions IPv6 DNS clients with Recursive DNS Server (RDNS) addresses and a DNS Search List so that the client can resolve its IPv6 DNS requests. Thus the firewall is acting like a DHCPv6 server for you.

STEP 1 | Select an interface and configure it with a security zone.

1. Select **Network > Interfaces** and either **Ethernet**, **VLAN**, **loopback**, or **Tunnel**, depending on what type of interface you want.
2. Select the interface to configure.
3. Select the **Interface Type—Layer3**.
4. On the **Config** tab, for **Virtual Router**, select the virtual router you are configuring, such as **default**.
5. For **Virtual System**, select the virtual system you are configuring if on a multi-virtual system firewall.
6. For **Security Zone**, select the zone to which the interface belongs or create a **New Zone**.
7. Click **OK**.

STEP 2 | Configure the interface with an IPv4 address.

You can assign an IPv4 address to a Layer 3 interface in one of three ways:

- **Static**
- **DHCP Client**—The firewall interface acts as a DHCP client and receives 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.

- PPPoE—Configure the interface as 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.
1. Select **Network > Interfaces** and either **Ethernet**, **VLAN**, **loopback**, or **Tunnel**, depending on what type of interface you want.
 2. Select the interface to configure.
 3. To configure the interface with a static IPv4 address, on the **IPv4** tab, set **Type** to **Static**.
 4. **Add a Name** and optional **Description** for the address.
 5. For **Type**, select one of the following:
 - **IP Netmask**—Enter the IP address and network mask to assign to the interface, for example, 208.80.56.100/24.



If you're using a /31 subnet mask for the Layer 3 interface address, the interface must be configured with the .1/31 address in order for utilities such as ping to work properly.



If you're configuring a loopback interface with an IPv4 address, it must have a /32 subnet mask; for example, 192.168.2.1/32.

- **IP Range**—Enter an IP address range, such as 192.168.2.1-192.168.2.4.
 - **FQDN**—Enter a Fully Qualified Domain Name.
6. Select **Tags** to apply to the address.
 7. Click **OK**.

STEP 3 | Configure an interface with Point-to-Point Protocol over Ethernet (PPPoE). See [Layer 3 Interfaces](#).



PPPoE is not supported in HA active/active mode.

1. Select **Network > Interfaces** and either **Ethernet**, **VLAN**, **loopback**, or **Tunnel**.
2. Select the interface to configure.
3. On the **IPv4** tab, set **Type** to **PPPoE**.
4. On the **General** tab, select **Enable** to activate the interface for PPPoE termination.
5. Enter the **Username** for the point-to-point connection.
6. Enter the **Password** for the username and **Confirm Password**.
7. Click **OK**.

STEP 4 | Configure an Interface as a DHCP Client so that it receives a dynamically-assigned IPv4 address.



DHCP client is not supported in HA active/active mode.

STEP 5 | Configure an interface with a static IPv6 address.

1. Select **Network > Interfaces** and either **Ethernet**, **VLAN**, **loopback**, or **Tunnel**.
2. Select the interface to configure.
3. On the **IPv6** tab, select **Enable IPv6 on the interface** to enable IPv6 addressing on the interface.
4. For **Interface ID**, enter the 64-bit extended unique identifier (EUI-64) in hexadecimal format (for example, 00:26:08:FF:FE:DE:4E:29). If you leave this field blank, the firewall uses the EUI-64

generated from the MAC address of the physical interface. If you enable the **Use interface ID as host portion** option when adding an address, the firewall uses the Interface ID as the host portion of that address.

5. **Add** the IPv6 **Address** or select an address group.
6. Select **Enable address on interface** to enable this IPv6 address on the interface.
7. Select **Use interface ID as host portion** to use the Interface ID as the host portion of the IPv6 address.
8. (Optional) Select **Anycast** to make the IPv6 address (route) an Anycast address (route), which means multiple locations can advertise the same prefix, and IPv6 sends the anycast traffic to the node it considers the nearest, based on routing protocol costs and other factors.
9. (Ethernet interface only) Select **Send Router Advertisement** (RA) to enable the firewall to send this address in Router Advertisements, in which case you must also enable the global **Enable Router Advertisement** option on the interface (next step).
10. (Ethernet interface only) Enter the **Valid Lifetime (sec)**, in seconds, that the firewall considers the address valid. The Valid Lifetime must equal or exceed the **Preferred Lifetime (sec)** (default is 2,592,000).
11. (Ethernet interface only) Enter the **Preferred Lifetime (sec)** (in seconds) that the valid address is preferred, which means the firewall can use it to send and received traffic. After the Preferred Lifetime expires, the firewall can't use the address to establish new connections, but any existing connections are valid until the **Valid Lifetime** expires (default is 604,800).
12. (Ethernet interface only) Select **On-link** if systems that have addresses within the prefix are reachable without a router.
13. (Ethernet interface only) Select **Autonomous** if systems can independently create an IP address by combining the advertised prefix with an Interface ID.
14. Click **OK**.

STEP 6 | (Ethernet or VLAN interface using IPv6 address only) Enable the firewall to send IPv6 Router Advertisements (RAs) from an interface, and optionally tune RA parameters.



Tune RA parameters for either of these reasons: To interoperate with a router/host that uses different values. To achieve fast convergence when multiple gateways are present. For example, set lower Min Interval, Max Interval, and Router Lifetime values so the IPv6 client/host can quickly change the default gateway after the primary gateway fails, and start forwarding to another default gateway in the network.

1. Select **Network > Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface you want to configure.
3. Select **IPv6**.
4. Select **Enable IPv6 on the interface**.
5. On the **Router Advertisement** tab, select **Enable Router Advertisement** (default is disabled).
6. (Optional) Set **Min Interval (sec)**, the minimum interval, in seconds, between RAs the firewall sends (range is 3-1,350; default is 200). The firewall sends RAs at random intervals between the minimum and maximum values you set.
7. (Optional) Set **Max Interval (sec)**, the maximum interval, in seconds, between RAs the firewall sends (range is 4-1,800; default is 600). The firewall sends RAs at random intervals between the minimum and maximum values you set.
8. (Optional) Set **Hop Limit** to apply to clients for outgoing packets (range is 1-255; default is 64). Enter 0 for no hop limit.
9. (Optional) Set **Link MTU**, the link maximum transmission unit (MTU) to apply to clients (range is 1,280-9,192; default is **unspecified**). Select **unspecified** for no link MTU.

10. (Optional) Set **Reachable Time (ms)**, the reachable time, in milliseconds, that the client will use to assume a neighbor is reachable after receiving a Reachability Confirmation message. Select **unspecified** for no reachable time value (range is 0-3,600,000; default is **unspecified**).
11. (Optional) Set **Retrans Time (ms)**, the retransmission timer that determines how long the client will wait, in milliseconds, before retransmitting Neighbor Solicitation messages. Select **unspecified** for no retransmission time (range is 0-4,294,967,295; default is **unspecified**).
12. (Optional) Set **Router Lifetime (sec)** to specify how long, in seconds, the client will use the firewall as the default gateway (range is 0-9,000; default is 1,800). Zero specifies that the firewall is not the default gateway. When the lifetime expires, the client removes the firewall entry from its Default Router List and uses another router as the default gateway.
13. Set **Router Preference**, which the client uses to select a preferred router if the network segment has multiple IPv6 routers. **High**, **Medium** (default), or **Low** is the priority that the RA advertises indicating the relative priority of firewall virtual router relative to other routers on the segment.
14. Select **Managed Configuration** to indicate to the client that addresses are available via DHCPv6.
15. Select **Other Configuration** to indicate to the client that other address information (such as DNS-related settings) is available via DHCPv6.
16. Select **Consistency Check** to have the firewall verify that RAs sent from other routers are advertising consistent information on the link. The firewall logs any inconsistencies.
17. Click **OK**.

STEP 7 | (Ethernet or VLAN interface using IPv6 address only) Specify the Recursive DNS Server addresses and DNS Search List the firewall will advertise in ND Router Advertisements from this interface.

The RDNS servers and DNS Search List are part of the DNS configuration for the DNS client so that the client can resolve IPv6 DNS requests.

1. Select **Network > Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface you are configuring.
3. Select **IPv6 > DNS Support**.
4. **Include DNS information in Router Advertisement** to enable the firewall to send IPv6 DNS information.
5. For **DNS Server**, **Add** the IPv6 address of a Recursive DNS Server. **Add** up to eight Recursive DNS servers. The firewall sends server addresses in an ICMPv6 Router Advertisement in order from top to bottom.
6. Specify the **Lifetime** in seconds, which is the maximum length of time the client can use the specific RDNS Server to resolve domain names.
 - The **Lifetime** range is any value equal to or between the **Max Interval** (that you configured on the **Router Advertisement** tab) and two times that **Max Interval**. For example, if your Max Interval is 600 seconds, the Lifetime range is 600-1,200 seconds.
 - The default **Lifetime** is 1,200 seconds.
7. For **DNS Suffix**, **Add** a **DNS Suffix** (domain name of a maximum of 255 bytes). **Add** up to eight DNS suffixes. The firewall sends suffixes in an ICMPv6 Router Advertisement in order from top to bottom.
8. Specify the **Lifetime** in seconds, which is the maximum length of time the client can use the suffix. The Lifetime has the same range and default value as the **Server**.
9. Click **OK**.

STEP 8 | (Ethernet or VLAN interface) Specify static ARP entries. Static ARP entries reduce ARP processing.

1. Select **Network > Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface you are configuring.
3. Select **Advanced > ARP Entries**.

-
4. **Add** an **IP Address** and its corresponding **MAC Address** (hardware or media access control address). For a VLAN interface, you must also select the **Interface**.



Static ARP entries do not time out. Auto learned ARP entries in the cache time out in 1,800 seconds by default; you can customize the ARP cache timeout; see [Configure Session Timeouts](#).

5. Click **OK**.

STEP 9 | (**Ethernet or VLAN interface**) Specify static Neighbor Discovery Protocol (NDP) entries. NDP for IPv6 performs functions similar to those provided by ARP for IPv4.

1. Select **Network > Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface you are configuring.
3. Select **Advanced > ND Entries**.
4. **Add** an **IPv6 Address** and its corresponding **MAC Address**.
5. Click **OK**.

STEP 10 | (**Optional**) Enable services on the interface.

1. To enable services on the interface, select **Network > Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface you are configuring.
3. Select **Advanced > Other Info**.
4. Expand the **Management Profile** list and select a profile or **New Management Profile**.
5. Enter a **Name** for the profile.
6. For **Permitted Services**, select services, such as **Ping**, and click **OK**.

STEP 11 | **Commit** your changes.

STEP 12 | Cable the interface.

Attach straight through cables from interfaces you configured to the corresponding switch or router on each network segment.

STEP 13 | Verify that the interface is 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**.

STEP 14 | Configure static routes and/or a dynamic routing protocol (RIP, OSPF, or BGP) so that the virtual router can route traffic.

- [Configure a Static Route](#)
- [RIP](#)
- [OSPF](#)
- [BGP](#)

STEP 15 | Configure a default route.

[Configure a Static Route](#) and set it as the default.

Manage IPv6 Hosts Using NDP

This topic describes how you can use NDP to provision IPv6 hosts; therefore, you don't need a separate DHCPv6 server for that purpose. It also explains how to use NDP to monitor IPv6 addresses, allowing you to quickly track the IPv6 address and MAC address of a device and the associated user who has violated a security rule.

- [IPv6 Router Advertisements for DNS Configuration](#)
- [Configure RDNS Servers and DNS Search List for IPv6 Router Advertisements](#)
- [NDP Monitoring](#)
- [Enable NDP Monitoring](#)

IPv6 Router Advertisements for DNS Configuration

The firewall implementation of [Neighbor Discovery](#) (ND) is enhanced so that you can provision IPv6 hosts with the Recursive DNS Server (RDNSS) Option and DNS Search List (DNSSL) Option per [RFC 6106](#), [IPv6 Router Advertisement Options for DNS Configuration](#). When you [Configure Layer 3 Interfaces](#), you configure these DNS options on the firewall so the firewall can provision your IPv6 hosts; therefore you don't need a separate DHCPv6 server to provision the hosts. The firewall sends IPv6 Router Advertisements (RAs) containing these options to IPv6 hosts as part of their DNS configuration to fully provision them to reach internet services. Thus, your IPv6 hosts are configured with:


- The addresses of RDNS servers that can resolve DNS queries.
- A list of domain names (suffixes) that the DNS client appends (one at a time) to an unqualified domain name before entering the domain name into a DNS query.

IPv6 Router Advertisement for DNS configuration is supported for Ethernet interfaces, subinterfaces, Aggregated Ethernet interfaces, and Layer 3 VLAN interfaces on all PAN-OS platforms.



The capability of the firewall to send IPv6 RAs for DNS configuration allows the firewall to perform a role similar to DHCP, and is unrelated to the firewall being a DNS proxy, DNS client or DNS server.

After you configure the firewall with the addresses of RDNS servers, the firewall provisions an IPv6 host (the DNS client) with those addresses. The IPv6 host uses one or more of those addresses to reach an RDNS server. Recursive DNS refers to a series of DNS requests by an RDNS Server, as shown with three pairs of queries and responses in the following figure. For example, when a user tries to access [www.paloaltonetworks.com](#), the local browser sees that it does not have the IP address for that domain name in its cache, nor does the client's operating system have it. The client's operating system launches a DNS query to a Recursive DNS Server belonging to the local ISP.



An IPv6 Router Advertisement can contain multiple DNS Recursive Server Address options, each with the same or different lifetimes. A single DNS Recursive DNS Server Address option can contain multiple Recursive DNS Server addresses as long as the addresses have the same lifetime.

A DNS Search List is a list of domain names (suffixes) that the firewall advertises to a DNS client. The firewall thus provisions the DNS client to use the suffixes in its unqualified DNS queries. The DNS client

appends the suffixes, one at a time, to an unqualified domain name before it enters the name into a DNS query, thereby using a fully qualified domain name (FQDN) in the DNS query. For example, if a user (of the DNS client being configured) tries to submit a DNS query for the name “quality” without a suffix, the router appends a period and the first DNS suffix from the DNS Search List to the name and transmits a DNS query. If the first DNS suffix on the list is “company.com”, the resulting DNS query from the router is for the FQDN “quality.company.com”.

If the DNS query fails, the client appends the second DNS suffix from the list to the unqualified name and transmits a new DNS query. The client uses the DNS suffixes in order until a DNS lookup succeeds (ignoring the remaining suffixes) or the router has tried all of the suffixes on the list.

You configure the firewall with the suffixes that you want to provide to the DNS client router in an ND DNSSL option; the DNS client receiving the DNS Search List option is provisioned to use the suffixes in its unqualified DNS queries.

To specify RDNS Servers and a DNS Search List, [Configure RDNS Servers and DNS Search List for IPv6 Router Advertisements](#).

Configure RDNS Servers and DNS Search List for IPv6 Router Advertisements

Perform this task to configure [IPv6 Router Advertisements for DNS Configuration](#) of IPv6 hosts.

STEP 1 | Enable the firewall to send IPv6 Router Advertisements from an interface.

1. Select **Network > Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface to configure.
3. On the **IPv6** tab, select **Enable IPv6 on the interface**.
4. On the **Router Advertisement** tab, select **Enable Router Advertisement**.
5. Click **OK**.

STEP 2 | Specify the Recursive DNS Server addresses and DNS Search List the firewall will advertise in ND Router Advertisements from this interface.

The RDNS servers and DNS Search List are part of the DNS configuration for the DNS client so that the client can resolve IPv6 DNS requests.

1. Select **Network > Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface you are configuring.
3. Select **IPv6 > DNS Support**.
4. **Include DNS information in Router Advertisement** to enable the firewall to send IPv6 DNS information.
5. For **DNS Server**, **Add** the IPv6 address of a Recursive DNS Server. **Add** up to eight Recursive DNS servers. The firewall sends server addresses in an ICMPv6 Router Advertisement in order from top to bottom.
6. Specify the **Lifetime** in seconds, which is the maximum length of time the client can use the specific RDNS Server to resolve domain names.
 - The **Lifetime** range is any value equal to or between the **Max Interval** (that you configured on the **Router Advertisement** tab) and two times that **Max Interval**. For example, if your Max Interval is 600 seconds, the Lifetime range is 600-1,200 seconds.
 - The default **Lifetime** is 1,200 seconds.
7. For **DNS Suffix**, **Add** a **DNS Suffix** (domain name of a maximum of 255 bytes). **Add** up to eight DNS suffixes. The firewall sends suffixes in an ICMPv6 Router Advertisement in order from top to bottom.
8. Specify the **Lifetime** in seconds, which is the maximum length of time the client can use the suffix. The Lifetime has the same range and default value as the **Server**.
9. Click **OK**.

STEP 3 | Commit your changes.

Click **Commit**.

NDP Monitoring

Neighbor Discovery Protocol (NDP) for IPv6 ([RFC 4861](#)) performs functions similar to ARP functions for IPv4. The firewall by default runs NDP, which uses ICMPv6 packets to discover and track the link-layer addresses and status of neighbors on connected links.

[Enable NDP Monitoring](#) so you can view the IPv6 addresses of devices on the link local network, their MAC address, associated username from User-ID (if the user of that device used the directory service to log in), reachability Status of the address, and Last Reported date and time the NDP monitor received a Router Advertisement from this IPv6 address. The username is on a best-case basis; there can be many IPv6 devices on a network with no username, such as printers, fax machines, servers, etc.

If you want to quickly track a device and user who has violated a security rule, it is very useful to have the IPv6 address, MAC address and username displayed all in one place. You need the MAC address that corresponds to the IPv6 address in order to trace the MAC address back to a physical switch or Access Point.



NDP monitoring is not guaranteed to discover all devices because there could be other networking devices between the firewall and the client that filter out NDP or Duplicate Address Detection (DAD) messages. The firewall can monitor only the devices that it learns about on the interface.

NDP monitoring also monitors Duplicate Address Detection (DAD) packets from clients and neighbors. You can also monitor IPv6 ND logs to make troubleshooting easier.

NDP monitoring is supported for Ethernet interfaces, subinterfaces, Aggregated Ethernet interfaces, and VLAN interfaces on all PAN-OS models.

Enable NDP Monitoring

Perform this task to enable [NDP Monitoring](#) for an interface.

STEP 1 | Enable NDP monitoring.

1. Select **Network > Interfaces** and **Ethernet** or **VLAN**.
2. Select the interface you are configuring.
3. Select **IPv6**.
4. Select **Address Resolution**.
5. Select **Enable NDP Monitoring**.



After you enable or disable NDP monitoring, you must Commit before NDP monitoring can start or stop.

6. Click **OK**.

STEP 2 | Commit your changes.

Click **Commit**.

STEP 3 | Monitor NDP and DAD packets from clients and neighbors.

1. Select **Network > Interfaces** and **Ethernet** or **VLAN**.
2. For the interface where you enabled NDP monitoring, in the Features column, hover over the NDP Monitoring  icon:

The NDP Monitoring summary for the interface displays the list of IPv6 **Prefixes** that this interface will send in the Router Advertisement (RA) if RA is enabled (they are the IPv6 prefixes of the interface itself).

The summary also indicates whether DAD, Router Advertisement, and DNS Support are enabled; IP addresses of any Recursive DNS Servers configured; and any DNS suffixes configured on the DNS Search List.

3. Click on the NDP Monitoring icon to display detailed information.

NDP Monitoring - vlan.128

IPv6 Address	MAC	User-ID	Status	Last Reported
☐ 2001:cb1d:12f2:350:2d3b:366:b5e5:8cc9	d8-bb-2c-8a-80-fa	unknown	STALE	2016/11/02 11:26:58
☐ 2001:cb1d:12f2:350:2d8a:aa88:cc69:c45f	88-27-93-4d-72-1e	unknown	REACHABLE	2016/11/01 22:34:15
☐ 2001:cb1d:12f2:350:2da1:12057:dd5f:1479	e0:b5-2d-2f-9e-18	unknown	STALE	2016/11/02 15:06:40
☐ 2001:cb1d:12f2:350:2dc5:cb1d:f471:9946	54-9f-13-32-f2-04	unknown	REACHABLE	2016/11/01 11:48:59
☐ 2001:cb1d:12f2:350:2da1:253b:e92b:a64b	68-d8-ca-7f-b4-9b	unknown	STALE	2016/11/01 17:33:19
☐ 2001:cb1d:12f2:350:2d00:b44e:f160-8f1e	b8-53-ac-df-ee-62	unknown	REACHABLE	2016/11/01 15:56:35

Clear All NDP Entries

Each row of the detailed NDP Monitoring table for the interface displays the IPv6 address of a neighbor the firewall has discovered, the corresponding MAC address, corresponding User ID (on a best-case basis), reachability Status of the address, and Last Reported date and time this NDP Monitor received an RA from this IP address. A User ID will not display for printers or other non-user-based hosts. If the status of the IP address is Stale, the neighbor is not known to be reachable, per RFC 4861.

At the bottom right is the count of **Total Devices Detected** on the link local network.

- Enter an IPv6 address in the filter field to search for an address to display.
- Select the check boxes to display or not display IPv6 addresses.
- Click the numbers, the right or left arrow, or the vertical scroll bar to advance through many entries.
- Click **Clear All NDP Entries** to clear the entire table.

STEP 4 | Monitor ND logs for reporting purposes.

1. Select **Monitor > Logs > System**.
2. In the Type column, view **ipv6nd** logs and corresponding descriptions.

For example, `inconsistent router advertisement received` indicates that the firewall received an RA different from the RA that it is going to send out.

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 VM-Series models support aggregate groups. You can add up to eight aggregate groups per firewall and each group can have up to eight interfaces.



PAN-OS firewall models support a maximum of 16,000 IP addresses assigned to physical or virtual Layer 3 interfaces; this maximum includes both IPv4 and IPv6 addresses.

Before configuring an aggregate group, you must configure its interfaces. Among the interfaces assigned to any particular aggregate group, the hardware media can differ (for example, you can mix fiber optic and copper), but the bandwidth and interface type must be the same. The bandwidth and interface type options are:

- **Bandwidth**—1Gbps, 10Gbps, 40Gbps, or 100Gbps
- **Interface type**—HA3, virtual wire, Layer 2, or Layer 3.



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.

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.

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.

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.

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.



When enabling access to a firewall interface using an Interface Management profile, do not enable management access (HTTP, HTTPS, SSH, or Telnet) from the internet or from other untrusted zones inside your enterprise security boundary, and never enable HTTP or Telnet access because those protocols transmit in cleartext. Follow the [Best Practices for Securing Administrative Access](#) to ensure that you are properly securing management access to your firewall.

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**.



Don't enable HTTP or Telnet because those protocols transmit in cleartext and therefore aren't secure.

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 [Allow Password Access to Certain Sites](#).
 - **User-ID**—Use to [Redistribute User Mappings and Authentication Timestamps](#).
 - **User-ID Syslog Listener-SSL or User-ID Syslog Listener-UDP**—Use to [Configure User-ID to Monitor Syslog Senders for User Mapping](#) 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**.

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**.

Virtual Routers

A virtual router is a function of the firewall that participates in Layer 3 routing. The firewall uses virtual routers to obtain routes to other subnets by you manually defining static routes or through participation in one or more Layer 3 routing protocols (dynamic routes). The routes that the firewall obtains through these methods populate the IP routing information base (RIB) on the firewall. When a packet is destined for a different subnet than the one it arrived on, the virtual router obtains the best route from the RIB, places it in the forwarding information base (FIB), and forwards the packet to the next hop router defined in the FIB. The firewall uses Ethernet switching to reach other devices on the same IP subnet. (An exception to one best route going in the FIB occurs if you are using ECMP, in which case all equal-cost routes go in the FIB.)

The Ethernet, VLAN, and tunnel interfaces defined on the firewall receive and forward Layer 3 packets. The destination zone is derived from the outgoing interface based on the forwarding criteria, and the firewall consults policy rules to identify the security policies that it applies to each packet. 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 Layer 3 interfaces on a virtual router to participate with dynamic routing protocols (BGP, OSPF, OSPFv3, or RIP) as well as add static routes. You can also create multiple virtual routers, each maintaining a separate set of routes that aren't shared between virtual routers, enabling you to configure different routing behaviors for different interfaces.

Each Layer 3 Ethernet, loopback, VLAN, and tunnel interface defined on the firewall must be associated with a virtual router. While each interface can belong to only one virtual router, you can configure multiple routing protocols and static routes for a virtual router. Regardless of the static routes and dynamic routing protocols you configure for a virtual router, one general configuration is required:

STEP 1 | Gather the required information from your network administrator.

- Interfaces on the firewall that you want to perform routing.
- Administrative distances for static, OSPF internal, OSPF external, IBGP, EBGP and RIP.

STEP 2 | Create a virtual router and apply interfaces to it.

The firewall comes with a virtual router named **default**. You can edit the **default** virtual router or add a new virtual router.

1. Select **Network > Virtual Routers**.
2. Select a virtual router (the one named **default** or a different virtual router) or **Add** the **Name** of a new virtual router.
3. Select **Router Settings > General**.
4. Click **Add** in the **Interfaces** box and select an already defined interface.

Repeat this step for all interfaces you want to add to the virtual router.

5. Click **OK**.

STEP 3 | Set Administrative Distances for static and dynamic routing.

Set Administrative Distances for types of routes as required for your network. When the virtual router has two or more different routes to the same destination, it uses administrative distance to choose the best path from different routing protocols and static routes, by preferring a lower distance.

- **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.



See [ECMP](#) if you want to leverage having multiple equal-cost paths for forwarding.

STEP 4 | Commit virtual router general settings.

Click **OK** and **Commit**.

STEP 5 | Configure Ethernet, VLAN, loopback, and tunnel interfaces as needed.

[Configure Layer 3 Interfaces](#).

Service Routes

The firewall uses the management (MGT) interface by default to access external services, such as DNS servers, external authentication servers, Palo Alto Networks services such as software, URL updates, licenses and AutoFocus. 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*. The service packets exit the firewall on the port assigned for the external service and the server sends its response to the configured source interface and source IP address.

You can configure service routes globally for the firewall (shown in the following task) or [Customize Service Routes for a Virtual System](#) on a firewall enabled for multiple virtual systems so that you have the flexibility to use interfaces associated with a virtual system. 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.

The following procedure enables you to change the interface the firewall uses to send requests to external services.

STEP 1 | Customize service routes.

1. Select **Device > Setup > Services > Global** (omit Global on a firewall without multiple virtual system capability), and in the Services Features section, click **Service Route Configuration**.



2. Select **Customize** and do one of the following to create a service route:

- For a predefined service:
 - Select **IPv4** or **IPv6** and click the link for the service for which you want customize the service route.



*To easily use the same source address for multiple services, select the checkbox for the services, click **Set Selected Routes**, and proceed to the next step.*

- To limit the list for Source Address, select a **Source Interface**; then select a **Source Address** (from that interface) as the service route. Selecting **Any** Source Interface makes all IP addresses on all interfaces available in the Source Address list from which you select an address. Selecting **Use default** causes the firewall to use the management interface for the service route, unless the packet destination IP address matches the configured Destination IP address, in which case the source IP address is set to the **Source Address** configured for the **Destination**. Selecting **MGT** causes the firewall to use the MGT interface for the service route, regardless of any destination service route.
 - Click **OK** to save the setting.
 - Repeat this step if you want to specify both an IPv4 and IPv6 address for a service.
 - For a destination service route:
 - Select **Destination** and **Add a Destination** IP address. In this case, if a packet arrives with a destination IP address that matches this configured **Destination** address, then the source IP address of the packet will be set to the **Source Address** configured in the next step.
 - To limit the list for Source Address, select a **Source Interface**; then select a **Source Address** (from that interface) as the service route. Selecting **Any** Source Interface makes all IP addresses on all interfaces available in the Source Address list from which you select an address. Selecting **MGT** causes the firewall to use the MGT interface for the service route.
 - Click **OK** to save the setting.
3. Repeat the prior steps for each service route you want to customize.

-
4. Click **OK** to save the service route configuration.

STEP 2 | Commit.

Click **Commit**.

Static Routes

Static routes are typically used in conjunction with dynamic routing protocols. You might configure a static route for a location that a dynamic routing protocol can't reach. Static routes require manual configuration on every router in the network, rather than the firewall entering dynamic routes in its route tables; even though static routes require that configuration on all routers, they may be desirable in small networks rather than configuring a routing protocol.

- [Static Route Overview](#)
- [Static Route Removal Based on Path Monitoring](#)
- [Configure a Static Route](#)
- [Configure Path Monitoring for a Static Route](#)

Static Route Overview

If you decide that you want specific Layer 3 traffic to take a certain route without participating in IP routing protocols, you can [Configure a Static Route](#) using IPv4 and IPv6 routes.

A default route is a specific static route. If you don't use dynamic routing to obtain a default route for your virtual router, you must configure a static default route. When the virtual router has an incoming packet and finds no match for the packet's destination in its route table, the virtual router sends the packet to the default route. The default IPv4 route is 0.0.0.0/0; the default IPv6 route is ::/0. You can configure both an IPv4 and IPv6 default route.

Static routes themselves don't change or adjust to changes in network environments, so traffic typically isn't rerouted if a failure occurs along the route to a statically defined endpoint. However, you have options to back up static routes in the event of a problem:

- You can configure a static route with a Bidirectional Forwarding Detection ([BFD](#)) profile so that if a BFD session between the firewall and the BFD peer fails, the firewall removes the failed static route from the RIB and FIB tables and uses an alternative route with a lower priority.
- You can [Configure Path Monitoring for a Static Route](#) so that the firewall can use an alternative route.

By default, static routes have an administrative distance of 10. When the firewall has two or more routes to the same destination, it uses the route with the lowest administrative distance. By increasing the administrative distance of a static route to a value higher than a dynamic route, you can use the static route as a backup route if the dynamic route is unavailable.

While you're configuring a static route, you can specify whether the firewall installs an IPv4 static route in the unicast or multicast route table (RIB), or both tables, or doesn't install the route at all. For example, you could install an IPv4 static route in the multicast route table only, because you want only multicast traffic to use that route. This option gives you more control over which route the traffic takes. You can specify whether the firewall installs an IPv6 static route in the unicast route table or not.

Static Route Removal Based on Path Monitoring

When you [Configure Path Monitoring for a Static Route](#), the firewall uses path monitoring to detect when the path to one or more monitored destination has gone down. The firewall can then reroute traffic using alternative routes. The firewall uses path monitoring for static routes much like path monitoring for HA or policy-based forwarding (PBF), as follows:

- ❑ The firewall sends ICMP ping messages (heartbeat messages) to one or more monitored destinations that you determine are robust and reflect the availability of the static route.
- ❑ If pings to any or all of the monitored destinations fail, the firewall considers the static route down too and removes it from the Routing Information Base (RIB) and Forwarding Information Base (FIB). The RIB

is the table of static routes the firewall is configured with and dynamic routes it has learned from routing protocols. The FIB is the forwarding table of routes the firewall uses for forwarding packets. The firewall selects an alternative static route to the same destination (based on the route with the lowest metric) from the RIB and places it in the FIB.

- ❑ The firewall continues to monitor the failed route. When the route comes back up, and (based on the **Any** or **All** failure condition) the path monitor returns to Up state, the preemptive hold timer begins. The path monitor must remain up for the duration of the hold timer; then the firewall considers the static route stable and reinstates it into the RIB. The firewall then compares metrics of routes to the same destination to decide which route goes in the FIB.

Path monitoring is a desirable mechanism to avoid silently dropping traffic for:


- A static or default route.
- A static or default route redistributed into a routing protocol.
- A static or default route when one peer does not support BFD. (The best practice is not to enable both BFD and path monitoring on a single interface.)
- A static or default route instead of using PBF path monitoring, which doesn't remove a failed static route from the RIB, FIB, or redistribution policy.



Path monitoring doesn't apply to static routes configured between virtual routers.

In the following figure, the firewall is connected to two ISPs for route redundancy to the internet. The primary default route 0.0.0.0 (metric 10) uses Next Hop 192.0.2.10; the secondary default route 0.0.0.0 (metric 50) uses Next Hop 198.51.100.1. The customer premises equipment (CPE) for ISP A keeps the primary physical link active, even after internet connectivity goes down. With the link artificially active, the firewall can't detect that the link is down and that it should replace the failed route with the secondary route in its RIB.

To avoid silently dropping traffic to a failed link, configure path monitoring of 192.0.2.20, 192.0.2.30, and 192.0.2.40 and if all (or any) of the paths to these destinations fail, the firewall presumes the path to Next Hop 192.0.2.10 is also down, removes the static route 0.0.0.0 (that uses Next Hop 192.0.2.10) from its RIB, and replaces it with the secondary route to the same destination 0.0.0.0 (that uses Next Hop 198.51.100.1), which also accesses the internet.



When you [Configure a Static Route](#), one of the required fields is the Next Hop toward that destination. The type of next hop you configure determines the action the firewall takes during path monitoring, as follows:

If Next Hop Type in Static Route is:	Firewall Action for ICMP Ping
IP Address	The firewall uses the source IP address and egress interface of the static route as the source address and egress interface in the ICMP ping. It uses the configured Destination IP address of the monitored destination as the ping's destination address. It uses the static route's next hop address as the ping's next hop address.
Next VR	The firewall uses the source IP address of the static route as the source address in the ICMP ping. The egress interface is based on the lookup result from the next hop's virtual router. The configured Destination IP address of the monitored destination is the ping's destination address.
None	The firewall uses the destination IP address of the path monitor as the next hop and sends the ICMP ping to the interface specified in the static route.

When path monitoring for a static or default route fails, the firewall logs a critical event (path-monitor-failure). When the static or default route recovers, the firewall logs another critical event (path-monitor-recovery).

Firewalls synchronize path monitoring configurations for an active/passive HA deployment, but the firewall blocks egress ICMP ping packets on a passive HA peer because it is not actively processing traffic. The firewall doesn't synchronize path monitoring configurations for active/active HA deployments.

Configure a Static Route

Perform the following task to configure [Static Routes](#) or a default route for a virtual router on the firewall.

STEP 1 | Configure a static route.

1. Select **Network > Virtual Router** and select the virtual router you are configuring, such as **default**.
2. Select the **Static Routes** tab.
3. Select **IPv4** or **IPv6**, depending on the type of static route you want to configure.
4. **Add a Name** for the route.
5. For **Destination**, enter the route and netmask (for example, 192.168.2.2/24 for an IPv4 address or 2001:db8:123:1::1/64 for an IPv6 address). If you're creating a default route, enter the default route (0.0.0.0/0 for an IPv4 address or ::/0 for an IPv6 address). Alternatively, you can create an address object of type IP Netmask.
6. (**Optional**) For **Interface**, specify the outgoing interface for packets to use to go to the next hop. Use this for stricter control over which interface the firewall uses rather than the interface in the route table for the next hop of this route.
7. For **Next Hop**, select one of the following:
 - **IP Address**—Enter the IP address (for example, 192.168.56.1 or 2001:db8:49e:1::1) when you want to route to a specific next hop. You must **Enable IPv6 on the interface** (when you [Configure Layer 3 Interfaces](#)) to use an IPv6 next hop address. If you're creating a default route, for **Next Hop** you must select **IP Address** and enter the IP address for your Internet gateway (for example, 192.168.56.1 or 2001:db8:49e:1::1). Alternatively, you can create an address object of type IP Netmask. The address object must have a netmask of /32 for IPv4 or /128 for IPv6.
 - **Next VR**—Select this option and then select a virtual router if you want to route internally to a different virtual router on the firewall.
 - **FQDN**—Enter an FQDN or select an address object that uses an FQDN, or create a new address object of type FQDN.



If you use an FQDN as a static route next hop, that FQDN must resolve to an IP address that belongs to the same subnet as the interface you configured for the static route; otherwise, the firewall rejects the resolution and the FQDN remains unresolved.



The firewall uses only one IP address (from each IPv4 or IPv6 family type) from the DNS resolution of the FQDN. If the DNS resolution returns more than one address, the firewall uses the preferred IP address that matches the IP family type (IPv4 or IPv6) configured for the next hop. The preferred IP address is the first address the DNS server returns in its initial response. The firewall retains this address as preferred as long as the address appears in subsequent responses, regardless of its order.

- **Discard**—Select to drop packets that are addressed to this destination.
 - **None**—Select if there is no next hop for the route. For example, a point-to-point connection does not require a next hop because there is only one way for packets to go.
8. Enter an **Admin Distance** for the route to override the default administrative distance set for static routes for this virtual router (range is 10 to 240; default is 10).
 9. Enter a **Metric** for the route (range is 1 to 65,535).

STEP 2 | Choose where to install the route.

Select the **Route Table** (the RIB) into which you want the firewall to install the static route:

- **Unicast**—Install the route in the unicast route table. Choose this option if you want the route used only for unicast traffic.
- **Multicast**—Install the route in the multicast route table (available for IPv4 routes only). Choose this option if you want the route used only for multicast traffic.
- **Both**—Install the route in the unicast and multicast route tables (available for IPv4 routes only). Choose this option if you want either unicast or multicast traffic to use the route.
- **No Install**—Do not install the route in either route table.

STEP 3 | (Optional) If your firewall model supports **BFD**, you can apply a **BFD Profile** to the static route so that if the static route fails, the firewall removes the route from the RIB and FIB and uses an alternative route. Default is **None**.

STEP 4 | Click **OK** twice.

STEP 5 | **Commit** the configuration.

Configure Path Monitoring for a Static Route

Use the following procedure to configure [Static Route Removal Based on Path Monitoring](#).

STEP 1 | Enable path monitoring for a static route.

1. Select **Network > Virtual Routers** and select a virtual router.
2. Select **Static Routes**, select **IPv4** or **IPv6**, and select the static route you want to monitor. You can monitor up to 128 static routes.
3. Select **Path Monitoring** to enable path monitoring for the route.

STEP 2 | Configure the monitored destination(s) for the static route.

1. **Add** a monitored destination by **Name**. You can add up to eight monitored destinations per static route.
2. Select **Enable** to monitor the destination.

3. For **Source IP**, select the IP address that the firewall uses in the ICMP ping to the monitored destination:
 - If the interface has multiple IP addresses, select one.
 - If you select an interface, the firewall uses the first IP address assigned to the interface by default.
 - If you select **DHCP (Use DHCP Client address)**, the firewall uses the address that DHCP assigned to the interface. To see the DHCP address, select **Network > Interfaces > Ethernet** and in the row for the Ethernet interface, click on **Dynamic DHCP Client**. The IP Address displays in the Dynamic IP Interface Status window.
4. For **Destination IP**, enter an IP address or address object to which the firewall will monitor the path. The monitored destination and static route destination must use the same address family (IPv4 or IPv6).



The destination IP address should belong to a reliable endpoint; you wouldn't want to base path monitoring on a device that itself is unstable or unreliable.

5. (Optional) Specify the **ICMP Ping Interval (sec)** in seconds to determine how frequently the firewall monitors the path (range is 1-60; default is 3).
6. (Optional) Specify the **ICMP Ping Count** of packets that don't return from the destination before the firewall considers the static route down and removes it from the RIB and FIB (range is 3-10; default is 5).
7. Click **OK**.

STEP 3 | Determine whether path monitoring for the static route is based on one or all monitored destinations, and set the preemptive hold time.

1. Select a **Failure Condition**, whether **Any** or **All** of the monitored destinations for the static route must be unreachable by ICMP for the firewall to remove the static route from the RIB and FIB and add the static route that has the next lowest metric going to the same destination to the FIB.



Select All to avoid the possibility of any single monitored destination signaling a route failure when the destination is simply offline for maintenance, for example.

2. (Optional) Specify the **Preemptive Hold Time (min)**, which is the number of minutes a downed path monitor must remain in Up state before the firewall reinstalls the static route into the RIB. The path monitor evaluates all of its monitored destinations for the static route and comes up based on the **Any** or **All** failure condition. If a link goes down or flaps during the hold time, when the link comes back up, the path monitor can come back up; the timer restarts when the path monitor returns to Up state.

A **Preemptive Hold Time** of zero causes the firewall to reinstall the route into the RIB immediately upon the path monitor coming up. Range is 0-1,440; default is 2.

3. Click **OK**.

STEP 4 | Commit.

Click **Commit**.

STEP 5 | Verify path monitoring on static routes.

1. Select **Network > Virtual Routers** and in the row of the virtual router you are interested in, select **More Runtime Stats**.
2. From the **Routing** tab, select **Static Route Monitoring**.
3. For a static route (Destination), view whether Path Monitoring is Enabled or Disabled. The Status column indicates whether the route is Up, Down, or Disabled. Flags for the static route are: A—active, S—static, E—ECMP.
4. Select **Refresh** periodically to see the latest state of the path monitoring (health check).

-
5. Hover over the Status of a route to view the monitored IP addresses and results of the pings sent to the monitored destinations for that route. For example, 3/5 means that a ping interval of 3 seconds and a ping count of 5 consecutive missed pings (the firewall receives no ping in the last 15 seconds) indicates path monitoring detects a link failure. Based on the **Any** or **All** failure condition, if path monitoring is in failed state and the firewall receives a ping after 15 seconds, the path can be deemed up and the **Preemptive Hold Time** starts.

The State indicates the last monitored ping results: success or failed. Failed indicates that the series of ping packets (ping interval multiplied by ping count) was not successful. A single ping packet failure does not reflect a failed ping state.

STEP 6 | View the RIB and FIB to verify that the static route is removed.

1. Select **Network > Virtual Routers** and in the row of the virtual router you are interested in, select **More Runtime Stats**.
2. From the **Routing** tab, select **Route Table (RIB)** and then the **Forwarding Table (FIB)** to view each, respectively.
3. Select **Unicast** or **Multicast** to view the appropriate route table.
4. For **Display Address Family**, select **IPv4 and IPv6**, **IPv4 Only**, or **IPv6 Only**.
5. (Optional) In the filter field, enter the route you are searching for and select the arrow, or use the scroll bar to move through pages of routes.
6. See if the route is removed or present.
7. Select **Refresh** periodically to see the latest state of the path monitoring (health check).



*To view the events logged for path monitoring, select **Monitor > Logs > System**. View the entry for **path-monitor-failure**, which indicates path monitoring for a static route destination failed, so the route was removed. View the entry for **path-monitor-recovery**, which indicates path monitoring for the static route destination recovered, so the route was restored.*

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.

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.

Clear **Reject Default Route** if you want to permit redistribution of default routes through RIP.

STEP 3 | Configure interfaces for RIP.

1. On the **Interfaces** tab, select an interface 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** list.
6. Select normal, passive or send-only from the **Mode** list.
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 1-3,600; 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 1-3,600; 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).

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 **Add** a name for the authentication profile to authenticate RIP messages.
2. Select **Simple Password** as the **Password Type**.
3. Enter a simple password and then confirm.

MD5 RIP authentication

1. Select **Auth Profiles** and **Add** a name for the authentication profile to authenticate RIP messages.
2. Select **MD5** as the **Password Type**.
3. **Add** one or more password entries, including:
 - Key-ID (range is 0-255)
 - Key
4. (Optional) Select **Preferred** status.
5. Click **OK** to specify the key to be used to authenticate outgoing message.
6. Click **OK** again in the Virtual Router - RIP Auth Profile dialog box.

STEP 6 | Commit your changes.

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.

The 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 must 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.
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 OSPF areas. ABRs gather topology information from their connected areas and distribute it to the backbone area.
- **Backbone Router**—A backbone router is a router that runs OSPF and has at least one interface connected to the OSPF backbone area. 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.

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. 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.

Clear **Reject Default Route** if you want to permit redistribution of default routes through OSPF.

STEP 3 | Configure Areas - Type for the OSPF protocol.

1. On the **Areas** tab, **Add** 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.
2. On the **Type** tab, select one of the following from the area **Type** list:
 - **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 you select NSSA, select **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**—**Add** ranges of external routes that you want to **Advertise** or for which you want to **Suppress** advertising.
3. Click **OK**.

STEP 4 | Configure Areas - Range for the OSPF protocol

1. On the **Range** tab, **Add** 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.

STEP 5 | Configure Areas - Interfaces for the OSPF protocol

1. On the **Interface** tab, **Add** the following information for each interface to be included in the area:
 - **Interface**—Select an interface.
 - **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 and **Add** the neighbor IP addresses for all neighbors that are reachable through this interface.
 - **Metric**—Enter an OSPF metric for this interface (range is 0-65,535; 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.
2. Click **OK**.

STEP 6 | Configure Areas - Virtual Links.

1. On the **Virtual Link** tab, **Add** the following information for each virtual link to be included in the backbone area:
 - **Name**—Enter a name for the virtual link.
 - **Enable**—Select to enable 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.
 - **Timing**—It is recommended that you keep the default timing settings.
 - **Auth Profile**—Select a previously-defined authentication profile.
2. Click **OK** to save virtual links.
3. Click **OK** to save area.

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. Select the **Auth Profiles** tab and **Add** a name for the authentication profile to authenticate OSPF messages.
2. Select **Simple Password** as the **Password Type**.
3. Enter a simple password and then confirm.

MD5 OSPF authentication

1. Select the **Auth Profiles** tab and **Add** a name for the authentication profile to authenticate OSPF messages.

2. Select **MD5** as the **Password Type** and **Add** 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.
3. Click **OK**.

STEP 8 | Configure Advanced OSPF options.

1. On the **Advanced** tab, select **RFC 1583 Compatibility** to ensure compatibility with RFC 1583.
2. Specify a value for the **SPF Calculation Delay (sec)** timer, which allows you to tune the delay time (in seconds) between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should use the same delay value to optimize convergence times.
3. Specify a value for the **LSA Interval (sec)** timer, which is 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. Click **OK**.

STEP 9 | Commit your changes.

Configure OSPFv3

OSPF supports both IPv4 and IPv6. You must use OSPFv3 if you are using IPv6.

STEP 1 | Configure general virtual router configuration settings.

See [Virtual Routers](#) for details.

STEP 2 | Configure general OSPFv3 configuration settings.

1. Select the **OSPFv3** tab.
2. Select **Enable** to enable the OSPF protocol.
3. Enter the **Router ID**.
4. 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.

STEP 3 | 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) (recommended) or IPv6 Authentication Header (AH).

ESP OSPFv3 authentication

1. On the **Auth Profiles** tab, **Add** a name for the authentication profile to authenticate OSPFv3 messages.
2. Specify a Security Policy Index (**SPI**) (hexadecimal value in the range from 00000000 to FFFFFFFF). The two ends of the OSPFv3 adjacency must have matching SPI values.
3. Select **ESP** for **Protocol**.
4. Select a **Crypto Algorithm**.

You can select **None** or one of the following algorithms: **SHA1**, **SHA256**, **SHA384**, **SHA512**, or **MD5**.

5. 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, **Add** a name for the authentication profile to authenticate OSPFv3 messages.
2. 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.
3. Select **AH** for **Protocol**.
4. Select a **Crypto Algorithm**.

You must enter one of the following algorithms: **SHA1**, **SHA256**, **SHA384**, **SHA512**, or **MD5**.

5. Enter a value for **Key** and then confirm.
6. Click **OK**.
7. Click **OK** again in the Virtual Router - OSPF Auth Profile dialog.

STEP 4 | Configure Areas - Type for the OSPFv3 protocol.

1. On the **Areas** tab, **Add** an **Area ID**. This is the identifier that each neighbor must accept to be part of the same area.
2. On the **General** tab, select one of the following from the area **Type** list:
 - **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**—**Add** ranges of external routes that you want to enable or suppress advertising for.

STEP 5 | 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** list.
3. Click **OK**.

To an Interface

1. On the **Areas** tab, select an existing area from the table.
2. Select the **Interface** tab and **Add** the authentication profile you want to associate with the OSPF interface from the **Auth Profile** list.
3. Click **OK**.

STEP 6 | Click **OK** again to save the area settings.

STEP 7 | (Optional) Configure Export Rules.

1. On the **Export Rules** tab, select **Allow Redistribute Default Route** to permit redistribution of default routes through OSPFv3.

2. Click **Add**.
3. Enter the **Name**; the value must be a valid IPv6 subnet or valid redistribution profile name.
4. Select **New Path Type**, **Ext 1** or **Ext 2**.
5. Specify a **New Tag** for the matched route, using has a 32-bit value in dotted-decimal notation.
6. Assign a **Metric** to the new rule (range is 1-16,777,215).
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. Specify a value for the **SPF Calculation Delay (sec)** timer, which allows you to tune the delay time (in seconds) between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should use the same delay value to optimize convergence times.
3. Specify a value for the **LSA Interval (sec)** timer, which is the minimum time (in seconds) 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**.
5. Click **OK**.

STEP 9 | Commit your changes.

Configure OSPF Graceful Restart

OSPF Graceful Restart directs OSPF neighbors to continue using routes through a firewall 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**—If 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 neighbor receives Grace LSAs informing it that the firewall will perform a graceful restart within a specified period of time defined as the **Grace Period**. During the grace period, the neighbor continues to forward routes through the firewall and to send LSAs that announce routes through the firewall. If the firewall resumes operation before expiration of the grace period, traffic forwarding will continue as before without network disruption. If the firewall does not resume operation after the grace period has expired, the neighbors will exit helper mode and resume normal operation, which will involve reconfiguring the routing table to bypass the firewall.
- **Firewall as a Graceful Restart Helper**—If neighboring routers may be down for short periods of time, the firewall can be configured to operate in Graceful Restart helper mode, in which case the firewall employs a **Max Neighbor Restart Time**. When the firewall receives the Grace LSAs from its OSPF neighbor, it continues 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 exits helper mode and resumes normal operation, which involves reconfiguring the routing table to bypass the neighbor.

STEP 1 | Select **Network > Virtual Routers** and select the virtual router you want to configure.

STEP 2 | Select **OSPF > Advanced** or **OSPFv3 > 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:

STEP 1 | Select **Network > Virtual Routers** and in the same row as the virtual router you are interested in, click the **More Runtime Stats** link.

STEP 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:

STEP 1 | Select **Network > Virtual Routers** and in the same row as the virtual router you are interested in, click the **More Runtime Stats** link.

STEP 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.

STEP 1 | Select **Monitor > System** and look for messages to confirm that OSPF adjacencies have been established.

STEP 2 | Select **OSPF > Neighbor** and examine the **Status** column to determine if OSPF adjacencies have been established (are full).

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.

- [BGP Overview](#)
- [MP-BGP](#)
- [Configure BGP](#)
- [Configure a BGP Peer with MP-BGP for IPv4 or IPv6 Unicast](#)
- [Configure a BGP Peer with MP-BGP for IPv4 Multicast](#)
- [BGP Confederations](#)

BGP Overview

BGP functions between autonomous systems (exterior BGP or eBGP) or within an AS (interior BGP or iBGP) to exchange routing and reachability information with BGP speakers. The firewall provides a complete BGP implementation, which includes the following features:

- Specification of one BGP routing instance per virtual router.
- BGP settings per virtual router, which include basic parameters such as local route ID and local AS, and advanced options such as path selection, route reflector, [BGP Confederations](#), route flap dampening, and graceful restart.
- Peer group and neighbor settings, which include neighbor address and remote AS, and advanced options such as neighbor attributes and connections.
- Route policies to control route import, export and advertisement; prefix-based filtering; and address aggregation.
- IGP-BGP interaction to inject routes to BGP using redistribution profiles.
- Authentication profiles, which specify the MD5 authentication key for BGP connections. Authentication helps prevent route leaking and successful DoS attacks.
- Multiprotocol BGP (MP-BGP) to allow BGP peers to carry IPv6 unicast routes and IPv4 multicast routes in Update packets, and to allow the firewall and a BGP peer to communicate with each other using IPv6 addresses.

MP-BGP

BGP supports IPv4 unicast prefixes, but a BGP network that uses IPv4 multicast routes or IPv6 unicast prefixes needs multiprotocol BGP (MP-BGP) in order to exchange routes of address types other than IPv4 unicast. MP-BGP allows BGP peers to carry IPv4 multicast routes and IPv6 unicast routes in Update packets, in addition to the IPv4 unicast routes that BGP peers can carry without MP-BGP enabled.

In this way, MP-BGP provides IPv6 connectivity to your BGP networks that use either native IPv6 or dual stack IPv4 and IPv6. Service providers can offer IPv6 service to their customers, and enterprises can use IPv6 service from service providers. The firewall and a BGP peer can communicate with each other using IPv6 addresses.

In order for BGP to support multiple network-layer protocols (other than BGP for IPv4), [Multiprotocol Extensions for BGP-4 \(RFC 4760\)](#) use Network Layer Reachability Information (NLRI) in a Multiprotocol Reachable NLRI attribute that the firewall sends and receives in BGP Update packets. That attribute contains information about the destination prefix, including these two identifiers:

- The Address Family Identifier (AFI), as defined by the IANA in [Address Family Numbers](#), indicates that the destination prefix is an IPv4 or IPv6 address. (PAN-OS supports IPv4 and IPv6 AFIs.)


- The Subsequent Address Family Identifier (SAFI) in PAN-OS indicates that the destination prefix is a unicast or multicast address (if the AFI is IPv4), or that the destination prefix is a unicast address (if the AFI is IPv6). PAN-OS does not support IPv6 multicast.

If you enable MP-BGP for IPv4 multicast or if you configure a multicast static route, the firewall supports separate unicast and multicast route tables for static routes. You might want to separate the unicast and multicast traffic going to the same destination. The multicast traffic can take a different path from unicast traffic because, for example, your multicast traffic is critical, so you need it to be more efficient by having it take fewer hops or undergo less latency.

You can also exercise more control over how BGP functions by configuring BGP to use routes from only the unicast or multicast route table (or both) when BGP imports or exports routes, sends conditional advertisements, or performs route redistribution or route aggregation.

You can decide to use a dedicated multicast RIB (route table) by enabling MP-BGP and selecting the Address Family of IPv4 and Subsequent Address Family of multicast or by installing an IPv4 static route in the multicast route table. After you do either of those methods to use the multicast RIB, the firewall uses the multicast RIB for all multicast routing and reverse path forwarding (RPF). If you prefer to use the unicast RIB for all routing (unicast and multicast), you should not enable the multicast RIB by either method.

In the following figure, a static route to 192.168.10.0/24 is installed in the unicast route table, and its next hop is 198.51.100.2. However, multicast traffic can take a different path to a private MPLS cloud; the same static route is installed in the multicast route table with a different next hop (198.51.100.4) so that its path is different.



Using separate unicast and multicast route tables gives you more flexibility and control when you configure these BGP functions:

- Install an IPv4 static route into the unicast or multicast route table, or both, as described in the preceding example. (You can install an IPv6 static route into the unicast route table only).
- Create an Import rule so that any prefixes that match the criteria are imported into the unicast or multicast route table, or both.
- Create an Export rule so that prefixes that match the criteria are exported (sent to a peer) from the unicast or multicast route table, or both.
- Configure a conditional advertisement with a Non Exist filter so that the firewall searches the unicast or multicast route table (or both) to ensure the route doesn't exist in that table, and so the firewall advertises a different route.
- Configure a conditional advertisement with an Advertise filter so that the firewall advertises routes matching the criteria from the unicast or multicast route table, or both.
- Redistribute a route that appears in the unicast or multicast route table, or both.

- Configure route aggregation with an advertise filter so that aggregated routes to be advertised come from the unicast or multicast route table, or both.
- Conversely, configure route aggregation with a suppress filter so that aggregated routes that should be suppressed (not advertised) come from the unicast or multicast route table, or both.

When you configure a peer with MP-BGP using an Address Family of IPv6, you can use IPv6 addresses in the Address Prefix and Next Hop fields of an Import rule, Export rule, Conditional Advertisement (Advertise Filter and Non Exist Filter), and Aggregate rule (Advertise Filter, Suppress Filter, and Aggregate Route Attribute).

Configure BGP

Perform the following task to configure BGP.

STEP 1 | Configure general virtual router configuration settings.

See [Virtual Routers](#) for details.

STEP 2 | Enable BGP for the virtual router, assign a router ID, and assign the virtual router to an AS.

1. Select **Network > Virtual Routers** and select a virtual router.
2. Select **BGP**.
3. **Enable BGP** for this virtual router.
4. Assign a **Router ID** to BGP for the virtual router, which is typically an IPv4 address to ensure the Router ID is unique.
5. Assign the **AS Number**—the number of the AS to which the virtual router belongs based on the router ID (range is 1 to 4,294,967,295).
6. Click **OK**.

STEP 3 | Configure general BGP configuration settings.

1. Select **Network > Virtual Routers** and select a virtual router.
2. Select **BGP > General**.
3. Select **Reject Default Route** to ignore any default routes that are advertised by BGP peers.
4. Select **Install Route** to install BGP routes in the global routing table.
5. Select **Aggregate MED** to enable route aggregation even when routes have different Multi-Exit Discriminator (MED) values.
6. Specify the **Default Local Preference** that can be used to determine preferences among different paths.
7. Select the **AS Format** for interoperability purposes:

- **2 Byte** (default)
- **4 Byte**



Runtime stats display BGP 4-byte AS numbers using asplain notation according to RFC 5396.

8. Enable or disable each of the following settings 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).
9. For **Auth Profiles**, **Add** an authentication profile:
 - **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.

10. Click **OK** twice.

STEP 4 | (Optional) Configure BGP settings.

1. Select **Network > Virtual Routers** and select a virtual router.
2. Select **BGP > Advanced**.
3. Select **ECMP Multiple AS Support** if you configured ECMP and you want to run ECMP over multiple BGP autonomous systems.
4. **Enforce First AS for EBGP** (enabled by default) to cause the firewall to drop an incoming Update packet from an eBGP peer that does not list the eBGP peer's own AS number as the first AS number in the AS_PATH attribute.
5. Select **Graceful Restart** and configure the following timers:
 - **Stale Route Time (sec)**—Specifies the length of time, in seconds, that a route can stay in the stale state (range is 1 to 3,600; 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 to 3,600; default is 120).
 - **Max Peer Restart Time (sec)**—Specifies the maximum length of time, in seconds, that the local device accepts as a grace period restart time for peer devices (range is 1 to 3,600; default is 120).
6. For **Reflector Cluster ID**, specify an IPv4 identifier to represent the reflector cluster.
7. For **Confederation Member AS**, specify the autonomous system number identifier (also called a sub-AS number), which is visible only within the BGP confederation. For more information, see [BGP Confederations](#).
8. **Add** 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 to 1,000.0; default is 1.25).
 - **Reuse**—Specify a route withdrawal threshold below which a suppressed route is used again (range is 0.0 to 1,000.0; 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 to 3,600; 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 to 3,600; 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 to 3,600; default is 300).
9. Click **OK** twice.

STEP 5 | Configure a BGP peer group.

1. Select **Network > Virtual Routers** and select a virtual router.
2. Select **BGP > Peer Group**, **Add a Name** for the peer group, and **Enable** it.
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. Select the **Type** of peer group:
 - **IBGP—Export Next Hop**: Select **Original** or **Use self**.
 - **EBGP Confed—Export Next Hop**: Select **Original** or **Use self**.
 - **EBGP Confed—Export Next Hop**: Select **Original** or **Use self**.
 - **EBGP—Import Next Hop**: Select **Original** or **Use self**; and **Export Next Hop**: Specify **Resolve** or **Use self**. Select **Remove Private AS** if you want to force BGP to remove private AS numbers from the AS_PATH attribute in Updates that the firewall sends to a peer in another AS.

6. Click **OK**.

STEP 6 | Configure a BGP peer that belongs to the peer group and specify its addressing.

1. Select **Network > Virtual Routers** and select a virtual router.
2. Select **BGP > Peer Group** and select the peer group you created.
3. For Peer, **Add** a peer by **Name**.
4. **Enable** the peer.
5. Enter the **Peer AS** to which the peer belongs.
6. Select **Addressing**.
7. For **Local Address**, select the **Interface** for which you are configuring BGP. If the interface has more than one IP address, enter the IP address for that interface to be the BGP peer.
8. For **Peer Address**, select either **IP** and enter the IP address or select or create an address object, or select **FQDN** and enter the FQDN or address object that is type FQDN.



The firewall uses only one IP address (from each IPv4 or IPv6 family type) from the DNS resolution of the FQDN. If the DNS resolution returns more than one address, the firewall uses the preferred IP address that matches the IP family type (IPv4 or IPv6) configured for the BGP peer. The preferred IP address is the first address the DNS server returns in its initial response. The firewall retains this address as preferred as long as the address appears in subsequent responses regardless of its order.

9. Click **OK**.

STEP 7 | Configure connection settings for the BGP peer.

1. Select **Network > Virtual Routers** and select a virtual router.
2. Select **BGP > Peer Group** and select the peer group you created.
3. Select the **Peer** you configured.
4. Select **Connection Options**.
5. Select an **Auth Profile** for the peer.
6. Set a **Keep Alive Interval (sec)**—The interval, in seconds, after which routes from the peer are suppressed according to the Hold Time setting (range is 0 to 1,200; default is 30).
7. Set **Multi Hop**—The time-to-live (TTL) value in the IP header (range is 0 to 255; default is 0). The default value of 0 means 1 for eBGP. The default value of 0 means 255 for iBGP.
8. Set **Open Delay Time (sec)**—The delay, in seconds, between a TCP handshake and the firewall sending the first BGP Open message to establish a BGP connection (range is 0 to 240; default is 0).
9. Set **Hold Time (sec)**—The length of time, in seconds, that may elapse between successive Keepalive or Update messages from the peer before the peer connection is closed (range is 3 to 3,600; default is 90).
10. Set **Idle Hold Time (sec)**—The length of time to wait, in seconds, before retrying to connect to the peer (range is 1 to 3,600; default is 15).
11. Set **Min Route Advertisement Interval (sec)**—The minimum amount of time, in seconds, between two successive Update messages that a BGP speaker (the firewall) sends to a BGP peer that advertise routes or withdrawal of routes (range is 1 to 600; default is 30).
12. For **Incoming Connections**, enter a **Remote Port** and select **Allow** to allow incoming traffic to this port.
13. For **Outgoing Connections**, enter a **Local Port** and select **Allow** to allow outgoing traffic from this port.
14. Click **OK**.

STEP 8 | Configure the BGP peer with settings for route reflector client, peering type, maximum prefixes, and Bidirectional Forwarding Detection (BFD).

1. Select **Network > Virtual Routers** and select a virtual router.

2. Select **BGP > Peer Group** and select the peer group you created.
3. Select the **Peer** you configured.
4. Select **Advanced**.
5. For **Reflector Client**, select one of the following:
 - **non-client** (default)—Peer is not a route reflector client.
 - **client**—Peer is a route reflector client.
 - **meshed-client**
6. For **Peering Type**, select one of the following:
 - **Bilateral**—The two BGP peers establish a peer connection.
 - **Unspecified** (default).
7. For **Max Prefixes**, enter the maximum number of supported IP prefixes (range is 1 to 100,000) or select **unlimited**.
8. To enable **BFD** for the peer (and thereby override the BFD setting for BGP, as long as BFD is not disabled for BGP at the virtual router level), select one of the following:
 - **default**—Peer uses only default BFD settings.
 - **Inherit-vr-global-setting** (default)—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](#).



Select **Disable BFD** to disable BFD for the BGP peer.

9. Click **OK**.

STEP 9 | Configure Import and Export rules.

The import and export rules are used to import and export routes from and to other routers (for example, importing the default route from your Internet Service Provider).

1. Select **Import**, Add a name in the **Rules** field, and **Enable** the import rule.
2. **Add** the **Peer Group** from which the routes will be imported.
3. Select **Match** 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. Select **Action** and define the action that should occur (allow or deny) based on the filtering options defined in the **Match** tab. If you select **Deny**, you don't need to define any additional options. If you select **Allow**, then define the other attributes.
5. Select **Export** 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**.

STEP 10 | 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, such as when you have links to the internet through multiple ISPs and you want traffic to be routed to one provider instead of the other except when there is a loss of connectivity to the preferred provider.

1. Select **Conditional Adv** and **Add a Policy** name.
2. **Enable** the conditional advertisement.
3. In the **Used By** section, **Add** the peer groups that will use the conditional advertisement policy.

4. Select **Non Exist Filter** and define the network prefixes of the preferred route. This specifies the route that you want to advertise when 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 **Advertise Filters** and define the prefixes 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 unavailable 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.
6. Click **OK**.

STEP 11 | 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 the creation of one aggregate address. This will result in an aggregate entry in the routing table when at least one specific route matching the address specified is learned.

1. Select **Aggregate** and **Add** a name for the aggregate address.
2. Enter the network **Prefix** that will be the primary prefix for the aggregated prefixes.
3. Select **Suppress Filters** and define the attributes that will cause the matched routes to be suppressed.
4. Select **Advertise Filters** and define the attributes that will cause the matched routes to always be advertised to peers.
5. Click **OK**.

STEP 12 | Configure redistribution rules.

This rule is used to redistribute host routes and unknown routes that are not on the local RIB to the peer routers.

1. Select **Redist Rules** and **Add** a new redistribution rule.
2. Enter the **Name** of an IP subnet or select a redistribution profile. You can also configure a new redistribution profile if needed.
3. **Enable** the rule.
4. Enter the route **Metric** that will be used for the rule.
5. In the **Set Origin** list, select **incomplete**, **igp**, or **egp**.
6. (Optional) Set MED, local preference, AS path limit, and community values.
7. Click **OK**.

STEP 13 | Commit your changes.

Configure a BGP Peer with MP-BGP for IPv4 or IPv6 Unicast

After you [Configure BGP](#), configure a BGP peer with **MP-BGP** for IPv4 or IPv6 unicast for either of the following reasons:

- To have your BGP peer carry IPv6 unicast routes, configure MP-BGP with the Address Family Type of **IPv6** and Subsequent Address Family of **Unicast** so that the peer can send BGP updates that include IPv6 unicast routes. BGP peering (Local Address and Peer Address) can still both be IPv4 addresses, or they can both be IPv6 addresses.
- To perform BGP peering over IPv6 addresses (**Local Address** and **Peer Address** use IPv6 addresses).

The following task shows how to enable a BGP peer with MP-BGP so it can carry IPv6 unicast routes, and so it can peer using IPv6 addresses.

The task also shows how to view the unicast or multicast route tables, and how to view the forwarding table, the BGP local RIB, and BGP RIB Out (routes sent to neighbors) to see routes from the unicast or multicast route table or a specific address family (IPv4 or IPv6).

STEP 1 | Enable MP-BGP Extensions for a peer.

Configure the following so that a BGP peer can carry IPv4 or IPv6 unicast routes in Updates packets and the firewall can use IPv4 or IPv6 addresses to communicate with its peer.

1. Select **Network > Virtual Routers** and select the virtual router you are configuring.
2. Select **BGP**.
3. Select **Peer Group** and select a peer group.
4. Select a BGP peer (router).
5. Select **Addressing**.
6. Select **Enable MP-BGP Extensions** for the peer.
7. For **Address Family Type**, select **IPv4** or **IPv6**. For example, select **IPv6**.
8. For **Subsequent Address Family**, **Unicast** is selected. If you chose **IPv4** for the Address Family, you can select **Multicast** also.
9. For **Local Address**, select an **Interface** and optionally select an **IP** address, for example, 2001:DB8:55::/32
10. For **Peer Address**, enter the peer's **IP** address, using the same address family (IPv4 or IPv6) as the Local Address, for example, 2001:DB8:58::/32.
11. Select **Advanced**.
12. (Optional) **Enable Sender Side Loop Detection**. When you enable Sender Side Loop Detection, the firewall will check the AS_PATH attribute of a route in its FIB before it sends the route in an update, to ensure that the peer AS number is not on the AS_PATH list. If it is, the firewall removes it to prevent a loop
13. Click **OK**.

STEP 2 | (Optional) Create a static route and install it in the unicast route table because you want the route to be used only for unicast purposes.

1. Select **Network > Virtual Routers** and select the virtual router you are configuring.
2. Select **Static Routes**, select **IPv4** or **IPv6**, and **Add** a route.
3. Enter a **Name** for the static route.
4. Enter the IPv4 or IPv6 **Destination** prefix and netmask, depending on whether you chose IPv4 or IPv6.
5. Select the egress **Interface**.
6. Select the **Next Hop** as **IPv6 Address** (or **IP Address** if you chose IPv4) and enter the address of the next hop to which you want to direct unicast traffic for this static route.
7. Enter an **Admin Distance**.
8. Enter a **Metric**.
9. For **Route Table**, select **Unicast**.
10. Click **OK**.

STEP 3 | Commit the configuration.

Click **Commit**.

STEP 4 | View the unicast or multicast route table.

1. Select **Network > Virtual Routers**.
2. In the row for the virtual router, click **More Runtime Stats**.
3. Select **Routing > Route Table**.
4. For **Route Table**, select **Unicast** or **Multicast** to display only those routes.
5. For **Display Address Family**, select **IPv4 Only**, **IPv6 Only**, or **IPv4 and IPv6** to display only routes for that address family.



Selecting Multicast with IPv6 Only is not supported.

STEP 5 | View the Forwarding Table.

1. Select **Network > Virtual Routers**.
2. In the row for the virtual router, click **More Runtime Stats**.
3. Select **Routing > Forwarding Table**.
4. For **Display Address Family**, select **IPv4 Only**, **IPv6 Only**, or **IPv4 and IPv6** to display only routes for that address family.

STEP 6 | View the BGP RIB tables.

1. View the BGP Local RIB, which shows the BGP routes that the firewall uses to route BGP packets.
 1. Select **Network > Virtual Routers**.
 2. In the row for the virtual router, click **More Runtime Stats**.
 3. Select **BGP > Local RIB**.
 4. For **Route Table**, select **Unicast** or **Multicast** to display only those routes.
 5. For **Display Address Family**, select **IPv4 Only**, **IPv6 Only**, or **IPv4 and IPv6** to display only routes for that address family.



Selecting Multicast with IPv6 Only is not supported.

2. View the BGP RIB Out table, which shows the routes that the firewall sends to BGP neighbors.
 1. Select **Network > Virtual Routers**.
 2. In the row for the virtual router, click **More Runtime Stats**.
 3. Select **BGP > RIB Out**.
 4. For **Route Table**, select **Unicast** or **Multicast** to display only those routes.
 5. For **Display Address Family**, select **IPv4 Only**, **IPv6 Only**, or **IPv4 and IPv6** to display only routes for that address family.



Selecting Multicast with IPv6 Only is not supported.

Configure a BGP Peer with MP-BGP for IPv4 Multicast

After you [Configure BGP](#), configure a BGP peer with MP-BGP for IPv4 multicast if you want your BGP peer to be able to learn and pass IPv4 multicast routes in BGP updates. You'll be able to separate unicast from multicast traffic, or employ the features listed in [MP-BGP](#) to use only routes from the unicast or multicast route table, or routes from both tables.

If you want to support multicast traffic only, you must use a filter to eliminate unicast traffic.

The firewall doesn't support ECMP for multicast traffic.

STEP 1 | Enable MP-BGP extensions so that a BGP peer can exchange IPv4 multicast routes.

1. Select **Network > Virtual Routers** and select the virtual router you are configuring.
2. Select **BGP**.
3. Select **Peer Group**, select a peer group and a BGP peer.
4. Select **Addressing**.
5. Select **Enable MP-BGP Extensions**.

-
6. For **Address Family Type**, select **IPv4**.
 7. For **Subsequent Address Family**, select **Unicast** and then **Multicast**.
 8. Click **OK**.

STEP 2 | (Optional) Create an IPv4 static route and install it in the multicast route table only.

You would do this to direct multicast traffic for a BGP peer to a specific next hop, as shown in the topology in [MP-BGP](#).

1. Select **Network > Virtual Routers** and select the virtual router you are configuring.
2. Select **Static Routes > IPv4** and **Add** a **Name** for the route.
3. Enter the IPv4 **Destination** prefix and netmask.
4. Select the egress **Interface**.
5. Select the **Next Hop** as **IP Address** and enter the IP address of the next hop to which you want to direct multicast traffic for this static route.
6. Enter an **Admin Distance**.
7. Enter a **Metric**.
8. For **Route Table**, select **Multicast**.
9. Click **OK**.

STEP 3 | Commit the configuration.

Click **Commit**.


STEP 4 | View the route table.

1. Select **Network > Virtual Routers**.
2. In the row for the virtual router, click **More Runtime Stats**.
3. Select **Routing > Route Table**.
4. For **Route Table**, select **Unicast** or **Multicast** to display only those routes.
5. For **Display Address Family**, select **IPv4 Only**, **IPv6 Only**, or **IPv4 and IPv6** to display only routes for that address family.

STEP 5 | To view the Forwarding table, BGP Local RIB, or BGP RIB Out table, see [Configure a BGP Peer with MP-BGP for IPv4 or IPv6 Unicast](#).

BGP Confederations

BGP confederations provide a way to divide an autonomous system (AS) into two or more sub-autonomous systems (sub-AS) to reduce the burden that the full mesh requirement for IBGP causes. The firewalls (or other routing devices) within a sub-AS must still have a full iBGP mesh with the other firewalls in the same sub-AS. You need BGP peering between sub-autonomous systems for full connectivity within the main AS. The firewalls peering with each other within a sub-AS form an IBGP confederation peering. The firewall in one sub-AS peering with a firewall in a different sub-AS form an EBGP confederation peering. Two firewalls from different autonomous systems that connect are EBGP peers.




Autonomous systems are identified with a public (globally-assigned) AS number, such as AS 24 and AS 25 in the preceding figure. In a PAN-OS environment, you assign each sub-AS a unique Confederation Member AS number, which is a private number seen only within the AS. In this figure, the confederations are AS 65100 and AS 65110. ([RFC6996](#), Autonomous System (AS) Reservation for Private Use, indicates that the IANA reserves AS numbers 64512-65534 for private use.)

The sub-AS confederations seem like full autonomous systems to each other within the AS. However, when the firewall sends an AS path to an EBGP peer, only the public AS number appears in the AS path; no private sub-AS (Confederation Member AS) numbers are included.

BGP peering occurs between the firewall and R2; the firewall in the figure has these relevant configuration settings:

- AS number—24
- Confederation Member AS—65100
- Peering Type—EBGP confed
- Peer AS—65110



Router 2 (R2) in AS 65110 is configured as follows:

- AS number—24
- Confederation Member AS—65110
- Peering Type—EBGP confed
- Peer AS—65100

BGP peering also occurs between the firewall and R1. The firewall has the following additional configuration:

- AS number—24
- Confederation Member AS—65100
- Peering Type—IBGP confed
- Peer AS—65110

R1 is configured as follows:

- AS number—24
- Confederation Member AS—65110
- Peering Type—IBGP confed
- Peer AS—65100


BGP peering occurs between the firewall and R5. The firewall has the following additional configuration:

- AS number—24
- Confederation Member AS—65100
- Peering Type—EBGP
- Peer AS—25



R5 is configured as follows:


- AS—25
- Peering Type—EBGP
- Peer AS—24

After the firewall is configured to peer with R1, R2, and R5, its peers are visible on the **Peer Group** tab:




The firewall shows the R1, R2, and R5 peers:






To verify that the routes from the firewall to the peers are established, on the virtual router's screen, select **More Runtime Stats** and select the **Peer** tab.



Select the **Local RIB** tab to view information about the routes stored in the Routing Information Base (RIB).



Then select the **RIB Out** tab.

Virtual Router - Default

Routing | BGP | OSPF | OSPFv3 | BGP | Multicast | BFD Summary Information

Summary | Peer | Peer Group | Local AS | All Out

0 items

Prefix	Next Hop	Peer	Local Pref	Adj. Path	Origin	ASID	Adm. Status	App. Status
15.1.1.0/24	101.1.1.11	R1		20			advertised	no aggregate
15.1.1.0/24	15.1.1.5	R2		(01110)			advertised	no aggregate
15.1.1.0/24	101.1.1.1	R5		24			advertised	no aggregate
15.1.1.0/24	101.1.1.1	R5		24			advertised	no aggregate
15.1.1.0/24	101.1.1.11	R2		(01100),20			advertised	no aggregate
15.1.1.0/24	101.1.1.11	R2		(01100)			advertised	no aggregate

Done

IP Multicast

IP multicast is a set of protocols that network appliances use to send multicast IP datagrams to a group of interested receivers using one transmission rather than unicasting the traffic to multiple receivers, thereby saving bandwidth. IP multicast is suitable for communication from one source (or many sources) to many receivers, such as audio or video streaming, IPTV, video conferencing, and distribution of other communication, such as news and financial data.

A multicast address identifies a group of receivers that want to receive the traffic going to that address. You should not use the multicast addresses reserved for special uses, such as the range 224.0.0.0 through 224.0.0.255 or 239.0.0.0 through 239.255.255.255. Multicast traffic uses UDP, which does not resend missed packets.

Palo Alto Networks® firewalls support IP multicast and Protocol Independent Multicast (PIM) on a Layer 3 interface that you configure for a [virtual router](#) on the firewall.

For multicast routing, the Layer 3 interface type can be Ethernet, Aggregate Ethernet (AE), VLAN, loopback, or tunnel. Interface groups allow you to configure more than one firewall interface at a time with the same Internet Group Management Protocol (IGMP) and PIM parameters, and with the same group permissions (multicast groups allowed to accept traffic from any source or from only a specific source). An interface can belong to only one interface group.

The firewall supports IPv4 multicast—it does not support IPv6 multicast. The firewall also does not support PIM Dense Mode (PIM-DM), IGMP proxy, IGMP static joins, Anycast RP, GRE, or multicast configurations on a Layer 2 or virtual wire interface type. However, a virtual wire interface can pass multicast packets. Also, a Layer 2 interface can switch Layer 3 IPv4 multicast packets between different VLANs and the firewall will retag the VLAN ID using the VLAN ID of the egress interface.

You must enable multicast for a virtual router and enable PIM for an ingress and an egress interface in order for the interfaces to receive or forward multicast packets. In addition to PIM, you must also enable IGMP on egress interfaces that face receivers. You must configure a Security policy rule to allow IP multicast traffic to a predefined Layer 3 destination zone named **multicast** or to **any** destination zone.

- [IGMP](#)
- [PIM](#)
- [Configure IP Multicast](#)
- [View IP Multicast Information](#)

IGMP

Internet Group Management Protocol (IGMP) is an IPv4 protocol that a multicast receiver uses to communicate with an interface on a Palo Alto Networks® firewall and that the firewall uses to track the membership of multicast groups. When a host wants to receive multicast traffic, its implementation of IGMP sends an IGMP Membership report message and the receiving router, in turn, sends a PIM Join message to the multicast group address of the group that the host wants to join. An IGMP-enabled router on the same physical network (such as an Ethernet segment) then uses PIM to communicate with other PIM-enabled routers to determine a path from the source to interested receivers.

Enable IGMP only on interfaces that face a multicast receiver. The receivers can be only one Layer 3 hop away from the virtual router. IGMP messages are Layer 2 messages that have a TTL value of one and, therefore, cannot go outside the LAN.

When you [Configure IP Multicast](#), specify whether an interface uses [IGMP Version 1](#), [IGMP Version 2](#), or [IGMP Version 3](#). You can enforce the IP Router Alert option, [RFC 2113](#), so that incoming IGMP packets that use IGMPv2 or IGMPv3 have the IP Router Alert option.

By default, an interface accepts IGMP Membership reports for all multicast groups. You can configure multicast group permissions to control the groups for which the virtual router accepts Membership reports from any source (Any-Source Multicast, or ASM), which is basically PIM Sparse Mode (PIM-SM). You can also specify the groups for which the virtual router accepts Membership reports from a specific source (PIM Source-Specific Multicast [PIM-SSM]). If you specify permissions for either ASM or SSM groups, the virtual router denies Membership reports from other groups. The interface must use IGMPv3 to pass PIM-SSM traffic.

You can specify the maximum number of sources and the maximum number of multicast groups that IGMP can process simultaneously for an interface.

The virtual router multicasts an IGMP Query at regular intervals to all receivers of a multicast group. A receiver responds to an IGMP Query with an IGMP Membership report that confirms the receiver still wants to receive multicast traffic for that group. The virtual router maintains a table of the multicast groups that have receivers; the virtual router forwards a multicast packet out the interface to the next hop only if there is still a receiver down that multicast distribution tree that is joined to the group. The virtual router does not track exactly which receivers are joined to a group. Only one router on a subnet responds to IGMP Queries and that is the IGMP Querier—the router with the lowest IP address.

You can configure an interface with an IGMP Query interval and the amount of time allowed for a receiver to respond to a query (the Max Query Response Time). When a virtual router receives an IGMP Leave message from a receiver to leave a group, the virtual router checks that the interface that received the Leave message is not configured with the Immediate Leave option. In the absence of the Immediate Leave option, the virtual router sends a Query to determine whether there are still receiver members for the group. The Last Member Query Interval specifies how many seconds are allowed for any remaining receivers for that group to respond and confirm that they still want multicast traffic for that group.

An interface supports the IGMP robustness variable, which you can adjust so that the firewall then tunes the Group Membership Interval, Other Querier Present Interval, Startup Query Count, and Last Member Query Count. A higher robustness variable can accommodate a subnet that is likely to drop packets.

[View IP Multicast Information](#) to see IGMP-enabled interfaces, the IGMP version, Querier address, robustness setting, limits on the number of multicast groups and sources, and whether the interface is configured for Immediate Leave. You can also see the multicast groups to which interfaces belong and other IGMP membership information.

PIM

IP multicast uses the Protocol Independent Multicast (PIM) routing protocol between routers to determine the path on the distribution tree that multicast packets take from the source to the receivers (multicast group members). A Palo Alto Networks® firewall supports PIM Sparse Mode (PIM-SM) ([RFC 4601](#)), PIM Any-Source Multicast (ASM) (sometimes referred to as PIM Sparse Mode), and PIM Source-Specific Multicast (SSM). In PIM-SM, the source does not forward multicast traffic until a receiver (user) belonging to a multicast group requests that the source send the traffic. When a host wants to receive multicast traffic, its implementation of IGMP sends an IGMP Membership report message, and the receiving router then sends a PIM Join message to the multicast group address of the group it wants to join.


- In **ASM**, the receiver uses IGMP to request traffic for a multicast group address; any source could have originated that traffic. Consequently, the receiver doesn't necessarily know the senders, and the receiver could receive multicast traffic in which it has no interest.
- In **SSM** ([RFC 4607](#)), the receiver uses IGMP to request traffic from one or more specific sources to a multicast group address. The receiver knows the IP address of the senders and receives only the multicast traffic it wants. SSM requires IGMPv3. You can override the default SSM address space, which is 232.0.0.0/8.

When you [Configure IP Multicast](#) on a Palo Alto Networks® firewall, you must enable PIM for an interface to forward multicast traffic, even on receiver-facing interfaces. This is unlike IGMP, which you enable only on receiver-facing interfaces.

ASM requires a *rendezvous point* (RP), which is a router located at the juncture or root of a shared distribution tree. The RP for a multicast domain serves as a single point to which all multicast group members send their Join messages. This behavior reduces the likelihood of a routing loop that would otherwise occur if group members sent their Join messages to multiple routers. (SSM doesn't need an RP because source-specific multicast uses a shortest-path tree and therefore has no need for an RP.)


In an ASM environment, there are two ways that the virtual router determines which router is the RP for a multicast group:

- **Static RP-to-Group Mapping**—configures the virtual router on the firewall to act as RP for multicast groups. You configure a local RP, either by configuring a static RP address or by specifying that the local RP is a candidate RP and the RP is chosen dynamically (based on lowest priority value). You can also statically configure one or more external RPs for different group address ranges not covered by the local RP, which helps you load-balance multicast traffic so that one RP is not overloaded.
- **Bootstrap Router (BSR)**—(RFC 5059)—defines the role of a BSR. First, candidates for BSR advertise their priority to each other and then the candidate with the largest priority is elected BSR, as shown in the following figure:



Next, the BSR discovers RPs when candidate RPs periodically unicast a BSR message to the BSR containing their IP address and the multicast group range for which they will act as RP. You can configure the local virtual router to be a candidate RP, in which case the virtual router announces its RP candidacy for a specific multicast group or groups. The BSR sends out RP information to the other RPs in the PIM domain.

When you configure PIM for an interface, you can select BSR Border when the interface on the firewall is at an enterprise boundary facing away from the enterprise network. The BSR Border setting prevents the firewall from sending RP candidacy BSR messages outside the LAN. In the following illustration, BSR Border is enabled for the interface facing the LAN and that interface has the highest priority. If the virtual router has both a static RP and a dynamic RP (learned from the BSR), you can specify whether the static RP should override the learned RP for a group when you configure the local, static RP.



In order for PIM Sparse Mode to notify the RP that it has traffic to send down a shared tree, the RP must be aware of the source. The host notifies the RP that it is sending traffic to a multicast group address when the *designated router* (DR) encapsulates the first packet from the host in a PIM Register message and unicasts the packet to the RP on its local network. The DR also forwards Prune messages from a receiver to the RP. The RP maintains the list of IP addresses of sources that are sending to a multicast group and the RP can forward multicast packets from sources.

Why do the routers in a PIM domain need a DR? When a router sends a PIM Join message to a switch, two routers could receive it and forward it to the same RP, causing redundant traffic and wasting bandwidth. To prevent unnecessary traffic, the PIM routers elect a DR (the router with the highest IP address), and only the DR forwards the Join message to the RP. Alternatively, you can assign a DR priority to an interface group, which takes precedence over IP address comparisons. As a reminder, the DR is forwarding (unicasting) PIM messages; it is not multicasting IP multicast packets.


You can specify the IP addresses of PIM neighbors (routers) that the interface group will allow to peer with the virtual router. By default, all PIM-enabled routers can be PIM neighbors, but the option to limit neighbors provides a step toward securing the virtual router in your PIM environment.

- [Shortest Path Tree \(SPT\) and Shared Tree](#)
- [PIM Assert Mechanism](#)
- [Reverse-Path Forwarding](#)


Shortest-Path Tree (SPT) and Shared Tree

After a receiver joins a multicast group, the routers in the multiaccess network build the routing paths necessary to send data to each receiver in the group. Each IP datagram sent to a multicast group is distributed (forwarded) to all members. The routing paths constitute a type of distribution tree for a multicast packet. The goal of a multicast distribution tree is for the router to duplicate a multicast packet when the packet reaches a divergence of paths and the router must send the packet down multiple paths to reach all group members, yet the distribution tree must refrain from sending packets down a path where no interested receivers exist. The distribution tree is one of the following:

- A **source tree**—A path from a multicast source (the root of the tree) through the network to the receivers in the multicast group. The source tree is the shortest path that a multicast packet can take from source to receiver, so it is also known as the **shortest-path tree (SPT)**. The sender and receiver are annotated as a source and multicast group pair, shortened to (S, G); for example, (192.168.1.1, 225.9.2.6). The following figure illustrates three shortest-path trees from the source to three receivers.



- A **shared tree**—A path rooted at the RP, not at the multicast source. A shared tree is also known as an RP tree or RPT. Routers forward multicast packets from various sources to the RP and the RP forwards the packets down the shared tree. A shared tree is annotated as (*, G), using a wildcard as the source because all sources belonging to the multicast group share the same distribution tree from the RP. An example shared tree annotation is (*, 226.3.1.5). The following figure illustrates a shared tree from the root at the RP to the receivers.



Source-Specific Multicast (SSM) uses source tree distribution. When you [Configure IP Multicast](#) to use Any Source Multicast (ASM), you can specify which distribution tree the virtual router on your Palo Alto Networks® firewall uses to deliver multicast packets to a group by setting an SPT threshold for the group:

- By default the virtual router switches multicast routing from shared tree to SPT when it receives the first multicast packet for a group or prefix (the **SPT Threshold** is set to 0).
- You can configure the virtual router to switch to SPT when the total number of kilobits in packets arriving for the specified multicast group or prefix at any interface over any length of time reaches a configured number.
- You can configure the virtual router to never switch to SPT for the group or prefix (it continues to use shared tree).

SPT requires more memory, so choose your setting based on your multicast traffic level to the group. If the virtual router switches to SPT, then packets will arrive from the source (rather than the RP) and the virtual

router sends a Prune message to the RP. The source sends subsequent multicast packets for that group down the shortest-path tree.

PIM Assert Mechanism

To prevent routers on a multiaccess network from forwarding the same multicast traffic to the same next hop (which would cause redundant traffic and wasted bandwidth), PIM uses the Assert mechanism to elect a single PIM Forwarder for the multiaccess network.

If the virtual router receives a multicast packet from a source on an interface that the virtual router already associates as the outgoing interface for the same (S,G) pair identified in the packet, that means this is a duplicate packet. Consequently, the virtual router sends an Assert message containing its metrics to the other routers on the multiaccess network. The routers then elect a PIM Forwarder in this manner:

1. The PIM Forwarder is the router with the lowest administrative distance to the multicast source.
2. In the event of a tie for lowest administrative distance, the PIM Forwarder is the router with the best unicast routing metric to the source.
3. In the event of a tie for best metric, the PIM Forwarder is the router with the highest IP address.

Routers that are not elected as the PIM Forwarder will stop forwarding traffic to the multicast group identified in the (S,G) pair.

When you [Configure IP Multicast](#), you can configure the interval at which the virtual router sends PIM Assert messages out an interface (the Assert interval). When you [View IP Multicast Information](#), the **PIM Interface** tab displays the Assert interval for an interface.

Reverse-Path Forwarding

PIM uses reverse-path forwarding (RPF) to prevent multicast routing loops by leveraging the unicast routing table on the virtual router. When the virtual router receives a multicast packet, it looks up the source of the multicast packet in its unicast routing table to see if the outgoing interface associated with that source IP address is the interface on which that packet arrived. If the interfaces match, the virtual router duplicates the packet and forwards it out the interfaces toward the multicast receivers in the group. If the interfaces don't match, the virtual router drops the packet. The unicast routing table is based on the underlying static routes or the interior gateway protocol (IGP) your network uses, such as OSPF.

PIM also uses RPF to build a [shortest-path tree](#) to a source, one PIM router hop at a time. The virtual router has the address of the multicast source, so the virtual router selects as its next hop back to the source the upstream PIM neighbor that the virtual router would use to forward unicast packets to the source. The next hop router does the same thing.

After RPF succeeds and the virtual router has a route entry in its multicast routing information base (mRIB), the virtual router maintains source-based tree entries (S,G) and shared tree entries (*,G) in its multicast forwarding information base (multicast forwarding table or mFIB). Each entry includes the source IP address, multicast group, incoming interface (RPF interface) and outgoing interface list. There can be multiple outgoing interfaces for an entry because the shortest path tree can branch at the router, and the router must forward the packet out multiple interfaces to reach receivers of the group that are located down different paths. When the virtual router uses the mFIB to forward a multicast packet, it matches an (S,G) entry before it attempts to match a (*,G) entry.

If you are advertising multicast source prefixes into BGP (you configured [MP-BGP](#) with the IPv4 Address Family and the multicast Subsequent Address Family), then the firewall always performs the RPF check on the BGP routes that the firewall received under the multicast Subsequent Address Family.

[View IP Multicast Information](#) to see how to view the mFIB and mRIB entries. Keep in mind that the multicast route table (mRIB) is a separate table from the unicast route table (RIB).

Configure IP Multicast


Configure interfaces on a virtual router of a Palo Alto Networks® firewall to receive and forward IP multicast packets. You must enable IP multicast for the virtual router, configure Protocol Independent Multicast (PIM) on the ingress and egress interfaces, and configure Internet Group Management Protocol (IGMP) on receiver-facing interfaces.

STEP 1 | Enable IP multicast for a virtual router.

1. Select **Network > Virtual Routers** and select a virtual router.
2. Select **Multicast** and **Enable IP multicast**.

STEP 2 | (ASM only) If the multicast domain in which the virtual router is located uses Any-Source Multicast (ASM), identify and configure the local and remote rendezvous points (RPs) for multicast groups.

1. Select **Rendezvous Point**.
2. Select a Local **RP Type**, which determines how the RP is chosen (the options are **Static**, **Candidate**, or **None**):
 - **Static**—Establishes a static mapping of an RP to multicast groups. Configuring a static RP requires you to explicitly configure the same RP on other PIM routers in the PIM domain.
 - Select the **RP Interface**. Valid interface types are Layer3, virtual wire, loopback, VLAN, Aggregate Ethernet (AE), and tunnel.
 - Select the **RP Address**. The IP addresses of the RP interface you selected populate the list.
 - Select **Override learned RP for the same group** so that this static RP serves as RP instead of the RP elected for the groups in the Group List.
 - **Add** one or more multicast **Groups** for which the RP acts as the RP.



- **Candidate**—Establishes a dynamic mapping of an RP to multicast groups based on priority so that each router in a PIM domain automatically elects the same RP.
 - Select the **RP Interface** of the candidate RP. Valid interface types are Layer 3, loopback, VLAN, Aggregate Ethernet (AE), and tunnel.
 - Select the **RP Address** of the candidate RP. The IP addresses for the RP interface you selected populate the list.

- (Optional) Change the **Priority** for the candidate RP. The firewall compares the priority of the candidate RP to the priority of other candidate RPs to determine which one acts as RP for the specified groups; the firewall selects the candidate RP with the lowest priority value (range is 0 to 255; default is 192).
 - (Optional) Change the **Advertisement Interval (sec)** (range is 1 to 26,214; default is 60).
 - Enter a **Group List** of multicast groups that communicate with the RP.
 - **None**—Select if this virtual router is not an RP.
3. **Add** a Remote Rendezvous Point and enter the **IP Address** of that remote (external) RP.
 4. **Add** the multicast **Group Addresses** for which the specified remote RP address acts as RP.
 5. Select **Override learned RP for the same group** so that the external RP you configured statically serves as RP instead of an RP that is dynamically learned (elected) for the groups in the Group Addresses list.
 6. Click **OK**.

STEP 3 | Specify a group of interfaces that share a multicast configuration (IGMP, PIM, and group permissions).

1. On the **Interfaces** tab, **Add a Name** for the interface group.
2. Enter a **Description**.
3. **Add an Interface** and select one or more Layer 3 interfaces that belong to the interface group.


STEP 4 | (Optional) Configure multicast group permissions for the interface group. By default, the interface group accepts IGMP membership reports and PIM join messages from all groups.

1. Select **Group Permissions**.
2. To configure Any-Source Multicast (ASM) groups for this interface group, in the Any Source window, **Add a Name** to identify a multicast group that accepts IGMP membership reports and PIM join messages from any source.
3. Enter the multicast **Group** address or group address and /prefix that can receive multicast packets from any source on these interfaces.
4. Select **Included** to include the ASM **Group** address in the interface group (default). De-select **Included** to easily exclude an ASM group from the interface group, such as during testing.
5. **Add** additional multicast **Groups** (for the interface group) that want to receive multicast packets from any source.
6. To configure Source-Specific Multicast (SSM) groups in this interface group, in the Source Specific window, **Add a Name** to identify a multicast group and source address pair. Don't use a name that you used for Any Source multicast. (You must use IGMPv3 to configure SSM.)
7. Enter the multicast **Group** address or group address and /prefix of the group that wants to receive multicast packets from the specified source only (and can receive the packets on these interfaces).



A Source Specific group for which you specify permissions is a group that the virtual router must treat as source-specific. Configure Source Specific Address Space (Step 9) that includes the source-specific groups for which you configured permission.

8. Enter the **Source** IP address from which this multicast group can receive multicast packets.
9. Select **Included** to include the SSM Group and source address pair in the interface group (default). De-select **Included** to easily exclude the pair from the interface group, such as during testing.
10. **Add** additional multicast **Groups** (for the interface group) that receive multicast packets from a specific source only.



STEP 5 | Configure IGMP for the interface group if an interface faces multicast receivers, which must use IGMP to join a group.

1. On the **IGMP** tab, **Enable IGMP** (default).
2. Specify **IGMP** parameters for interfaces in the interface group:
 - **IGMP Version**—1, 2, or 3 (default).
 - **Enforce Router-Alert IP Option** (disabled by default)—Select this option if you require incoming IGMP packets that use IGMPv2 or IGMPv3 to have the [IP Router Alert Option](#), RFC 2113.
 - **Robustness**—A variable that the firewall uses to tune the Group Membership Interval, Other Querier Present Interval, Startup Query Count, and Last Member Query Count (range is 1 to 7; default is 2). Increase the value if the subnet on which this firewall is located is prone to losing packets.
 - **Max Sources**—Maximum number of sources that IGMP can process simultaneously for an interface (range is 1 to 65,535; default is **unlimited**).
 - **Max Groups**—Maximum number of groups that IGMP can process simultaneously for an interface (range is 1 to 65,535; default is **unlimited**).
 - **Query Interval**—Number of seconds between IGMP membership Query messages that the virtual router sends to a receiver to determine whether the receiver still wants to receive the multicast packets for a group (range is 1 to 31,744; default is 125).
 - **Max Query Response Time (sec)**—Maximum number of seconds allowed for a receiver to respond to an IGMP membership Query message before the virtual router determines that the receiver no longer wants to receive multicast packets for the group (range is 0 to 3,174.4; default is 10).
 - **Last Member Query Interval (sec)**—Number of seconds allowed for a receiver to respond to a Group-Specific Query that the virtual router sends after a receiver sends a Leave Group message (range is 0.1 to 3,174.4; default is 1).
 - **Immediate Leave** (disabled by default)—When there is only one member in a multicast group and the virtual router receives an IGMP Leave message for that group, the Immediate Leave setting causes the virtual router to remove that group and outgoing interface from the multicast routing information base (mRIB) and multicast forwarding information base (mFIB) immediately, rather than waiting for the Last Member Query Interval to expire. The Immediate Leave setting saves network resources. You cannot select Immediate Leave if the interface group uses IGMPv1.

STEP 6 | Configure PIM Sparse Mode (PIM-SM) for the interface group.

1. On the **PIM** tab, **Enable PIM** (enabled by default).

2. Specify PIM parameters for the interface group:

- **Assert Interval**—Number of seconds between [PIM Assert messages](#) that the virtual router sends to other PIM routers on the multiaccess network when they are electing a PIM forwarder (range is 0 to 65,534; default is 177).
- **Hello Interval**—Number of seconds between PIM Hello messages that the virtual router sends to its PIM neighbors from each interface in the interface group (range is 0 to 18,000; default is 30).
- **Join Prune Interval**—Number of seconds between PIM Join messages (and between PIM Prune messages) that the virtual router sends upstream toward a multicast source (range is 0 to 18,000; default is 60).
- **DR Priority**—Designated Router (DR) priority that controls which router in a multiaccess network forwards PIM Join and Prune messages to the RP (range is 0 to 429,467,295; default is 1). The DR priority takes precedence over IP address comparisons to elect the DR.
- **BSR Border**—Select this option if the interfaces in the interface group are on a virtual router that is the BSR located at the border of an enterprise LAN. This will prevent RP candidacy BSR messages from leaving the LAN.

3. **Add** one or more **Permitted PIM Neighbors** by specifying the **IP Address** of each router from which the virtual router accepts multicast packets.

STEP 7 | Click **OK** to save the interface group settings.

STEP 8 | (Optional) Change the Shortest-Path Tree (SPT) threshold, as described in [Shortest-Path Tree \(SPT\) and Shared Tree](#).

1. Select **SPT Threshold** and **Add a Multicast Group/Prefix**, the multicast group or prefix for which you are specifying the distribution tree.
2. Specify the **Threshold (kb)**—The point at which routing to the specified multicast group or prefix switches from shared tree (sourced from the RP) to SPT distribution:
 - **0 (switch on first data packet)** (default)—The virtual router switches from shared tree to SPT for the group or prefix when the virtual router receives the first data packet for the group or prefix.
 - **never (do not switch to spt)**—The virtual router continues to use the shared tree to forward packets to the group or prefix.
 - Enter the total number of kilobits from multicast packets that can arrive for the multicast group or prefix at any interface and over any time period, upon which the virtual router changes to SPT distribution for that multicast group or prefix.

STEP 9 | Identify the multicast groups or groups and prefixes that accept multicast packets only from a specific source.

1. Select **Source Specific Address Space** and **Add the Name** for the space.
2. Enter the multicast **Group** address with prefix length to identify the address space that receives multicast packets from a specific source. If the virtual router receives a multicast packet for an SSM group but the group is not covered by a **Source Specific Address Space**, the virtual router drops the packet.
3. Select **Included** to include the source-specific address space as a multicast group address range from which the virtual router will accept multicast packets that originated from an allowed specific source. De-select **Included** to easily exclude a group address space for testing.
4. Add other source-specific address spaces to include all those groups for which you specified SSM group permission.



STEP 10 | (Optional) Change the length of time that a multicast route remains in the mRIB after the session ends between a multicast group and a source.

1. Select the **Advanced** tab.
2. Specify the **Multicast Route Age Out Time (sec)** (range is 210 to 7,200; default is 210).

STEP 11 | Click **OK** to save the multicast configuration.

STEP 12 | Create a Security policy rule to allow multicast traffic to the destination zone.

1. [Create a Security Policy Rule](#) and on the **Destination** tab, select **multicast** or **any** for the **Destination Zone**. The **multicast** zone is a predefined Layer 3 zone that matches all multicast traffic. The **Destination Address** can be a multicast group address.
2. Configure the rest of the Security policy rule.

STEP 13 | (Optional) Enable buffering of multicast packets before a route is set up.

1. Select **Device > Setup > Session** and edit Session Settings.
2. Enable **Multicast Route Setup Buffering** (disabled by default). The firewall can preserve the first packet(s) from a multicast flow if an entry for the corresponding multicast group does not yet exist in the multicast forwarding table (mFIB). The **Buffer Size** controls how many packets the firewall buffers from a flow. After the route is installed in the mFIB, the firewall automatically forwards the buffered first packet(s) to the receiver. (You need to enable multicast route setup buffering only if your content servers are directly connected to the firewall and your multicast application cannot withstand the first packet of the flow being dropped.)
3. (Optional) Change the **Buffer Size**. Buffer size is the number of packets per multicast flow that the firewall can buffer until the mFIB entry is set up (range is 1 to 2,000; default is 1,000). The firewall can buffer a maximum of 5,000 packets total (for all flows).
4. Click **OK**.

STEP 14 | **Commit** your changes.

STEP 15 | [View IP Multicast Information](#) to view mRIB and mFIB entries, IGMP interface settings, IGMP group memberships, PIM ASM and SSM modes, group mappings to RPs, DR addresses, PIM settings, PIM neighbors, and more.

STEP 16 | If you [Configure a Static Route](#) for multicast traffic, you can install the route only in the multicast routing table (not the unicast routing table) so that the route is used for multicast traffic only.


STEP 17 | If you enable IP multicast, it is not necessary to [Configure BGP with MP-BGP for IPv4 Multicast](#) unless you have a logical multicast topology separate from a logical unicast topology. You configure MP-BGP extensions with the IPv4 address family and multicast

subsequent address family only when you want to advertise multicast source prefixes into BGP under multicast subsequent address family.

View IP Multicast Information

After you [Configure IP Multicast](#) routing, view multicast routes, forwarding entries, and information about your IGMP and PIM interfaces.

- Select **Network > Virtual Routers** and in the row for the virtual router you configured, click **More Runtime Stats**.
 1. Select **Routing > Route Table** and then the **Multicast** radio button to display only multicast routes (destination IP multicast group, the next hop toward that group, and outgoing interface). This information comes from the mRIB.
 2. Select **Multicast > FIB** to view multicast route information from the mFIB: multicast groups to which the virtual router belongs, the corresponding source, incoming interfaces, and outgoing interfaces toward the receivers.




3. Select **Multicast > IGMP > Interface** to view IGMP-enabled interfaces, the associated IGMP version, IP address of the IGMP Querier, Querier up time and expiry time, the robustness setting, limits on numbers of multicast groups and sources, and whether the interface is configured for immediate leave.




4. Select **Multicast > IGMP > Membership** to see IGMP-enabled interfaces and the multicast groups to which they belong, the source, and other IGMP information.




5. Select **Multicast > PIM > Group Mapping** to view multicast groups mapped to an RP, the origin of the RP mapping, the PIM mode for the group (ASM or SSM), and whether the group is inactive. Groups in SSM mode don't use an RP, so the RP address displayed is 0.0.0.0. The default SSM group is 232.0.0.0/8.



6. Select **Multicast > PIM > Interface** to view the IP address of the DR for an interface; the DR priority; the Hello, Join/Prune, and Assert intervals; and whether the interface is a bootstrap router (BSR).



7. Select **Multicast > PIM > Neighbor** to view information about routers that are PIM neighbors to the virtual router.



Route Redistribution

Route redistribution on the firewall is the process of making routes that the firewall learned from one routing protocol (or a static or connected route) available to a different routing protocol, thereby increasing accessibility of network traffic. Without route redistribution, a router or virtual router advertises and shares routes only with other routers that run the same routing protocol. You can redistribute IPv4 or IPv6 BGP, connected, or static routes into the OSPF RIB and redistribute OSPFv3, connected, or static routes into the BGP RIB.

This means, for example, you can make specific networks that were once available only by manual static route configuration on specific routers available to BGP autonomous systems or OSPF areas. You can also advertise locally connected routes, such as routes to a private lab network, into BGP autonomous systems or OSPF areas.

You might want to give users on your internal OSPFv3 network access to BGP so they can access devices on the internet. In this case you would redistribute BGP routes into the OSPFv3 RIB.

Conversely, you might want to give your external users access to some parts of your internal network, so you make internal OSPFv3 networks available through BGP by redistributing OSPFv3 routes into the BGP RIB.

STEP 1 | Create a redistribution profile.

1. Select **Network > Virtual Routers** and select a virtual router.
2. Select **Redistribution Profile** and **IPv4** or **IPv6** and **Add** a profile.
3. Enter a **Name** for the profile, which must start with an alphanumeric character and can contain zero or more underscores (_), hyphens (-), dots (.), or spaces (up to 16 characters).
4. Enter a **Priority** for the profile in the range 1 to 255. The firewall matches routes to profiles in order using the profile with the highest priority (lowest priority value) first. Higher priority rules take precedence over lower priority rules.
5. For **Redistribute**, select one of the following:
 - **Redist**—Select for redistribution the routes that match this filter.
 - **No Redist**—Select for redistribution routes that match the redistribution profiles except the routes that match this filter. This selection treats the profile as a block list that specifies which routes not to select for redistribution. For example, if you have multiple redistribution profiles for BGP, you can create a **No Redist** profile to exclude several prefixes, and then a general redistribution profile with a lower priority (higher priority value) after it. The two profiles combine and the higher priority profile takes precedence. You can't have only **No Redist** profiles; you would always need at least one **Redist** profile to redistribute routes.
6. On the **General Filter** tab, for Source Type, select one or more types of route to redistribute:
 - **bgp**—Redistribute BGP routes that match the profile.
 - **connect**—Redistribute connected routes that match the profile.
 - **ospf (IPv4 only)**—Redistribute OSPF routes that match the profile.
 - **rip (IPv4 only)**—Redistribute RIP routes that match the profile.
 - **ospfv3 (IPv6 only)**—Redistribute OSPFv3 routes that match the profile.
 - **static**—Redistribute static routes that match the profile.
7. (Optional) For **Interface**, **Add** one or more egress interfaces of associated routes to match for redistribution. To remove an entry, click **Delete**.
8. (Optional) For **Destination**, **Add** one or more IPv4 or IPv6 destinations of routes to match for redistribution. To remove an entry, click **Delete**.

-
9. (Optional) For **Next Hop, Add** one or more next hop IPv4 or IPv6 addresses of routes to match for redistribution. To remove an entry, click **Delete**.
 10. Click **OK**.

STEP 2 | (Optional—When General Filter includes *ospf* or *ospfv3*) Create an OSPF filter to further specify which OSPF or OSPFv3 routes to redistribute.

1. Select **Network > Virtual Routers** and select the virtual router.
2. Select **Redistribution Profile** and **IPv4** or **IPv6** and select the profile you created.
3. Select **OSPF Filter**.
4. For **Path Type**, select one or more of the following types of OSPF path to redistribute: **ext-1**, **ext-2**, **inter-area**, or **intra-area**.
5. To specify an **Area** from which to redistribute OSPF or OSPFv3 routes, **Add** an area in IP address format.
6. To specify a **Tag**, **Add** a tag in IP address format.
7. Click **OK**.

STEP 3 | (Optional—When General Filter includes *bgp*) Create a BGP filter to further specify which BGP routes to redistribute.

1. Select **Network > Virtual Routers** and select the virtual router.
2. Select **Redistribution Profile** and **IPv4** or **IPv6** and select the profile you created.
3. Select **BGP Filter**.
4. For **Community**, **Add** to select from the list of communities, such as well-known communities: **local-as**, **no-advertise**, **no-export**, or **nopeer**. You can also enter a 32-bit value in decimal or hexadecimal or in AS:VAL format, where AS and VAL are each in the range 0 to 65,535. Enter a maximum of 10 entries.
5. For **Extended Community**, **Add** an extended community as a 64-bit value in hexadecimal or in TYPE:AS:VAL or TYPE:IP:VAL format. TYPE is 16 bits; AS or IP is 16 bits; VAL is 32 bits. Enter a maximum of five entries.
6. Click **OK**.

STEP 4 | Select the protocol into which you are redistributing routes, and set the attributes for those routes.

This task illustrates redistributing routes into BGP.

1. Select **Network > Virtual Routers** and select the virtual router.
2. Select **BGP > Redist Rules**.
3. Select **Allow Redistribute Default Route** to allow the firewall to redistribute the default route.
4. Click **Add**.
5. Select **Address Family Type: IPv4** or **IPv6** to specify in which route table the redistributed routes will be put.
6. Select the **Name** of the Redistribution profile you created, which selects the routes to redistribute.
7. **Enable** the redistribution rule.
8. (Optional) Enter any of the following values, which the firewall applies to the routes being redistributed:
 - **Metric** in the range 1 to 65,535.
 - **Set Origin**—Origin of the route: **igp**, **egp**, or **incomplete**.
 - **Set MED**—MED value in the range 0 to 4,294,967,295.
 - **Set Local Preference**—Local preference value in the range 0 to 4,294,967,295.
 - **Set AS Path Limit**—Maximum number of autonomous systems in the AS_PATH in the range 1 to 255.

-
- **Set Community**—Select or enter a 32-bit value in decimal or hexadecimal, or enter a value in AS:VAL format, where AS and VAL are each in the range 0 to 65,525. Enter a maximum of 10 entries.
 - **Set Extended Community**—Select or enter an extended community as a 64-bit value in hexadecimal or in TYPE:AS:VAL or TYPE:IP:VAL format. TYPE is 16 bits; AS or IP is 16 bits; VAL is 32 bits. Enter a maximum of five entries.
9. Click **OK**.

STEP 5 | Commit your changes.

GRE Tunnels

The Generic Routing Encapsulation (GRE) tunnel protocol is a carrier protocol that encapsulates a payload protocol. The GRE packet itself is encapsulated in a transport protocol (IPv4 or IPv6).


- [GRE Tunnel Overview](#)
- [Create a GRE Tunnel](#)

GRE Tunnel Overview

A Generic Routing Encapsulation (GRE) tunnel connects two endpoints (a firewall and another appliance) in a point-to-point, logical link. The firewall can terminate GRE tunnels; you can route or forward packets to a GRE tunnel. GRE tunnels are simple to use and often the tunneling protocol of choice for point-to-point connectivity, especially to services in the cloud or to partner networks.

[Create a GRE tunnel](#) when you want to direct packets that are destined for an IP address to take a certain point-to-point path, for example to a cloud-based proxy or to a partner network. The packets travel through the GRE tunnel (over a transit network such as the internet) to the cloud service while on their way to the destination address. This enables the cloud service to enforce its services or policies on the packets.

The following figure is an example of a GRE tunnel connecting the firewall across the internet to a cloud service.



For better performance and to avoid single points of failure, split multiple connections to the firewall among multiple GRE tunnels rather than use a single tunnel. Each GRE tunnel needs a tunnel interface.

When the firewall allows a packet to pass (based on a policy match) and the packet egresses to a GRE tunnel interface, the firewall adds GRE encapsulation; it doesn't generate a session. The firewall does not perform a Security policy rule lookup for the GRE-encapsulated traffic, so you don't need a Security policy rule for the GRE traffic that the firewall encapsulates. However, when the firewall receives GRE traffic, it generates a session and applies all policies to the GRE IP header in addition to the encapsulated traffic. The firewall treats the received GRE packet like any other packet. Therefore:

- If the firewall receives the GRE packet on an interface that has the same zone as the tunnel interface associated with the GRE tunnel (for example, tunnel.1), the source zone is the same as the destination zone. By default, traffic is allowed within a zone (intrazone traffic), so the ingress GRE traffic is allowed by default.
- However, if you configured your own intrazone Security policy rule to deny such traffic, you must explicitly allow GRE traffic.

- Likewise, if the zone of the tunnel interface associated with the GRE tunnel (for example, tunnel.1) is a different zone from that of the ingress interface, you must configure a Security policy rule to allow the GRE traffic.

Because the firewall encapsulates the tunneled packet in a GRE packet, the additional 24 bytes of GRE header automatically result in a smaller [Maximum Segment Size \(MSS\)](#) in the maximum transmission unit (MTU). If you don't change the IPv4 MSS Adjustment Size for the interface, the firewall reduces the MTU by 64 bytes by default (40 bytes of IP header + 24 bytes of GRE header). This means if the default MTU is 1,500 bytes, the MSS will be 1,436 bytes ($1,500 - 40 - 24 = 1,436$). If you configure an MSS Adjustment Size of 300 bytes, for example, the MSS will be only 1,176 bytes ($1,500 - 300 - 24 = 1,176$).

The firewall does not support routing a GRE or IPSec tunnel to a GRE tunnel, but you can route a GRE tunnel to an IPSec tunnel. Additionally:

- A GRE tunnel does not support QoS.
- The firewall does not support a single interface acting as both a GRE tunnel endpoint and a decryption broker.
- GRE tunneling does not support NAT between GRE tunnel endpoints.



If you need to connect to another vendor's network, we recommend you [Set Up an IPSec Tunnel](#), not a GRE tunnel; you should use a GRE tunnel only if that is the only point-to-point tunnel mechanism that the vendor supports. You can also enable GRE over IPSec if the remote endpoint requires that (Add GRE Encapsulation). Add GRE encapsulation in cases where the remote endpoint requires traffic to be encapsulated within a GRE tunnel before IPSec encrypts the traffic. For example, some implementations require multicast traffic to be encapsulated before IPSec encrypts it. If this is a requirement for your environment and the GRE tunnel and IPSec tunnel share the same IP address, Add GRE Encapsulation when you set up the IPSec tunnel.



If you aren't planning to terminate a GRE tunnel on the firewall, but you want the ability to inspect and control traffic passing through the firewall inside a GRE tunnel, don't create a GRE tunnel. Instead, perform [Tunnel Content Inspection](#) of GRE traffic. With tunnel content inspection, you are inspecting and enforcing policy on GRE traffic passing through the firewall, not creating a point-to-point, logical link for the purpose of directing traffic.

Create a GRE Tunnel

Create a [Generic Routing Encapsulation \(GRE\) tunnel](#) to connect two endpoints in a point-to-point, logical link.

STEP 1 | Create a tunnel interface.

1. Select **Network > Interfaces > Tunnel**.
2. **Add** a tunnel and enter the tunnel **Interface Name** followed by a period and a number (range is 1 to 9,999). For example, **tunnel.1**.
3. On the **Config** tab, assign the tunnel interface to a **Virtual Router**.
4. Assign the tunnel interface to a **Virtual System** if the firewall supports multiple virtual systems.
5. Assign the tunnel interface to a **Security Zone**.



6. Assign an IP address to the tunnel interface. (You must assign an IP address if you want to route to this tunnel or monitor the tunnel endpoint.) Select **IPv4** or **IPv6** or configure both.



This address and the corresponding address of the tunnel interface of the peer should be on the same subnet because it is a point-to-point, logical link.

- (**IPv4 only**) On the **IPv4** tab, **Add** an IPv4 address, select an address object, or click **New Address** and specify the **Type** of address and enter it. For example, enter **192.168.2.1**.
- (**IPv6 only**) On the **IPv6** tab, **Enable IPv6 on the interface**.
 1. For **Interface ID**, select **EUI-64 (default 64-bit Extended Unique Identifier)**.
 2. **Add** a new **Address**, select an IPv6 address object, or click **New Address** and specify an address **Name**. **Enable address on interface** and click **OK**.
 3. Select **Type** of address and enter the IPv6 address or FQDN and click **OK** to save the new address.
 4. Select **Enable address on interface** and click **OK**.

7. Click **OK**.

STEP 2 | Create a GRE tunnel to force packets to traverse a specific point-to-point path.

1. Select **Network > GRE Tunnels** and **Add** a tunnel by **Name**.
2. Select the **Interface** to use as the local GRE tunnel endpoint (source interface), which is an Ethernet interface or subinterface, an Aggregate Ethernet (AE) interface, a loopback interface, or a VLAN interface.
3. Select the **Local Address** to be **IP** and select the IP address of the interface you just selected.
4. Enter the **Peer Address**, which is the IP address of the opposite endpoint of the GRE tunnel.
5. Select the **Tunnel Interface** that you created in Step 1. (This identifies the tunnel when it is the egress **Interface** for routing.)
6. Enter the **TTL** for the IP packet encapsulated in the GRE packet (range is 1 to 255; default is 64).
7. Select **Copy ToS Header** to copy the Type of Service (ToS) field from the inner IP header to the outer IP header of the encapsulated packets to preserve the original ToS information. Select this option if your network uses QoS and depends on the ToS bits for enforcing QoS policies.



STEP 3 | (Best Practice) Enable the Keep Alive function for the GRE tunnel.



If Keep Alive is enabled, by default it takes three unreturned keepalive packets (Retries) at 10-second intervals for the GRE tunnel to go down and it takes five Hold Timer intervals at 10-second intervals for the GRE tunnel to come back up.

1. Select **Keep Alive** to enable the keepalive function for the GRE tunnel (default is disabled).
2. (Optional) Set the **Interval (sec)** (in seconds) between keepalive packets that the local end of the GRE tunnel sends to the tunnel peer. This is also the interval that, when multiplied by the **Hold Timer**, is the length of time that the firewall must see successful keepalive packets before the GRE tunnel comes back up (range is 1 to 50; default is 10). Setting an interval too small will cause many keepalive packets that might be unnecessary in your environment and will require extra bandwidth and processing. Setting an interval too large can delay failover because error conditions might not be identified immediately.
3. (Optional) Enter the **Retry** setting, which is the number of intervals that keepalive packets are not returned before the firewall considers the tunnel peer down (range is 1 to 255; default is 3). When the tunnel is down, the firewall removes routes associated with the tunnel from the forwarding table. Configuring a retry setting helps avoid taking measures on a tunnel that is not really down.
4. (Optional) Set the **Hold Timer**, which is the number of **Intervals** that keepalive packets are successful, after which the firewall re-establishes communication with the tunnel peer (range is 1 to 64; default is 5).

STEP 4 | Click **OK**.

STEP 5 | Configure a routing protocol or static route to route traffic to the destination by way of the GRE tunnel. For example, [Configure a Static Route](#) to the network of the destination server and specify the egress **Interface** to be the local tunnel endpoint (tunnel.1). Configure the Next Hop to be the IP address of the tunnel at the opposite end. For example, 192.168.2.3.

STEP 6 | **Commit** your changes.

STEP 7 | Configure the opposite end of the tunnel with its public IP address, its local and peer IP addresses (that correspond to the peer and local IP addresses, respectively, of the GRE tunnel on the firewall), and its routing protocol or static route.

STEP 8 | Verify that the firewall can communicate with the tunnel peer over the GRE tunnel.

1. [Access the CLI](#).
2. `> ping source 192.168.2.1 host 192.168.2.3`

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.

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.

DHCP Message	Description
	- 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.

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 don't 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
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 DHCP 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 DHCP OFFER and DHCP ACK 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.
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.

Capacities are as follows:

- For firewall models other than PA-5200 Series and PA-7000 Series firewalls, see the [Product Selection tool](#).
- On PA-5220 firewalls, you can configure a maximum of 500 DHCP servers and a maximum of 2,048 DHCP relay agents minus the number of DHCP servers configured. For example, if you configure 500 DHCP servers, you can configure 1,548 DHCP relay agents.
- On PA-5250, PA-5260, and PA-7000 Series firewalls, you can configure a maximum of 500 DHCP servers, and a maximum of 4,096 DHCP relay agents minus the number of DHCP servers configured. For example, if you configure 500 DHCP servers, you can configure 3,596 DHCP relay agents.

Perform the following task to configure an interface on the firewall to act as a DHCP server.

STEP 1 | Select an interface to be a DHCP Server.

1. Select **Network > DHCP > DHCP Server** and **Add** an **Interface** name or select one.
2. 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.
3. (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.*

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**.

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 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, **Add** 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**.

STEP 4 | (Optional) Add another vendor-specific or custom DHCP option.

1. Repeat the prior step 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, **Add** 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 this step 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 you just specified.
4. (Optional) Repeat the prior two steps to reserve another address.

STEP 7 | Commit your changes.

Click **OK** and **Commit**.

Configure an Interface as a DHCP Client

Before configuring a firewall interface as a DHCP client, make sure you have configured a Layer 3 interface (Ethernet, Ethernet subinterface, VLAN, VLAN subinterface, aggregate, or aggregate subinterface) and the interface is assigned to a virtual router and a zone. Configure an interface as a DHCP client if you need to use DHCP to request an IPv4 address for the interface.



You can also [Configure the Management 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, **Add** a Layer 3 interface or select a configured Layer 3 interface that you want to be a DHCP client.
3. Select the **IPv4** tab and, for **Type**, select **DHCP Client**.
4. Select **Enable**.
5. (Optional) Enable the option to **Automatically create default route pointing to default gateway provided by server** (enabled by default). Enabling this option causes the firewall to create a static route to the default gateway, which is useful when clients try to access many destinations that do not need to have routes maintained in a route table on the firewall.
6. (Optional) Enable the option to **Send Hostname** to assign a hostname to the DHCP client interface and send that hostname (Option 12) to a DHCP server, which can then register the hostname with the DNS server. The DNS server can then automatically manage hostname-to-dynamic IP address resolutions. External hosts can identify the interface by its hostname. The default value indicates **system-hostname**, which is the firewall hostname that you set in **Device > Setup > Management > General Settings**. Alternatively, enter a hostname for the interface, which can be a maximum of 64 characters, including uppercase and lowercase letters, numbers, period (.), hyphen (-), and underscore (_).



7. (Optional) Enter a **Default Route Metric** (priority level) for the route between the firewall and the DHCP server (range is 1 to 65,535; default is 10). 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.



The Default Route Metric for the route between the firewall and the DHCP server is 10 by default. If the static default route 0.0.0.0/0 uses the DHCP interface as its egress interface, that route's default Metric is also 10. Therefore, there are two routes with a metric of 10 and the firewall can randomly choose one of the routes one time and the other route another time.



Suppose you enable the option to Automatically create default route pointing to default gateway provided by server, select a virtual router, add a static route for a Layer 3 interface, change the Metric (which defaults to 10) to a value greater than 10 (for this example, 100) and Commit your changes. In the route table, the route's metric will not indicate 100. Instead, it will indicate the default value of 10, as expected, because 10 takes precedence over the configured value of 100. However, if you change the static route's Metric to a value less than 10 (such as 6), the route in the route table is updated to indicate the configured metric of 6.

8. (Optional) Enable the option to **Show DHCP Client Runtime Info** to see all of the settings the client inherited from its DHCP server.

STEP 2 | Commit your changes.

Click **OK** and **Commit**.

The Ethernet interface should now indicate **Dynamic-DHCP Client** as its **IP Address** on the **Ethernet** tab.

STEP 3 | (Optional) See which interfaces on the firewall are configured as DHCP clients.

1. Select **Network > Interfaces > Ethernet** and check the **IP Address** to see which interfaces indicate DHCP Client.
2. Select **Network > Interfaces > VLAN** and check the **IP Address** to see which interfaces indicate DHCP Client.

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 models do not support this DHCP functionality.



- For hardware-based firewall models (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.

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**.

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**. Don't 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**. Don't select this option if you want to manually configure a domain.
3. Click **OK**.

STEP 3 | Commit your changes.

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**.

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.

Capacities are as follows:

- You can configure a combined total of 500 DHCP servers (IPv4) and DHCP relay agents (IPv4 and IPv6) on all firewall models except for PA-5200 Series and PA-7000 Series firewalls
- On PA-5220 firewalls, you can configure a maximum of 500 DHCP servers and a maximum of 2,048 DHCP relay agents minus the number of DHCP servers configured. For example, if you configure 500 DHCP servers, you can configure 1,548 DHCP relay agents.
- On PA-5250, PA-5260, and PA-7000 Series firewalls, you can configure a maximum of 500 DHCP servers, and a maximum of 4,096 DHCP relay agents minus the number of DHCP servers configured. For example, if you configure 500 DHCP servers, you can configure 3,596 DHCP relay agents.

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.

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 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, **Add** 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, **Add** 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 the prior three steps to enter a maximum of eight DHCP server addresses per IP address family.

STEP 3 | Commit the configuration.

Click **OK** and **Commit**.

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 DHCP Leases](#)
- [View DHCP Client Information](#)
- [Gather Debug Output about DHCP](#)

View DHCP Server Information

Perform this task 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. 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.

- View DHCP pool statistics, IP address the DHCP server assigned, MAC address, state and duration of lease, and lease start time.

```
admin@PA-220> show dhcp server lease interface 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-220>
```

- View the options that a DHCP server has assigned to clients.

```
admin@PA-220> 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-220>
```

Clear DHCP Leases

You have several options for clearing DHCP leases.

- Release expired [DHCP Leases](#) of an interface (server), such as ethernet1/2, before the hold timer releases them automatically. Those addresses will be available in the IP pool again.

```
admin@PA-220> clear dhcp lease interface ethernet1/2 expired-only
```

- Release the lease of a particular IP address, for example, 192.168.3.1.

```
admin@PA-220> clear dhcp lease interface ethernet1/2 ip 192.168.3.1
```

- Release the lease of a particular MAC address, for example, f0:2c:ae:29:71:34.

```
admin@PA-220> 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 either of these CLI commands.

- **admin@PA-220>show dhcp client state <interface_name>**
- **admin@PA-220> 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-220>				

Gather Debug Output about DHCP

To gather debug output about DHCP, use one of the following commands:

- **admin@PA-220> debug dhcpd**
- **admin@PA-220> debug management-server dhcpd**

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](#)
- [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](#)
- [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.

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.

A DNS record of an FQDN includes a time-to-live (TTL) value, and by default the firewall refreshes each FQDN in its cache based on that individual TTL provided the DNS server, as long as the TTL is greater than or equal to the [Minimum FQDN Refresh Time](#) you configure on the firewall, or the default setting of 30 seconds if you don't configure a minimum. Refreshing an FQDN based on its TTL value is especially helpful for securing access to cloud platform services, which often require frequent FQDN refreshes to ensure highly available services. For example, cloud environments that support autoscaling depend on FQDN resolutions for dynamically scaling services up and down, and fast resolutions of FQDNs are critical in such time-sensitive environments.

By configuring a minimum FQDN refresh time, you limit how small a TTL value the firewall honors. If your IP addresses don't change very often you may want to set a higher Minimum FQDN Refresh Time so that the firewall doesn't refresh entries unnecessarily. The firewall uses the higher of the DNS TTL time and the configured Minimum FQDN Refresh Time.

For example, two FQDNs have the following TTL values. The Minimum FQDN Refresh Time overrides smaller (faster) TTL values.

	TTL	If Minimum FQDN Refresh = 26	Actual Refresh Time
FQDN A	20		26
FQDN B	30		30

The FQDN refresh timer starts when the firewall receives a DNS response from the DNS server or DNS proxy object that is resolving the FQDN.

Additionally, you can set a [stale timeout](#) to configure how long the firewall continues to use stale (expired) FQDN resolutions in the event of an unreachable DNS Server. At the end of the stale timeout period, if the DNS server is still unreachable, the stale FQDN entries become unresolved (the firewall removes stale FQDN entries).

The following firewall tasks are related to DNS:

- 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](#).
- Customize how the firewall handles DNS resolution initiated by Security policy rules, reporting, and management services (such as email, Kerberos, SNMP, syslog, and more) for each virtual system, 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](#).
- Configure the firewall to act as a DNS server for a client, as shown in [Use Case 3: Firewall Acts as DNS Proxy Between Client and Server](#).
- Configure an Anti-Spyware profile to [Use DNS Queries to Identify Infected Hosts on the Network](#).
- Enable [Passive DNS Monitoring](#), which allows the firewall to automatically share domain-to-IP address mappings based on your network traffic with Palo Alto Networks. The Palo Alto Networks threat research team uses this information to gain insight into malware propagation and evasion techniques that abuse the DNS system.
- [Enable Evasion Signatures](#) and then enable evasion signatures for threat prevention.
- [Configure an Interface as a DHCP Server](#). This enables the firewall to act 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

When configured as a DNS proxy, 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 multiple tenants (ISP subscribers) with DNS services, each tenant should have its own DNS proxy defined, which keeps the tenant's DNS service separate from other tenants' services.

In the proxy object, you specify the interfaces for which the firewall is acting as DNS proxy. The DNS proxy for the interface does not use the service route; responses to the DNS requests are always sent to the interface assigned to the virtual router where the DNS request arrived.

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. You must enable **Cache** and **Cache EDNS Responses** (under **Network > DNS Proxy > Advanced**) if this DNS proxy object is assigned to **Device > Setup > Services > DNS** or **Device > Virtual Systems > vsys > General > DNS Proxy**. Furthermore, if this DNS proxy object has **DNS proxy rules** configured, those rules also need to have cache enabled (**Turn on caching of domains resolved by this mapping**).

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. [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.

DNS Server Profile

To simplify configuration for a virtual system, a DNS server profile allows you to specify the virtual system that is being configured, an inheritance source or the primary and secondary IP addresses for DNS servers, and a source interface and source address (service route) that will be used in packets sent to the DNS server. The source interface determines the virtual router, which has a route table. The destination IP address is looked up in the 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 of the destination IP egress interface determined by the route table lookup, but the source IP address would be the address configured. The source address is used as the destination address in the reply from the DNS server.

The virtual system report and virtual system server profile send their queries to the DNS server specified for the virtual system, if there is one. (The DNS server used is defined in **Device > Virtual Systems > General > DNS Proxy**.) If there is no DNS server specified for the virtual system, the DNS server specified for the firewall is queried.

You [Configure a DNS Server Profile](#) for a virtual system only; it is 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 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 this 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](#)
- [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.



When the firewall is enabled to act as a DNS proxy, evasion signatures that detected crafted HTTP or TLS requests can alert to instances where a client connects to a domain other than the domains specified in the original DNS query. As a best practice, [Enable Evasion Signatures](#) after configuring DNS proxy to trigger an alert if crafted requests are detected.

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.

STEP 2 | (Optional) Specify DNS Proxy rules.

1. On the **DNS Proxy Rules** tab, **Add 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**, **Add** 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.

[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, **Add a Name**.
2. Enter the Fully Qualified Domain Name (**FQDN**).
3. For **Address**, **Add** the IP address to which the FQDN should be mapped.

You can provide additional IP addresses for an entry. The firewall will provide all of the IP addresses in its DNS response and the client chooses which address to use.
4. Click **OK**.

STEP 4 | Enable caching and configure other advanced settings for the DNS Proxy.

1. On the **Advanced** tab, 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).
2. For **UDP Queries Retries**, enter:
 - **Interval (sec)**—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**—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.)
3. Select **Cache** to enable the firewall to cache FQDN-to-address mappings that it learns. You must enable **Cache** (enabled by default) if this DNS proxy object is used for queries that the firewall

generates (that is, under **Device > Setup > Services > DNS**, or under **Device > Virtual Systems** and you select a virtual system and **General > DNS Proxy**).

- Select **Enable TTL** to limit the length of time the firewall caches DNS resolution entries for the proxy object. Disabled by default.
 - Enter **Time to Live (sec)**, the number of seconds after which all cached entries for the proxy object are removed. After the entries are removed, new DNS requests must be resolved and cached again. Range is 60-86,400. There is no default TTL; entries remain until the firewall runs out of cache memory.
- **Cache EDNS Responses**—You must enable this setting if this DNS proxy object is used for queries that the firewall generates (that is, under **Device > Setup > Services > DNS**, or under **Device > Virtual Systems** and you select a virtual system and **General > DNS Proxy**).

STEP 5 | Commit your changes.

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.

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 **Add** a **Name** for the DNS server profile.
2. For **Location**, select the virtual system to which the profile applies.
3. For **Inheritance Source**, 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.
4. 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**.*

5. Specify the IP address of the **Secondary DNS** server, or leave as **inherited** if you chose an **Inheritance Source**.

STEP 2 | Configure the service route that the firewall automatically uses, based on whether the target DNS Server has an IP address family type of IPv4 or IPv6.

1. Click **Service Route IPv4** to enable the subsequent interface and IPv4 address to be used as the service route, if the target DNS address is an IPv4 address.
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 | Commit the configuration.

Click **OK** and **Commit**.

Use Case 1: Firewall Requires DNS Resolution

In this use case, the firewall is the client requesting DNS resolutions of FQDNs for Security policy rules, reporting, management services (such as email, Kerberos, SNMP, syslog, and more), and management events such as software update services, dynamic software updates, and WildFire. In dynamic environments, FQDNs change more frequently; accurate DNS resolutions allow the firewall to enforce accurate policing, provide reporting and management services, and handle management events. The shared, global DNS services perform the DNS resolution for the management plane functions.



STEP 1 | Configure the primary and secondary DNS servers you want the firewall to use for DNS resolutions.



You must manually configure at least one DNS server on the firewall or it won't be able to resolve hostnames; the firewall cannot use DNS server settings from another source, such as an ISP.

1. Edit the Services settings (**Device > Setup > Services > Global** for firewalls that support multiple virtual systems; **Device > Setup > Services** for those that don't).
2. On the **Services** tab, for **DNS**, select **Servers** and enter the **Primary DNS Server** address and **Secondary DNS Server** address.
3. Proceed to Step 3.

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. Edit the Services settings (**Device > Setup > Services > Global** for firewalls that support multiple virtual systems; **Device > Setup > Services** for those that don't).
2. On the **Services** tab, for **DNS**, select **DNS Proxy Object**.
3. From the **DNS Proxy** list, select the DNS proxy that you want to use to configure global DNS services, or select **DNS Proxy** to configure a new DNS proxy object as follows:
 1. **Enable** and then enter a **Name** for the DNS proxy object.
 2. On firewalls that support multiple virtual systems, for **Location**, select **Shared** for global, firewall-wide DNS proxy services.



Shared DNS proxy objects don't use DNS server profiles because they don't require a specific service route belonging to a tenant virtual system.

3. Enter the **Primary** DNS server IP address. Optionally enter a **Secondary** DNS server IP address.
4. Select the **Advanced** tab. Ensure that **Cache** is enabled and **Cache EDNS Responses** is enabled (both are enabled by default).

5. Click **OK**.

STEP 3 | (Optional) Set a **Minimum FQDN Refresh Time (sec)** to limit how frequently the firewall refreshes FQDN cache entries.

By default, the firewall refreshes each FQDN in its cache based on the individual TTL for the **FQDN in a DNS record**, as long as the TTL is greater than or equal to this minimum FQDN refresh setting (or as long as the TTL is greater than or equal to the default setting of 30 seconds if you don't configure a minimum FQDN refresh time). To set a minimum FQDN refresh time, enter a value in seconds (range is 0 to 14,400; default is 30). A setting of 0 means the firewall refreshes FQDNs based on the TTL value in the DNS records; the firewall doesn't enforce a minimum FQDN refresh time. The firewall uses the higher of the DNS TTL time and the minimum FQDN refresh time.



If the TTL for the FQDN in DNS is short, but your FQDN resolutions don't change as frequently as the TTL timeframe so don't need a faster refresh, you should set a Minimum FQDN Refresh Time to avoid making FQDN refresh attempts more often than necessary.

STEP 4 | (Optional) Specify an **FQDN Stale Entry Timeout (min)**, which is the number of minutes that the firewall continues to use stale FQDN resolutions in the event of an unreachable DNS server (range is 0 to 10,080; default is 1,440).

A setting of 0 means the firewall does not continue to use a stale FQDN entry.




Make sure the FQDN stale entry timeout is short enough not to allow incorrect traffic forwarding (which can pose a security risk), but long enough to allow traffic continuity without causing an unplanned network outage.

STEP 5 | 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.



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.



STEP 1 | For each virtual system, specify the DNS Proxy to use.

1. Select **Device > Virtual Systems** and **Add** the **ID** of the virtual system (range is 1-255), and an optional **Name**, in this example, Corp1 Corporation.
2. On the **General** tab, choose a **DNS Proxy** or create a new one. In this example, Corp1 DNS Proxy is selected as the proxy for Corp1 Corporation's virtual system.
3. For **Interfaces**, click **Add**. In this example, Ethernet1/20 is dedicated to this tenant.
4. For **Virtual Routers**, click **Add**. A virtual router named Corp1 VR is assigned to the virtual system in order to separate routing functions.
5. Click **OK**.

STEP 2 | Configure a DNS Proxy and a server profile to support DNS resolution for a virtual system.

1. Select **Network > DNS Proxy** and click **Add**.
2. Click **Enable** and enter a **Name** for the DNS Proxy.
3. For **Location**, select the virtual system of the tenant, in this example, Corp1 Corporation (vsys6). (You could choose the **Shared DNS Proxy** resource instead.)
4. For **Server Profile**, choose or create a profile to customize DNS servers to use for DNS resolutions for this tenant's security policy, reporting, and server profile services.

If the profile is not already configured, in the **Server Profile** field, click **DNS Server Profile** to [Configure a DNS Server Profile](#).

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. Select the **Advanced** tab. Ensure that **Cache** is enabled and **Cache EDNS Responses** is enabled (both are enabled by default). This is required if this DNS proxy object is under Device > Virtual Systems > vsys > General > DNS Proxy.
7. Click **OK**.
8. Click **OK** and **Commit**.



Optional advanced features such as split DNS can be configured using DNS Proxy Rules. A separate DNS server profile can be used to redirect DNS resolutions matching the Domain Name in a DNS Proxy Rule to another set of DNS servers, if required. Use Case 3 illustrates split DNS.

If you use two separate DNS server profiles in the same DNS Proxy object, one for the DNS Proxy and one for the DNS proxy rule, the following behaviors occur:


- If a service route is defined in the DNS server profile used by the DNS Proxy, it takes precedence and is used.
- If a service route is defined in the DNS server profile used in the DNS proxy rules, it is not used. If the service route differs from the one defined in the DNS server profile used by the DNS Proxy, the following warning message is displayed during the **Commit** process:

Warning: The DNS service route defined in the DNS proxy object is different from the DNS proxy rule's service route. Using the DNS proxy object's service route.

- If no service route is defined in any DNS server profile, the global service route is used if needed.

Use Case 3: Firewall Acts as DNS Proxy Between Client and Server

In this use case, the firewall is located between a DNS client and a DNS server. A DNS Proxy on the firewall is configured to act as the DNS server for the hosts that reside on the tenant's network connected to the firewall interface. In such a scenario, the firewall performs DNS resolution on its dataplane.



This scenario happens to use *split DNS*, a configuration where DNS Proxy rules are configured to redirect DNS requests to a set of DNS servers based on a domain name match. If there is no match, the server profile determines the DNS servers to which 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 172.16.1.1 to the DNS proxy at 192.168.1.1, then the request to the DNS server (at 10.10.10.10) would use a source of 192.168.1.1 and a destination of 10.10.10.10.

STEP 1 | Select **Network > DNS Proxy** and click **Add**.

-
- STEP 2** | Click **Enable** and enter a **Name** for the DNS Proxy.
- STEP 3** | For **Location**, select the virtual system of the tenant, in this example, Corp1 Corporation (vsys6).
- STEP 4** | For **Interface**, select the interface that will receive the DNS requests from the tenant's hosts, in this example, Ethernet1/20.
- STEP 5** | Choose or create a **Server Profile** to customize DNS servers to resolve DNS requests for this tenant.
- STEP 6** | On the **DNS Proxy Rules** tab, **Add a Name** for the rule.
- STEP 7** | (Optional) Select **Turn on caching of domains resolved by this mapping**.
- STEP 8** | **Add** one or more **Domain Name(s)**, one entry per row. [DNS Proxy Rule and FQDN Matching](#) describes how the firewall matches FQDNs to domain names in a DNS proxy rule.
- STEP 9** | For **DNS Server profile**, select a profile. 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.
- STEP 10** | In this example, if the domain in the request matches myweb.corp1.com, the DNS server defined in the myweb DNS Server Profile is used. If there is no match, the DNS server defined in the **Server Profile** (Corp1 DNS Server Profile) is used.
- STEP 11** | Click **OK** twice.

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 domain 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.	Rule: *.boat.com FQDN: www.boat.com — Match FQDN: www.blue.boat.com — Match FQDN: boat.com — Not a Match

FQDN Comparison to DNS Proxy Rule	For Example
This means a rule consisting of only a wildcard (*) matches any FQDN with one or more tokens.	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)
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 Comparison to DNS Proxy Rule	For Example
	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.)
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 Comparison to DNS Proxy Rule	For Example
	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
Best Practices for Creating DNS Proxy Rules to Avoid Ambiguity and Unexpected Results	
Include a top-level domain in the domain name to avoid invoking an Implicit-tail-match that may match the FQDN to more than one rule.	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.	
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

Dynamic DNS Overview

When you have services hosted behind the firewall and use destination NAT policies on the firewall to access those services or when you need to provide remote access to the firewall, you can register IPv4 address changes (whether the interface is a DHCP client receiving a dynamic address or has a static address) or IPv6 address changes (static address only) for the interface with a dynamic DNS (DDNS) service provider. The DDNS service automatically updates the domain name-to-IP address mappings to provide accurate IP addresses to DNS clients, which, in turn, can access the firewall and services behind the firewall. DDNS is often used in branch deployments that are hosting services. Without DDNS support for firewall interfaces, you would need external components to provide accurate IP addresses to clients.

The firewall supports the following [DDNS service providers](#): DuckDNS, DynDNS, FreeDNS Afraid.org, Dynamic API, FreeDNS Afraid.org, and No-IP. The individual DDNS service provider determines the services it provides, such as how many IP addresses it supports for a hostname and whether it supports IPv6 addresses. Palo Alto Networks uses content updates to add new DDNS service providers and to provide updates to their services.




For high availability (HA) configurations, make sure that content versions on the HA firewall peers (active/passive or active/active) are in sync because the firewall maintains the DDNS configuration based on the current Palo Alto Networks content release version. Palo Alto Networks can change or deprecate existing DDNS services through a content release. Additionally, a DDNS service provider can change the services it provides. A mismatch in content release versions between the HA peers can cause issues with their ability to use the DDNS service.



The firewall does not support DDNS over an interface that is a Point-to-Point Protocol over Ethernet (PPPoE) termination point.

In the following example, the firewall is a DDNS client of a DDNS service provider. Initially, the DHCP server assigns IP address 10.1.1.1 to the Ethernet 1/2 interface. A destination NAT policy translates the public-facing 10.1.1.1 to the real address of Server A (192.168.10.1) behind the firewall.




1. When a user attempts to contact `www.serverA.companyx.com`, the user queries its local DNS server for the IP address. The `www.serverA.companyx.com` (set, for example, as a CNAME to your `duckdns.org`

record: serverA.companyx.duckdns.org) is a domain belonging to the DDNS provider (DuckDNS in this example). The DNS server checks for the record with the DDNS provider to resolve the query.

2. The DNS server responds to the user with 10.1.1.1, which is the IP address for www.serverA.companyx.com.
3. The user packet with destination 10.1.1.1 goes to firewall interface Ethernet 1/2.
4. In this example, the firewall performs destination NAT and translates 10.1.1.1 to 192.168.10.1 before sending the packet to the destination.

After some time passes, DHCP assigns a new IP address to the firewall interface, which triggers a DDNS update, as follows:



1. The DHCP Server assigns a new IP address (10.1.2.2) to Ethernet 1/2.
2. When the firewall receives the new address, it sends an update to the DDNS service with the new address for www.serverA.companyx.com, which the DDNS service registers. (The firewall also sends regular updates based on the update interval you configure. The firewall sends DDNS updates over HTTPS port 443.)

Consequently, the next time the client queries the DNS server for the IP address of www.serverA.companyx.com and the DNS server checks the DDNS service, the DDNS service sends the updated address (10.1.2.2). Thus, the user successfully accesses a service or application through the firewall interface using the updated interface address.



If your firewall is configured for HA active/passive mode, be aware that the firewall sends DDNS updates to the DDNS service while the two HA firewall states are converging. After the HA states converge, DDNS is disabled on the passive firewall. For example, when two HA firewalls first boot up, they both send DDNS updates until they establish whether they are in HA active or passive mode. During this interval, you still see DDNS updates in system logs. After the HA states converge and each firewall notifies its clients that it is active or passive, the passive firewall no longer sends DDNS updates. (In HA active/active mode, each firewall has an independent DDNS configuration and doesn't synchronize the DDNS configuration.)

Configure Dynamic DNS for Firewall Interfaces

Before you configure **DDNS** for a firewall interface:

- Determine the hostname that you registered with your DDNS provider.
- Obtain the public SSL certificate from the DDNS service and import it in to the firewall.
- (If you use [FreeDNS Afraid.org v1](#) or [FreeDNS Afraid.org Dynamic API v1](#)) On the DDNS server, the Dynamic DNS service tab includes the following option: **Link updates of the same IP together?** When this option is enabled, the DDNS service updates all hostnames in DNS records that contain the old IP address that is changing, not just the DNS record for a single hostname and IP address. To avoid updating DNS records of hosts you didn't intend to update, you should disable the **Link updates of the same IP together?** option so that the DDNS server updates only the DNS record that contains the specific hostname with the new IP address that is in the DDNS update.


STEP 1 | Configure DDNS.

1. Select **Network > Interfaces > Ethernet** and select a Layer 3 interface, subinterface, or Aggregate Ethernet (AE) interface; or select **Network > Interfaces > VLAN** and select an interface or subinterface.
2. Select **Advanced > DDNS** and select **Settings**.
3. **Enable** DDNS. You must initially enable DDNS to configure it. (If your DDNS configuration is unfinished, you can save it without enabling it so that you do not lose your partial configuration.)
4. Enter the **Update Interval (days)**, which is the number of days between updates that the firewall sends to the DDNS service to update IP addresses mapped to FQDNs (default is 1; range is 1 to 30). Choose an interval based on how frequently your IP addresses change. (The updates that the firewall sends at regular intervals are in addition to the updates the firewall sends upon receiving an address change. The updates sent at regular intervals are to ensure that updates sent per address change are not lost, for example.)
5. Enter the **Hostname** for the interface, which is already registered with the DDNS service (for example, `www.serverA.companyx.com` or `serverA`).



Make sure this hostname matches the hostname you registered with your DDNS service. You should enter an FQDN for the hostname; the firewall doesn't validate the hostname except to confirm that the syntax uses only valid characters allowed by DNS for a domain name.

6. Select **IPv4** and select one or more IPv4 addresses assigned to the interface or **Add** an IPv4 address to associate with the hostname (for example, `10.1.1.1`). You can select only as many IPv4 addresses as the DDNS service allows. All selected IPv4 addresses are registered with the DDNS service. Select at least one IPv4 or one IPv6 address.
7. Select **IPv6** and select one or more IPv6 addresses assigned to the interface or **Add** an IPv6 address to associate with the hostname. You can select only as many IPv6 addresses as the DDNS service allows. All selected IPv6 addresses are registered with the DDNS service. Select at least one IPv4 or one IPv6 address.
8. Select or [create a new certificate profile \(Certificate Profile\)](#) using the imported SSL certificate from the DDNS service to verify the SSL certificate of the DDNS service when the firewall first connects to a DDNS service to register an IP address and at every update. When the firewall connects to the DDNS service to send updates, the DDNS service presents the firewall with an SSL certificate signed by the certificate authority (CA) so that the firewall can authenticate the DDNS service.
9. Select the **Vendor** (and version number) you are using for DDNS service.



Palo Alto Networks may change the supported DDNS service providers via a content update.

10. The vendor choice determines the vendor-specific **Name** and **Value** fields below the Vendor field. Some Value fields are read-only to notify you of the parameters the firewall uses to connect to the DDNS service. Configure the remaining Value fields, such as a password that the DDNS service provides to you and a timeout that the firewall uses if it doesn't receive an update from the DDNS service.
11. Click **OK**.

STEP 2 | (Optional) If you want the firewall to communicate with the DDNS service using an interface other than the management interface, configure a service route for DDNS ([Set Up Network Access for External Services](#)).

STEP 3 | Commit your changes.

STEP 4 | View DDNS information for the interface.

1. Select **Network > Interfaces > Ethernet** or **Network > Interfaces > VLAN** and select the interface you configured. (Interfaces with DDNS configured display the DDNS icon——in the Features field.)
2. Select **Advanced > DDNS** and **Settings**.
3. **Show Runtime Info** to see the DDNS information for the interface, including the Last return code (result of the last FQDN update) and Last time (date and time) the DDNS service received an FQDN update.

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](#)
- [Destination NAT with DNS Rewrite Use Cases](#)
- [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.

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 selecting **Device > Troubleshooting** and testing the traffic matches for the NAT rule. For example:

The screenshot displays the 'Test Configuration' and 'Test Result' sections of the Palo Alto Networks firewall configuration interface. The 'Test Configuration' section on the left includes fields for 'Select Test' (NAT Policy Match), 'From' (G-vlan-trust), 'To' (G-vlan-trust), 'Source' (10.54.21.28), 'Destination' (9.9.9.9), 'Source Port' (11-65535), 'Destination Port' (443), 'Protocol' (8), 'To Interface' (None), and 'Use Device ID' (0 + 12). Below these fields are 'Execute' and 'Reset' buttons. The 'Test Result' section on the right shows the 'NAT Policy Match Result' with a table titled 'Result Detail' containing columns for 'Name' and 'Value'. The table has one row with 'Result' and 'success-002'.

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 10.1.1.1 to the address in the NAT pool, 192.168.2.2. The translated packet is sent on to a router.

For the return traffic, the router does not know how to reach 192.168.2.2 (because that IP address is just an address in the NAT address pool), so it sends an ARP request packet to the firewall.

- If the address pool (192.168.2.2) is in the same subnet as the egress/ingress interface IP address (192.168.2.3/24), 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 (192.168.2.2) 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 192.168.2.2, in order to ensure the return traffic is routed back to the firewall, as shown in the figure below.



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](#).



(Affects only PA-7000 Series firewalls that do not use second-generation PA-7050-SMC-B or PA-7080-SMC-B Switch Management Cards) When you use Point-to-Point Tunnel Protocol (PPTP) with DIPP NAT, the firewall is limited to using a translated IP address-and-port pair for only one connection; the firewall does not support DIPP NAT. The workaround is to upgrade the PA-7000 Series firewall to a second-generation SMC-B card.

- **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 destination address to a different destination address; for example, it translates a public destination address to a private destination address. Destination NAT also offers the option to perform port forwarding or port translation.

Destination NAT allows static and dynamic translation:

- **Static IP**—You can configure a one-to-one, [static translation](#) in several formats. You can specify that the original packet have a single destination IP address, a range of IP addresses, or an IP netmask, as long as the translated packet is in the same format and specifies the same number of IP addresses. The firewall statically translates an original destination address to the same translated destination address each time. That is, if there is more than one destination address, the firewall translates the first destination address configured for the original packet to the first destination address configured for the translated packet, and translates the second original destination address configured to the second translated destination address configured, and so on, always using the same translation.

If you use destination NAT to translate a static IPv4 address, you might also use DNS services on one side of the firewall to resolve FQDNs for a client on the other side. When the DNS response containing the IPv4 address traverses the firewall, the DNS server provides an internal IP address to an external device, or vice versa. Beginning with PAN-OS 9.0.2 and in later 9.0 releases, you can configure the firewall to rewrite the IP address in the DNS response (that matches the rule) so that the client receives the appropriate address to reach the destination service. The applicable [DNS rewrite use case](#) determines how you configure such a rewrite.

- **Dynamic IP (with session distribution)**—Destination NAT allows you to translate the original destination address to a destination host or server that has a [dynamic IP address](#), such as an address group or address object that uses an IP netmask, IP range, or FQDN, any of which can return multiple addresses from DNS. Dynamic IP (with session distribution) supports IPv4 addresses only. Destination NAT using a dynamic IP address is especially helpful in cloud deployments that use dynamic IP addressing.

If the translated destination address resolves to more than one address, the firewall distributes incoming NAT sessions among the multiple addresses to provide improved session distribution. Distribution is based on one of several methods: round-robin (the default method), source IP hash, IP modulo, IP hash, or least sessions. If a DNS server returns more than 32 IPv4 addresses for an FQDN, the firewall uses the first 32 addresses in the packet.



If the translated address is an address object of type FQDN that resolves to only IPv6 addresses, the destination NAT policy rule considers the FQDN as unresolved.

Using **Dynamic IP (with session distribution)** allows you to translate multiple pre-NAT destination IP addresses M to multiple post-NAT destination IP addresses N . A many-to-many translation means there can be $M \times N$ destination NAT translations using a single NAT rule.



For destination NAT, the best practice is to:

- Use Static IP address translation for static IP addresses, which allows the firewall to check and ensure that the number of original destination IP addresses equals the number of translated destination IP addresses.
- Use Dynamic IP (with session distribution) address translation only for FQDN-based dynamic addresses (the firewall does not perform an IP address number check).

The following are common examples of destination NAT translations that the firewall allows:

Translation Type	Original Packet's Destination Address	Maps to Translated Packet's Destination Address	Notes
Static IP	192.168.1.1	2.2.2.2	Original packet and translated packet each have one possible destination address.
	192.168.1.1-192.168.1.4	2.2.2.1-2.2.2.4	Original packet and translated packet each have four possible destination addresses: 192.168.1.1 always maps to 2.2.2.1 192.168.1.2 always maps to 2.2.2.2 192.168.1.3 always maps to 2.2.2.3 192.168.1.4 always maps to 2.2.2.4
	192.168.1.1/30	2.2.2.1/30	Original packet and translated packet each have four possible destination addresses: 192.168.1.1 always maps to 2.2.2.1 192.168.1.2 always maps to 2.2.2.2 192.168.1.3 always maps to 2.2.2.3 192.168.1.4 always maps to 2.2.2.4
Dynamic IP (with session distribution)	192.168.1.1/30	domainname.com	Original packet has four destination addresses and if, for example, the FQDN in the translated destination address resolves to five IP addresses, then there are 20 possible destination NAT translations in a single NAT rule.

One common use for 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 actual port number private. The port translation 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](#).

Destination NAT with DNS Rewrite Use Cases

When you use destination NAT to perform a static translation from one IPv4 address to a different IPv4 address, you may also be using DNS services on one side of the firewall to resolve FQDNs for a client. When the DNS response containing the IP address traverses the firewall to go to the client, the firewall doesn't perform NAT on that IP address, so the DNS server provides an internal IP address to an external device, or vice versa, resulting in the DNS client being unable to connect to the destination service.

Beginning with PAN-OS 9.0.2 and in later 9.0 releases, you can [configure the firewall to rewrite the IP address in the DNS response](#) (from the A Record) based on the translated IP address configured for the NAT policy rule. The firewall performs NAT on the IPv4 address (the FQDN resolution) in the DNS response before forwarding the response to the client; thus, the client receives the appropriate address to reach the destination service. A single NAT policy rule causes the firewall to perform NAT on packets that match the rule, and also causes the firewall to perform NAT on IP addresses in DNS responses that match the original destination address or translated destination address in the rule. You must specify how the firewall performs NAT on the IP address in the DNS response relative to the NAT rule: **reverse** or **forward**. For example:

- If you enable DNS rewrite with the **reverse** setting in a NAT rule that performs static translation of IP address **1.1.1.10 to 192.168.1.10**, the firewall rewrites a DNS response in the reverse way, translating **192.168.1.10 to 1.1.1.10**.
- If you enable DNS rewrite with the **forward** setting in a NAT rule that performs static translation of IP address **1.1.1.10 to 192.168.1.10**, the firewall rewrites a DNS response in the same way, translating **1.1.1.10 to 192.168.1.10**.



If you have an overlapping NAT rule with DNS Rewrite disabled, and a NAT rule below it that has DNS Rewrite enabled and is included in the overlap, the firewall rewrites the DNS response according to the overlapped NAT rule (in either reverse or forward setting). The rewrite takes precedence and the order of the NAT rules is ignored.

Consider the use cases for configuring DNS rewrite:

- [Destination NAT with DNS Rewrite Reverse Use Cases](#)
- [Destination NAT with DNS Rewrite Forward Use Cases](#)


Destination NAT with DNS Rewrite Reverse Use Cases

The following use cases illustrate [destination NAT with DNS rewrite](#) enabled in the **reverse** direction. The difference between these two use cases is simply whether the DNS client, DNS server, and destination server are on the public or internal side of the firewall. In either case, the DNS client is on the opposite side of the firewall from its ultimate destination server. (If your DNS client and its ultimate destination server are on the same side of the firewall, consider [Destination NAT with DNS Rewrite Forward Use Cases](#) 3 and 4.)

Use case 1 illustrates the DNS client on the public side of the firewall, while the DNS server and the ultimate destination server are both on the internal side. This case requires DNS rewrite in the reverse direction. The DNS client queries for the IP address of red.com. Based on the NAT rule, the firewall translates the query (originally going to public address 1.1.2.1) to internal address 192.168.2.1. The DNS


server responds that red.com has IP address 192.168.2.10. The rule includes **Enable DNS Rewrite - reverse** and the DNS response of 192.168.2.10 matches the destination Translated Address of 192.168.2.0/24 in the rule, so the firewall translates the DNS response using the **reverse** translation that the rule uses. The rule says translate 1.1.2.0/24 to 192.168.2.0/24, so the firewall rewrites the DNS response of 192.168.2.10 to 1.1.2.10. The DNS client receives the response and sends to 1.1.2.10, which the rule translates to 192.168.2.10 to reach server red.com.

Use case 1 summary: DNS client and destination server are on opposite sides of the firewall. The DNS server provides an address that matches the translated destination address in the NAT rule, so translate the DNS response using the **reverse** translation of the NAT rule.



Use case 2 illustrates the DNS client on the internal side of the firewall, while the DNS server and the ultimate destination server are both on the public side. This case requires DNS rewrite in the reverse direction. The DNS client queries for the IP address of red.com. Based on the NAT rule, the firewall translates the query (originally going to internal address 192.168.2.1) to the public address 1.1.2.1. The DNS server responds that red.com has IP address 1.1.2.10. The rule includes **Enable DNS Rewrite - reverse** and the DNS response of 1.1.2.10 matches the destination Translated Address of 1.1.2.0/24 in the rule, so the firewall translates the DNS response using the **reverse** translation that the rule uses. The rule says translate 192.168.2.0/24 to 1.1.2.0/24, so the firewall rewrites the DNS response 1.1.2.10 to 192.168.2.10. The DNS client receives the response and sends to 192.168.2.10, which the rule translates to 1.1.2.10 to reach server red.com.

Use case 2 summary is the same as Use case 1 summary: DNS client and destination server are on opposite sides of the firewall. The DNS server provides an address that matches the translated destination address in the NAT rule, so translate the DNS response using the **reverse** translation of the NAT rule.




To implement DNS rewrite, [Configure Destination NAT with DNS Rewrite](#).

Destination NAT with DNS Rewrite Forward Use Cases

The following use cases illustrate [destination NAT with DNS rewrite](#) enabled in the **forward** direction. The difference between these two use cases is simply whether the DNS client, DNS server, and destination server are on the public or internal side of the firewall. In either case, the DNS client is on the same side of the firewall as its ultimate destination server. (If your DNS client and its ultimate destination server are on opposite sides of the firewall, consider [Destination NAT with DNS Rewrite Reverse Use Cases 1 and 2](#).)


Use case 3 illustrates the DNS client and the ultimate destination server both on the internal side of the firewall, while the DNS server is on the public side. This case requires DNS rewrite in the forward direction. The DNS client queries for the IP address of red.com. Based on Rule 1, the firewall translates the query (originally going to internal address 192.168.1.1) to 1.1.1.1. The DNS server responds that red.com has IP address 1.1.2.10. Rule 2 includes **Enable DNS Rewrite - forward** and the DNS response of 1.1.2.10 matches the original destination address of 1.1.2.0/24 in Rule 2, so the firewall translates the DNS response using the **same** translation the rule uses. Rule 2 says translate 1.1.2.0/24 to 192.168.2.0/24, so the firewall rewrites DNS response 1.1.2.10 to 192.168.2.10. The DNS client receives the response and sends to 192.168.2.10 to reach server red.com.

Use case 3 summary: DNS client and destination server are on the same side of the firewall. The DNS server provides an address that matches the original destination address in the NAT rule, so translate the DNS response using the same (**forward**) translation as the NAT rule.



Use case 4 illustrates the DNS client and the ultimate destination server both on the public side of the firewall, while the DNS server is on the internal side. This case requires DNS Rewrite in the forward direction. The DNS client queries for the IP address of red.com. Based on Rule 2, the firewall translates the query (originally going to public destination 1.1.2.1) to 192.168.2.1. The DNS server responds that red.com has IP address 192.168.2.10. Rule 1 includes **Enable DNS Rewrite - forward** and the DNS response of 192.168.2.10 matches the original destination address of 192.168.2.0/24 in Rule 1, so the firewall translates the DNS response using the **same** translation the rule uses. Rule 1 says translate 192.168.2.0/24 to 1.1.2.0/24, so the firewall rewrites DNS response 192.168.2.10 to 1.1.2.10. The DNS client receives the response and sends to 1.1.2.10 to reach server red.com.

Use case 4 summary is the same as Use case 3 summary: DNS client and destination server are on the same side of the firewall. The DNS server provides an address that matches the original destination address in the NAT rule, so translate the DNS response using the same (**forward**) translation as the NAT rule.



To implement DNS rewrite, [Configure Destination NAT with DNS Rewrite](#).

NAT Rule Capacities

The number of NAT rules allowed is based on the firewall model. 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 model-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 model'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 model. 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 model 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 model.

Model	Default Oversubscription Rate
PA-220	2
PA-820	2
PA-850	2
PA-3220	4
PA-3250	4
PA-3260	4
PA-5220	4

Model	Default Oversubscription Rate
PA-5250	8
PA-5260	8
PA-5280	8
PA-7050	8
PA-7080	8
VM-50	2
VM-100	1
VM-200	1
VM-300	2
VM-500	8
VM-700	8
VM-1000-HV	2

The firewall supports a maximum of 256 translated IP addresses per NAT rule, and each model 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 model, 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:

```
admin@PA-7050-HA-0 (active-primary) > show running global-ippool
```

Idx	Type	From	To	Num	Ref.Cnt	Size	Ratio
1	Dynamic IP	201.0.0.0-201.0.255.255	210.0.0.0	4096	2	657072	N/A
2	Dynamic IP	202.0.0.0-202.0.0.255	220.0.0.0	256	1	41232	N/A
3	Dynamic IP/Port	200.0.2.100-200.0.2.100	200.0.1.11	1	1	68720	8

Usable NAT DIP/DIPP shared memory size: 56490064	← Total physical NAT memory (bytes)
Used NAT DIP/DIPP shared memory size: 767024 (1.3%)	← Bytes and % of usable NAT memory
Dynamic IP NAT Pool: 2 (1.19%)	← Number of DIP pools in use and % of total usable memory that all DIP pools use
Dynamic IP/Port NAT Pool: 1 (0.12%)	← Number of DIPP pools in use and % of total usable memory that all DIPP pools use

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.


```
admin@PA-7050-1A-0(vsys1|active-primary)> show running ippool
```

VSYS 1 has 4 NAT rules, DiP and DiPP rules:

Rule	Type	Used	Available	Mem Size	Ratio
nat1	Dynamic IP	0	4096	788144	0
nat2	Dynamic IP	0	256	49424	0
nat3	Dynamic SP/Port	0	638976	100976	4
nat11	Dynamic IP	0	4096	788144	0

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.

```
admin@PA-7050-1A-0(vsys1|active-primary)> show running nat-rule-ippool rule nat1
```

VSYS 1 Rule nat1:

Rule: nat1, Pool index: 1, memory usage: 788144

Reserve IP: no

201.0.0.0-201.0.255.255 <=>

255.0.0.0-255.0.15.255

Source	Xlat-Source	Ref.Cnt(F)	TTL(s)

Total IPs in use: 0

Total entries in time-reserve cache: 0

Total free list left: 4096

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 \(Static Source NAT\)](#)
- [Configure Destination NAT with DNS Rewrite](#)
- [Configure Destination NAT Using Dynamic IP Addresses](#)
- [Modify the Oversubscription Rate for DIPP NAT](#)
- [Reserve Dynamic IP NAT Addresses](#)
- [Disable NAT for a Specific Host or Interface](#)

The first three NAT examples in this section are 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 \(Static Source NAT\)](#).

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.

STEP 1 | Create an address object for the external IP address you plan to use.

1. Select **Objects > Addresses** and **Add** a **Name** and optional **Description** for the object.
2. Select **IP Netmask** from the **Type** and then enter the IP address of the external interface on the firewall, 203.0.113.100 in this example.
3. 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** list.
6. On the **Translated Packet** tab, select **Dynamic IP And Port** from the **Translation Type** list 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**.

STEP 3 | Commit your changes.

Click **Commit**.

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.

STEP 1 | Create an address object for the web server.

1. Select **Objects > Addresses** and **Add** a **Name** and optional **Description** for the address object.

2. For **Type**, select **IP Netmask** and enter the public IP address of the web server, 203.0.113.11 in this example.

You can switch the address object type from **IP Netmask** to **FQDN** by clicking **Resolve**, and when the FQDN appears, click **Use this FQDN**. Alternatively, for **Type**, select **FQDN** and enter the FQDN to use for the address object. If you enter an FQDN and click **Resolve**, the IP address to which the FQDN resolves appears in the field. To switch the address object **Type** from an FQDN to an IP Netmask using this IP address, click **Use this address** and the **Type** will switch to **IP Netmask** with the IP address appearing in the field.

3. 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** list.
4. In the **Destination Address** section, **Add** the address object you created for your public web server.
5. On the **Translated Packet** tab, for Destination Address Translation, for **Translation Type**, select **Static IP** 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. Alternatively, you can select **Translation Type** to be **Dynamic IP (with session distribution)** and enter the **Translated Address** to be an address object or address group that uses an IP netmask, IP range, or FQDN. Any of these can return multiple addresses from DNS. If the translated destination address resolves to more than one address, the firewall distributes incoming NAT sessions among the multiple addresses based on one of several methods you can select: **Round Robin** (the default method), **Source IP Hash**, **IP Modulo**, **IP Hash**, or **Least Sessions**.
6. Click **OK**.

STEP 3 | Click **Commit**.

Enable Bi-Directional Address Translation for Your Public-Facing Servers (Static Source NAT)

When your public-facing servers have private IP addresses assigned on the network segment where they are physically located, you need a source NAT rule to translate the source address of the server to the external address upon egress. You create a static NAT rule to translate the internal source address, 10.1.1.11, to the external web server address, 203.0.113.11 in our example.

However, a public-facing server must be able to both send and receive packets. You need a reciprocal policy that translates the public address (the destination IP address in incoming packets from Internet users) into the private address so that the firewall can route the packet to your DMZ network. You create a bi-directional static NAT rule, as described in the following procedure. Bi-directional translation is an option for static NAT only.

STEP 1 | Create an address object for the web server's internal IP address.

1. Select **Objects > Addresses** and **Add** a **Name** and optional **Description** for the object.
2. Select **IP Netmask** from the **Type** list and enter the IP address of the web server on the DMZ network, 10.1.1.11 in this example.
3. 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** list.
4. In the **Source Address** section, **Add** the address object you created for your internal web server address.
5. On the **Translated Packet** tab, select **Static IP** from the **Translation Type** list in the **Source Address Translation** section and then select the address object you created for your external web server address from the **Translated Address** list.
6. In the **Bi-directional** field, select **Yes**.
7. Click **OK**.

STEP 3 | Commit.

Click **Commit**.

Configure Destination NAT with DNS Rewrite

Beginning with PAN-OS 9.0.2 and in later 9.0 releases, when you configure a destination NAT policy rule that performs static translation of IPv4 addresses, you can also configure the rule so that the firewall rewrites the IPv4 address in a DNS response based on the original or translated IP address configured for the rule. The firewall performs NAT on the IPv4 address (the FQDN resolution) in a DNS response (that matches the rule) before forwarding the response to the client; thus, the client receives the appropriate address to reach the destination service.

View the [DNS rewrite use cases](#) to help you determine whether to specify that the rewrite occur in the **reverse** or **forward** direction.



You cannot enable Bi-directional source address translation in the same NAT rule where you enable DNS rewrite.

STEP 1 | Create a destination NAT policy rule that specifies the firewall perform static translation of IPv4 addresses that match the rule, and also specifies the firewall rewrite IP addresses in DNS responses when that IPv4 address (from the A Record) matches the original or translated destination address in the NAT rule.

1. Select **Policies > NAT** and **Add** a NAT policy rule.
2. (Optional) On the **General** tab, enter a descriptive **Name** for the rule.
3. For **NAT Type**, select **ipv4**.
4. On the **Original Packet** tab, for **Source Zone**, **Add** the appropriate zone
5. For **Destination Zone**, select the appropriate zone.
6. (Optional) **Add Source Address**, **Destination Address**, **Destination Interface**, and/or **Service** to further define the rule.
7. On the **Translated Packet** tab, for Destination Address Translation, select **Translation Type** to be **Static IP**.
8. On the **Translated Packet** tab, for Destination Address Translation, select **Translation Type** to be **Static IP**.
9. **Enable DNS Rewrite** and select a **Direction**:
 - Select **reverse** (default) when the IP address in the DNS response requires the opposite translation that the NAT rule specifies.


- Select **forward** when the IP address in the DNS response requires the same translation that the NAT rule specifies.
10. Click **OK**.

STEP 2 | Commit your changes.

Configure Destination NAT Using Dynamic IP Addresses

Use **Destination NAT** to translate the original destination address to a destination host or server that has a dynamic IP address and uses an FQDN. Destination NAT using a dynamic IP address is especially helpful in cloud deployments, which typically use dynamic IP addressing. When the host or server in the cloud has new (dynamic) IP addresses, you don't need 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.

In the following example topology, clients want to reach servers that are hosting web applications in the cloud. An external Elastic Load Balancer (ELB) connects to firewalls, which connect to internal ELBs that connect to the servers. Over time, Amazon Web Services (AWS), for example, adds (and removes) IP addresses for the FQDN assigned to the internal ELBs based on the demand for services. The flexibility of using an FQDN for NAT to the internal ELB helps the policy to resolve to different IP addresses at different times, making destination NAT easier to use because the updates are dynamic.



STEP 1 | Create an address object using the FQDN of the server to which you want to translate the address. (The address object could also be an IP netmask or IP range.)

1. Select **Objects > Addresses** and **Add** an address object by **Name**, such as **post-NAT-Internal-ELB**.
2. Select **FQDN** as the **Type** and enter the FQDN. In this example, the FQDN is **ielb.appweb.com**.
3. Click **OK**.

STEP 2 | Create the destination NAT policy.

1. Select **Policies > NAT** and **Add** a NAT policy rule by **Name** on the **General** tab.
2. Select **ipv4** as the **NAT Type**.
3. On the **Original Packet** tab, **Add** the **Source Zone** and **Destination Zone**.
4. On the **Translated Packet** tab, in the Destination Address Translation section, select **Dynamic IP (with session distribution)** as the **Translation Type**.
5. For **Translated Address**, select the address object you created for the FQDN, IP netmask, or IP range. In this example, the FQDN is **post-NAT-Internal-ELB**.
6. For **Session Distribution Method**, select one of the following:

- **Round Robin** (default)—Assigns new sessions to IP addresses in rotating order. Unless you have a reason to change the distribution method, round robin distribution is likely suitable.
- **Source IP Hash**—Assigns new sessions based on hash of source IP address. If you have traffic coming from a single source IP address, don't select Source IP Hash; select a different method.
- **IP Modulo**—The firewall takes into consideration the source and destination IP address from the incoming packet; the firewall performs an XOR operation and a modulo operation; the result determines to which IP address the firewall assigns new sessions.
- **IP Hash**—Assigns new sessions based on hash of source and destination IP addresses.
- **Least Sessions**—Assigns new sessions to the IP address with the fewest concurrent sessions. If you have many short-lived sessions, **Least Sessions** provides you with a more balanced distribution of sessions.



The firewall does not remove duplicate IP addresses from the list of destination IP addresses before it distributes sessions among the multiple IP addresses. The firewall distributes sessions to the duplicate addresses in the same way it distributes sessions to non-duplicate addresses. (Duplicate addresses in the translation pool can occur, for example, if the translated address is an address group of address objects, and one address object is an FQDN that resolves to an IP address, while another address object is a range that includes the same IP address.)

7. Click **OK**.

STEP 3 | Commit your changes.

STEP 4 | (Optional) You can configure the frequency at which the firewall refreshes an FQDN ([Use Case 1: Firewall Requires DNS Resolution](#)).

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.

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** list, select **1x**, **2x**, **4x**, or **8x**, depending on which ratio you want.



The Platform Default setting applies the default oversubscription setting for the model. If you want no oversubscription, select 1x.

3. Click **OK** and **Commit** the change.

Reserve Dynamic IP NAT Addresses

You can reserve Dynamic IP NAT addresses (for a configurable period of time) to prevent them from being allocated as translated addresses to a different source IP address that needs translation. When configured, the reservation applies to all of the translated Dynamic IP addresses in progress and any new translations.

For both translations in progress and new translations, when a source IP address is translated to an available translated IP address, that pairing is retained even after all sessions related to that specific source IP are expired. The reservation timer for each source IP address begins after all sessions that use that source IP address translation expire. Dynamic IP NAT is a one-to-one translation; one source IP address translates to one translated IP address that is chosen dynamically from those addresses available in the configured pool.

Therefore, a translated IP address that is reserved is not available for any other source IP address until the reservation expires because a new session has not started. The timer is reset each time a new session for a source IP/translated IP mapping begins, after a period when no sessions were active.

By default, no addresses are reserved. You can reserve Dynamic IP NAT addresses for the firewall or for a virtual system.

- Reserve dynamic IP NAT addresses for a firewall.

Enter the following commands:

```
admin@PA-3250# set setting nat reserve-ip yes
```

```
admin@PA-3250# set setting nat reserve-time <1-604800 secs>
```

- Reserve dynamic IP NAT addresses for a virtual system.

Enter the following commands:

```
admin@PA-3250# set vsys <vsysid> setting nat reserve-ip yes
```

```
admin@PA-3250# 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.

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.

STEP 1 | Create the NAT policy.

1. Select **Policies > NAT** and click **Add** a descriptive **Name** for the policy.

2. 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** list.
3. For **Source Address**, click **Add** and enter the host address. Click **OK**.
4. On the **Translated Packet** tab, select **None** from the **Translation Type** list in the Source Address Translation section of the screen.
5. Click **OK**.

STEP 2 | Commit your changes.

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.

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 192.0.2.100.



Before configuring the NAT rules, consider the sequence of events for this scenario.

- ❑ Host 192.0.2.250 sends an ARP request for the address 192.0.2.100 (the public address of the destination server).

- ❑ The firewall receives the ARP request packet for destination 192.0.2.100 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 192.0.2.100 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 192.0.2.100.

- ❑ 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 (192.0.2.100). The configured NAT rule would look like this:

Name	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Port	Destination Port
192.0.2.100 to 10.1.1.100	Untrust-L3	Untrust-L3	Ethernet1/1	192.0.2.100	10.1.1.100	any	any	any

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:

Name	Zone	Address	Zone	Address	Application	Service	Action	Profile	Options
Webserver access	Untrust-L3	any	DMZ	10.1.1.100	Webserver-public	any	allow	none	

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 192.0.2.100 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 (192.0.2.100).



The following NAT and security rules must be configured on the firewall:

Name	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Port	Destination Port
192.0.2.100 to 10.1.1.100	Untrust-L3	Untrust-L3	Ethernet1/1	192.0.2.100	10.1.1.100	any	any	any

Name	Zone	Address	Zone	Address	Application	Service
Webserver access	Untrust-L3	any	DMZ	10.1.1.100	Webserver-public	any

Use the **show session all** CLI command to verify the translation.

Destination NAT Example—One-to-Many Mapping

In this example, one IP address maps to two different internal hosts. The firewall uses the application to identify the internal host to which the firewall forwards the traffic.



All HTTP traffic is sent to host 10.1.1.100 and SSH traffic is sent to server 10.1.1.101. The following address objects are required:

- Address object for the one pre-translated IP address of the server
- Address object for the real IP address of the SSH server
- Address object for the real IP address of the web server

The corresponding address objects are created:

- Servers-public: 192.0.2.100
- SSH-server: 10.1.1.101
- webserver-private: 10.1.1.100

The NAT rules would look like this:

Name	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Port	Destination Port	Source Port	Destination Port
1. NAT rule webserver	Untrust L3	Trust L3	eth0	192.0.2.100	10.1.1.100	HTTP	*	*	*	*
2. NAT rule SSH	Untrust L3	Trust L3	eth0	192.0.2.100	10.1.1.101	SSH	*	*	*	*


The security rules would look like this:

Name	Zone	Address	Zone	Address	Application	Service	Action
1. Webserver access	Untrust L3	any	Trust L3	10.1.1.100	HTTP	allow	allow
2. SSH access	Untrust L3	any	Trust L3	10.1.1.101	SSH	allow	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.



To verify the translations, use the CLI command **show session all filter destination 80.80.80.80**. 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 example 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 192.0.2.0/24 and 172.16.1.0/24. The link between the routers is configured in subnet 198.51.100.0/30. 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
172.16.1.0/24	198.51.100.2

Route on R2:

Destination	Next Hop
192.0.2.0/24	198.51.100.1

Now the firewall is deployed in virtual wire mode between the two Layer 3 devices. A NAT IP address pool with range 198.51.100.9 to 198.51.100.14 is configured on the firewall. All communications from clients in subnet 192.0.2.0/24 accessing servers in network 172.16.1.0/24 will arrive at R2 with a translated source address in the range 198.51.100.9 to 198.51.100.14. The response from servers will be directed to these addresses.



In order for source NAT to work, you must configure proper routing on R2, so that packets destined for other addresses are not dropped. The routing table below shows the modified routing table on R2; the route ensures traffic to the destinations 198.51.100.9-198.51.100.14 (that is, hosts on subnet 198.51.100.8/29) will be sent back through the firewall to R1.

Route on R2:

Destination	Next Hop
198.51.100.8/29	198.51.100.1

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 192.0.2.100 is statically translated to address 198.51.100.100. 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 198.51.100.100, which the firewall translates to 198.0.2.100. Any connections initiated by the server at 192.0.2.100 are translated to source IP address 198.51.100.100.



Route on R2:

Destination	Next Hop
198.51.100.100/32	198.51.100.1

Serial	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation
1	Trust	Untrust	eth1	192.0.2.100	198.51.100.100	any	nat	198.51.100.100

Virtual Wire Destination NAT Example

Clients in the Untrust zone access the server using the IP address 198.51.100.100, which the firewall translates to 192.0.2.100. Both the NAT and security policies must be configured from the Untrust zone to the Trust zone.



Route on R2:

Destination	Next Hop
198.51.100.100/32	198.51.100.1

Serial	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation
1	Untrust	Trust	eth1	198.51.100.100	192.0.2.100	any	nat	192.0.2.100

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](#)
- [Model Support for NPTv6](#)
- [Unique Local Addresses](#)
- [Reasons to Use NPTv6](#)

NPTv6 Does Not Provide Security

It is important to understand that NPTv6 does not provide security. In general, stateless network address translation does not provide any security; it provides an address translation function. NPTv6 does not hide or translate port numbers. You must set up firewall security policies correctly in each direction to ensure that traffic is controlled as you intended.

Model Support for NPTv6

NPTv6 is supported on the following models (NPTv6 with hardware lookup but packets go through the CPU): PA-7000 Series, PA-5200 Series, PA-800 firewall and PA-220 firewall. Models supported with no ability to have hardware perform a session look-up: PA-3020 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](#)
- [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).

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 in 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.

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**.

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 **Add** 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. The destination address (an address object is allowed) must be a netmask, not just an IPv6 address and not a range. The prefix length must be a value from /32 to /64, inclusive. For example, `2001:db8::/32`.

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 the prior step.



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 or enter your internal destination address.
5. Click **OK**.

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 | Commit the configuration.

Click **OK** and **Commit**.

NAT64

NAT64 provides a way to transition to IPv6 while you still need to communicate with IPv4 networks. When you need to communicate from an IPv6-only network to an IPv4 network, you use NAT64 to translate source and destination addresses from IPv6 to IPv4 and vice versa. NAT64 allows IPv6 clients to access IPv4 servers and allows IPv4 clients to access IPv6 servers. You should understand [NAT](#) before configuring NAT64.

- [NAT64 Overview](#)
- [IPv4-Embedded IPv6 Address](#)
- [DNS64 Server](#)
- [Path MTU Discovery](#)
- [IPv6-Initiated Communication](#)
- [Configure NAT64 for IPv6-Initiated Communication](#)
- [Configure NAT64 for IPv4-Initiated Communication](#)
- [Configure NAT64 for IPv4-Initiated Communication with Port Translation](#)

NAT64 Overview

You can configure two types of NAT64 translation on a Palo Alto Networks firewall; each one is doing a bidirectional translation between the two IP address families:

- The firewall supports stateful NAT64 for [IPv6-Initiated Communication](#), which maps multiple IPv6 addresses to one IPv4 address, thus preserving IPv4 addresses. (It does not support stateless NAT64, which maps one IPv6 address to one IPv4 address and therefore does not preserve IPv4 addresses.) [Configure NAT64 for IPv6-Initiated Communication](#).
- The firewall supports IPv4-initiated communication with a static binding that maps an IPv4 address and port number to an IPv6 address. [Configure NAT64 for IPv4-Initiated Communication](#). It also supports port rewrite, which preserves even more IPv4 addresses by translating an IPv4 address and port number to an IPv6 address with multiple port numbers. [Configure NAT64 for IPv4-Initiated Communication with Port Translation](#).

A single IPv4 address can be used for NAT44 and NAT64; you don't reserve a pool of IPv4 addresses for NAT64 only.


NAT64 operates on Layer 3 interfaces, subinterfaces, and tunnel interfaces. To use NAT64 on a Palo Alto Networks firewall for IPv6-initiated communication, you must have a third-party [DNS64 Server](#) or a solution in place to separate the DNS query function from the NAT function. The DNS64 server translates between your IPv6 host and an IPv4 DNS server by encoding the IPv4 address it receives from a public DNS server into an IPv6 address for the IPv6 host.

Palo Alto Networks supports the following NAT64 features:

- Hairpinning (NAT U-Turn); additionally, NAT64 prevents hairpinning loop attacks by dropping all incoming IPv6 packets that have a source prefix of 64::/n.
- Translation of TCP/UDP/ICMP packets per [RFC 6146](#) and the firewall makes a best effort to translate other protocols that don't use an application-level gateway (ALG). For example, the firewall can translate a GRE packet. This translation has the same limitation as NAT44: if you don't have an ALG for a protocol that can use a separate control and data channel, the firewall might not understand the return traffic flow.
- Translation between IPv4 and IPv6 of the ICMP length attribute of the original datagram field, per [RFC 4884](#).

IPv4-Embedded IPv6 Address

NAT64 uses an IPv4-embedded IPv6 address as described in [RFC 6052, IPv6 Addressing of IPv4/IPv6 Translators](#). An IPv4-embedded IPv6 address is an IPv6 address in which 32 bits have an IPv4 address encoded in them. The IPv6 prefix length (PL in the figure) determines where in the IPv6 address the IPv4 address is encoded, as follows:



The firewall supports translation for /32, /40, /48, /56, /64, and /96 subnets using these prefixes. A single firewall supports multiple prefixes; each NAT64 rule uses one prefix. The prefix can be the Well-Known Prefix (64:FF9B::/96) or a Network-Specific Prefix (NSP) that is unique to the organization that controls the address translator (the DNS64 device). An NSP is usually a network within the organization's IPv6 prefix. The DNS64 device typically sets the u field and suffix to zeros; the firewall ignores those fields.

DNS64 Server

If you need to use a DNS and you want to perform NAT64 translation using [IPv6-Initiated Communication](#), you must use a third-party DNS64 server or other DNS64 solution that is set up with the Well-Known Prefix or your NSP. When an IPv6 host attempts to access an IPv4 host or domain on the internet, the DNS64 server queries an authoritative DNS server for the IPv4 address mapped to that host name. The DNS server returns an Address record (A record) to the DNS64 server containing the IPv4 address for the host name.

The DNS64 server in turn converts the IPv4 address to hexadecimal and encodes it into the appropriate octets of the IPv6 prefix it is set up to use (the Well-Known Prefix or your NSP) based on the prefix length, which results in an [IPv4-Embedded IPv6 Address](#). The DNS64 server sends an AAAA record to the IPv6 host that maps the IPv4-embedded IPv6 address to the IPv4 host name.

Path MTU Discovery

IPv6 does not fragment packets, so the firewall uses two methods to reduce the need to fragment packets:

- When the firewall is translating IPv4 packets in which the DF (don't fragment) bit is zero, that indicates the sender expects the firewall to fragment packets that are too large, but the firewall doesn't fragment packets for the IPv6 network (after translation) because IPv6 doesn't fragment packets. Instead, you can configure the minimum size into which the firewall will fragment IPv4 packets before translating them. The **NAT64 IPv6 Minimum Network MTU** value is this setting, which complies with [RFC 6145, IP/ICMP Translation Algorithm](#). You can set the **NAT64 IPv6 Minimum Network MTU** to its maximum value (**Device** > **Setup** > **Session**), which causes the firewall to fragment IPv4 packets to the IPv6 minimum

size before translating them to IPv6. (The **NAT64 IPv6 Minimum Network MTU** does not change the interface MTU.)

- The other method the firewall uses to reduce fragmentation is Path MTU Discovery (PMTUD). In an IPv4-initiated communication, if an IPv4 packet to be translated has the DF bit set and the MTU for the egress interface is smaller than the packet, the firewall uses PMTUD to drop the packet and return an ICMP 'Destination Unreachable - fragmentation needed' message to the source. The source lowers the path MTU for that destination and resends the packet until successive reductions in the path MTU allow packet delivery.

IPv6-Initiated Communication

IPv6-initiated communication to the firewall is similar to source NAT for an IPv4 topology. [Configure NAT64 for IPv6-Initiated Communication](#) when your IPv6 host needs to communicate with an IPv4 server.

In the NAT64 policy rule, configure the original source to be an IPv6 host address or Any. Configure the destination IPv6 address as either the Well-Known Prefix or the NSP that the DNS64 server uses. (You do not configure the full IPv6 destination address in the rule.)

If you need to use a DNS, you need to use a [DNS64 Server](#) to convert an IPv4 DNS "A" result into an "AAAA" result merged with the NAT64 prefix. If you don't use a DNS, you need to create the address using the IPv4 destination address and the NAT64 prefix configured on the firewall, following [RFC 6052](#) rules.

For environments that use a DNS, the example topology below illustrates communication with the DNS64 server. The DNS64 server must be set up to use the Well-Known Prefix 64:FF9B::/96 or your Network-Specific Prefix, which must comply with RFC 6052 (/32, /40, /48, /56, /64, or /96).

On the translated side of the firewall, the translation type must be Dynamic IP and Port in order to implement stateful NAT64. You configure the source translated address to be the IPv4 address of the egress interface on the firewall. You do not configure the destination translation field; the firewall translates the address by first finding the prefix length in the original destination address of the rule and then based on the prefix, extracting the encoded IPv4 address from the original destination IPv6 address in the incoming packet.

Before the firewall looks at the NAT64 rule, the firewall must do a route lookup to find the destination security zone for an incoming packet. You must ensure that the NAT64 prefix can be reached through the destination zone assignment because the NAT64 prefix should not be routable by the firewall. The firewall would likely assign the NAT64 prefix to the default route or drop the NAT64 prefix because there is no route for it. The firewall will not find a destination zone because the NAT64 prefix is not in its routing table, associated with an egress interface and zone.

You must also configure a tunnel interface (with no termination point). You apply the NAT64 prefix to the tunnel and apply the appropriate zone to ensure that IPv6 traffic with the NAT64 prefix is assigned to the proper destination zone. The tunnel also has the advantage of dropping IPv6 traffic with the NAT64 prefix if the traffic does not match the NAT64 rule. Your configured routing protocol on the firewall looks up the IPv6 prefix in its routing table to find the destination zone and then looks at the NAT64 rule.

The following figure illustrates the role of the DNS64 server in the name resolution process. In this example, the DNS64 server is configured to use Well-Known Prefix 64:FF9B::/96.

1. A user at the IPv6 host enters the URL `www.abc.com`, which generates a name server lookup (nslookup) to the DNS64 server.
2. The DNS64 Server sends an nslookup to the public DNS server for `www.abc.com`, requesting its IPv4 address.
3. The DNS server returns an A record that provides the IPv4 address to the DNS64 server.
4. The DNS64 server sends an AAAA record to the IPv6 user, converting the IPv4 dotted decimal address `198.51.100.1` into `C633:6401` hexadecimal and embedding it into its own IPv6 prefix, `64:FF9B::/96`.

[198 = C6 hex; 51 = 33 hex; 100 = 64 hex; 1 = 01 hex.] The result is [IPv4-Embedded IPv6 Address](#) 64:FF9B::C633:6401.

Keep in mind that in a /96 prefix, the IPv4 address is the last four octets encoded in the IPv6 address. If the DNS64 server uses a /32, /40, /48, /56 or /64 prefix, the IPv4 address is encoded as shown in RFC 6052.



Upon the transparent name resolution, the IPv6 host sends a packet to the firewall containing its IPv6 source address and destination IPv6 address 64:FF9B::C633:6401 as determined by the DNS64 server. The firewall performs the NAT64 translation based on your NAT64 rule.



Configure NAT64 for IPv6-Initiated Communication

This configuration task and its addresses correspond to the figures in [IPv6-Initiated Communication](#).

STEP 1 | Enable IPv6 to operate on the firewall.

1. Select **Device > Setup > Session** and edit the Session Settings.
2. Select **Enable IPv6 Firewalling**.
3. Click **OK**.

STEP 2 | Create an address object for the IPv6 destination address (pre-translation).

1. Select **Objects > Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64-IPv4 Server.
3. For **Type**, select **IP Netmask** and enter the IPv6 prefix with a netmask that is compliant with RFC 6052 (/32, /40, /48, /56, /64, or /96). This is either the Well-Known Prefix or your Network-Specific Prefix that is configured on the [DNS64 Server](#).

For this example, enter 64:FF9B::/96.



The source and destination must have the same netmask (prefix length).

(You don't enter a full destination address because, based on the prefix length, the firewall extracts the encoded IPv4 address from the original destination IPv6 address in the incoming packet. In this example, the prefix in the incoming packet is encoded with C633:6401 in hexadecimal, which is the IPv4 destination address 198.51.100.1.)

4. Click **OK**.

STEP 3 | (Optional) Create an address object for the IPv6 source address (pre-translation).

1. Select **Objects > Addresses** and click **Add**.
2. Enter a **Name** for the object.

3. For **Type**, select **IP Netmask** and enter the address of the IPv6 host, in this example, 2001:DB8::5/96.
4. Click **OK**.

STEP 4 | (Optional) Create an address object for the IPv4 source address (translated).

1. Select **Objects > Addresses** and click **Add**.
2. Enter a **Name** for the object.
3. For **Type**, select **IP Netmask** and enter the IPv4 address of the firewall's egress interface, in this example, 192.0.2.1.
4. Click **OK**.

STEP 5 | Create the NAT64 rule.

1. Select **Policies > NAT** and click **Add**.
2. On the **General** tab, enter a **Name** for the NAT64 rule, for example, nat64_ipv6_init.
3. (Optional) Enter a **Description**.
4. For **NAT Type**, select **nat64**.

STEP 6 | Specify the original source and destination information.

1. For the **Original Packet**, **Add** the **Source Zone**, likely a trusted zone.
2. Select the **Destination Zone**, in this example, the Untrust zone.
3. (Optional) Select a **Destination Interface** or the default (**any**).
4. For **Source Address**, select **Any** or **Add** the address object you created for the IPv6 host.
5. For **Destination Address**, **Add** the address object you created for the IPv6 destination address, in this example, nat64-IPv4 Server.
6. (Optional) For **Service**, select **any**.

STEP 7 | Specify the translated packet information.

1. For the **Translated Packet**, in **Source Address Translation**, for **Translation Type**, select **Dynamic IP and Port**.
2. For **Address Type**, do one of the following:
 - Select **Translated Address** and **Add** the address object you created for the IPv4 source address.
 - Select **Interface Address**, in which case the translated source address is the IP address and netmask of the firewall's egress interface. For this choice, select an **Interface** and optionally an **IP Address** if the interface has more than one IP address.
3. Leave **Destination Address Translation** unselected. (The firewall extracts the IPv4 address from the IPv6 prefix in the incoming packet, based on the prefix length specified in the original destination of the NAT64 rule.)
4. Click **OK** to save the NAT64 policy rule.

STEP 8 | Configure a tunnel interface to emulate a loopback interface with a netmask other than 128.

1. Select **Network > Interfaces > Tunnel** and **Add** a tunnel.
2. For **Interface Name**, enter a numeric suffix, such as .2.
3. On the **Config** tab, select the **Virtual Router** where you are configuring NAT64.
4. For **Security Zone**, select the destination zone associated with the IPv4 server destination (Trust zone).
5. On the **IPv6** tab, select **Enable IPv6 on the interface**.
6. Click **Add** and for the **Address**, select **New Address**.
7. Enter a **Name** for the address.
8. (Optional) Enter a **Description** for the tunnel address.
9. For **Type**, select **IP Netmask** and enter your IPv6 prefix and prefix length, in this example, 64:FF9B::/96.

10. Click **OK**.
11. Select **Enable address on interface** and click **OK**.
12. Click **OK**.
13. Click **OK** to save the tunnel.

STEP 9 | Create a security policy to allow NAT traffic from the trust zone.

1. Select **Policies > Security** and **Add** a rule **Name**.
2. Select **Source** and **Add a Source Zone**; select **Trust**.
3. For **Source Address**, select **Any**.
4. Select **Destination** and **Add a Destination Zone**; select **Untrust**.
5. For **Application**, select **Any**.
6. For **Actions**, select **Allow**.
7. Click **OK**.

STEP 10 | Commit your changes.

Click **Commit**.


STEP 11 | Troubleshoot or view a NAT64 session.

```
> show session id <session-id>
```

Configure NAT64 for IPv4-Initiated Communication

IPv4-initiated communication to an IPv6 server is similar to destination NAT in an IPv4 topology. The destination IPv4 address maps to the destination IPv6 address through a one-to-one, static IP translation (not a many-to-one translation).

The firewall encodes the source IPv4 address into Well-Known Prefix 64:FF9B::/96 as defined in RFC 6052. The translated destination address is the actual IPv6 address. The use case for IPv4-initiated communication is typically when an organization is providing access from the public, untrust zone to an IPv6 server in the organization's DMZ zone. This topology does not use a DNS64 server.



STEP 1 | Enable IPv6 to operate on the firewall.

1. Select **Device > Setup > Session** and edit the Session Settings.
2. Select **Enable IPv6 Firewalling**.
3. Click **OK**.

STEP 2 | (Optional) When an IPv4 packet has its DF bit set to zero (and because IPv6 does not fragment packets), ensure the translated IPv6 packet does not exceed the path MTU for the destination IPv6 network.

1. Select **Device > Setup > Session** and edit Session Settings.
2. For **NAT64 IPv6 Minimum Network MTU**, enter the smallest number of bytes into which the firewall will fragment IPv4 packets for translation to IPv6 (range is 1280-9216, default is 1280).



If you don't want the firewall to fragment an IPv4 packet prior to translation, set the MTU to 9216. If the translated IPv6 packet still exceeds this value, the firewall drops the packet and issues an ICMP packet indicating destination unreachable - fragmentation needed.

3. Click **OK**.

STEP 3 | Create an address object for the IPv4 destination address (pre-translation).

1. Select **Objects > Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64_ip4server.
3. For **Type**, select **IP Netmask** and enter the IPv4 address of the firewall interface in the Untrust zone. The address must use no netmask or a netmask of /32 only. This example uses 198.51.19.1/32.
4. Click **OK**.

STEP 4 | Create an address object for the IPv6 source address (translated).

1. Select **Objects > Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64_ip6source.
3. For **Type**, select **IP Netmask** and enter the NAT64 IPv6 address with a netmask that is compliant with RFC 6052 (/32, /40, /48, /56, /64, or /96).

For this example, enter 64:FF9B::/96.

(The firewall encodes the prefix with the IPv4 source address 192.1.2.8, which is C001:0208 in hexadecimal.)

4. Click **OK**.

STEP 5 | Create an address object for the IPv6 destination address (translated).

1. Select **Objects > Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64_server_2.
3. For **Type**, select **IP Netmask** and enter the IPv6 address of the IPv6 server (destination). The address must use no netmask or a netmask of /128 only. This example uses 2001:DB8::2/128.
4. Click **OK**.

STEP 6 | Create the NAT64 rule.

1. Select **Policies > NAT** and click **Add**.
2. On the **General** tab, enter a **Name** for the NAT64 rule, for example, nat64_ip4_init.
3. For **NAT Type**, select **nat64**.

STEP 7 | Specify the original source and destination information.

1. For the **Original Packet**, **Add** the **Source Zone**, likely an untrust zone.
2. Select the **Destination Zone**, likely a trust or DMZ zone.
3. For **Source Address**, select **Any** or **Add** the address object for the IPv4 host.
4. For **Destination Address**, **Add** the address object for the IPv4 destination, in this example, nat64_ip4server.
5. For **Service**, select **any**.

STEP 8 | Specify the translated packet information.

1. For the **Translated Packet**, in the **Source Address Translation**, **Translation Type**, select **Static IP**.
2. For **Translated Address**, select the source translated address object you created, nat64_ip6source.
3. For **Destination Address Translation**, for **Translated Address**, specify a single IPv6 address (the address object, in this example, nat64_server_2, or the IPv6 address of the server).
4. Click **OK**.

STEP 9 | Create a security policy to allow the NAT traffic from the Untrust zone.

1. Select **Policies > Security** and **Add** a rule **Name**.
2. Select **Source** and **Add** a **Source Zone**; select **Untrust**.
3. For **Source Address**, select **Any**.
4. Select **Destination** and **Add** a **Destination Zone**; select **DMZ**.
5. For **Actions**, select **Allow**.
6. Click **OK**.

STEP 10 | Commit your changes.


Click **Commit**.

STEP 11 | Troubleshoot or view a NAT64 session.

```
> show session id <session-id>
```

Configure NAT64 for IPv4-Initiated Communication with Port Translation

This task builds on the task to [Configure NAT64 for IPv4-Initiated Communication](#), but the organization controlling the IPv6 network prefers to translate the public destination port number to an internal destination port number and thereby keep it private from users on the IPv4 untrust side of the firewall. In this example, port 8080 is translated to port 80. To do that, in the Original Packet of the NAT64 policy rule, create a new Service that specifies the destination port is 8080. For the Translated Packet, the translated port is 80.



STEP 1 | Enable IPv6 to operate on the firewall.

1. Select **Device > Setup > Session** and edit the Session Settings.
2. Select **Enable IPv6 Firewalling**.
3. Click **OK**.

STEP 2 | (Optional) When an IPv4 packet has its DF bit set to zero (and because IPv6 does not fragment packets), ensure the translated IPv6 packet does not exceed the path MTU for the destination IPv6 network.

1. Select **Device > Setup > Session** and edit Session Settings.
2. For **NAT64 IPv6 Minimum Network MTU**, enter the smallest number of bytes into which the firewall will fragment IPv4 packets for translation to IPv6 (range is 1280-9216, default is 1280).



If you don't want the firewall to fragment an IPv4 packet prior to translation, set the MTU to 9216. If the translated IPv6 packet still exceeds this value, the firewall drops the packet and issues an ICMP packet indicating destination unreachable - fragmentation needed.

3. Click **OK**.

STEP 3 | Create an address object for the IPv4 destination address (pre-translation).

1. Select **Objects > Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64_ip4server.
3. For **Type**, select **IP Netmask** and enter the IPv4 address and netmask of the firewall interface in the Untrust zone. This example uses 198.51.19.1/24.
4. Click **OK**.

STEP 4 | Create an address object for the IPv6 source address (translated).

1. Select **Objects > Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64_ip6source.
3. For **Type**, select **IP Netmask** and enter the NAT64 IPv6 address with a netmask that is compliant with RFC 6052 (/32, /40, /48, /56, /64, or /96).

For this example, enter 64:FF9B::/96.

(The firewall encodes the prefix with the IPv4 source address 192.1.2.8, which is C001:0208 in hexadecimal.)

4. Click **OK**.

STEP 5 | Create an address object for the IPv6 destination address (translated).

1. Select **Objects > Addresses** and click **Add**.
2. Enter a **Name** for the object, for example, nat64_server_2.
3. For **Type**, select **IP Netmask** and enter the IPv6 address of the IPv6 server (destination). This example uses 2001:DB8::2/64.



The source and destination must have the same netmask (prefix length).

4. Click **OK**.

STEP 6 | Create the NAT64 rule.

1. Select **Policies > NAT** and click **Add**.
2. On the **General** tab, enter a **Name** for the NAT64 rule, for example, nat64_ipv4_init.
3. For **NAT Type**, select **nat64**.

STEP 7 | Specify the original source and destination information, and create a service to limit the translation to a single ingress port number.

1. For the **Original Packet**, Add the **Source Zone**, likely an untrust zone.
2. Select the **Destination Zone**, likely a trust or DMZ zone.
3. For **Service**, select **New Service**.
4. Enter a **Name** for the Service, such as Port_8080.
5. Select **TCP** as the **Protocol**.
6. For **Destination Port**, enter 8080.
7. Click **OK** to save the Service.
8. For **Source Address**, select **Any** or **Add** the address object for the IPv4 host.
9. For **Destination Address**, **Add** the address object for the IPv4 destination, in this example, nat64_ip4server.

STEP 8 | Specify the translated packet information.

1. For the **Translated Packet**, in the **Source Address Translation**, **Translation Type**, select **Static IP**.
2. For **Translated Address**, select the source translated address object you created, nat64_ip6source.

-
3. For **Destination Address Translation**, for **Translated Address**, specify a single IPv6 address (the address object, in this example, `nat64_server_2`, or the IPv6 address of the server).
 4. Specify the private destination **Translated Port** number to which the firewall translates the public destination port number, in this example, 80.
 5. Click **OK**.

STEP 9 | Create a security policy to allow the NAT traffic from the Untrust zone.

1. Select **Policies > Security** and **Add** a rule **Name**.
2. Select **Source** and **Add** a **Source Zone**; select **Untrust**.
3. For **Source Address**, select **Any**.
4. Select **Destination** and **Add** a **Destination Zone**; select **DMZ**.
5. For **Actions**, select **Allow**.
6. Click **OK**.

STEP 10 | Commit your changes.

Click **Commit**.

STEP 11 | Troubleshoot or view a NAT64 session.

```
> show session id <session-id>
```

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 Model, 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, the hash can be based on the source and destination addresses, or the hash can be based on the source address only (in PAN-OS 8.0.3 and later releases). Using an IP hash based on only the source address causes all sessions belonging to the same source IP address to always take the same path from available multiple paths. Thus the path is considered sticky and is easier to troubleshoot if necessary. 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 Model, Interface, and IP Routing Support

ECMP is supported on all Palo Alto Networks firewall models, with hardware forwarding support on the PA-7000 Series, PA-5200 Series, PA-3200 Series firewalls. 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.

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.

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.

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** list:

- **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**—There are two IP hash methods that determine which ECMP route to use (select hash options in Step 5):
 - Use a hash of the source address (available in PAN-OS 8.0.3 and later releases).
 - Use a hash of the source and destination IP addresses (the default IP hash method).
- **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.

STEP 5 | (IP Hash only) Configure IP Hash options.

If you selected **IP Hash** as the **Method**:

1. Select **Use Source Address Only** (available in PAN-OS 8.0.3 and later releases) if you want to ensure all sessions belonging to the same source IP address always take the same path from available multiple paths. This IP hash option provides path stickiness and eases troubleshooting. If you don't select this option or you're using a release prior to PAN-OS 8.0.3, the IP hash is based on the source and destination IP addresses (the default IP hash method).



*If you select **Use Source Address Only**, you shouldn't push the configuration from Panorama to firewalls running PAN-OS 8.0.2, 8.0.1, or 8.0.0.*

2. Select **Use Source/Destination Ports** if you want to use source or destination port numbers in the **IP Hash** calculation.



*Enabling this option along with **Use Source Address Only** will randomize path selection even for sessions belonging to the same source IP address.*

3. 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**.
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 | Commit your changes.


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.



In the following figure, two ECMP paths to a destination go through two firewalls belonging to two different ISPs in different BGP autonomous systems.



STEP 1 | Configure ECMP.

See [Configure ECMP on a Virtual Router](#).

STEP 2 | For BGP routing, enable ECMP over multiple autonomous systems.

1. Select **Network > Virtual Routers** and select the virtual router on which to enable ECMP for multiple BGP autonomous systems.
2. Select **BGP > Advanced** and select **ECMP Multiple AS Support**.

STEP 3 | Commit your changes.

Click **OK** and **Commit**.

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.

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, multiple routes to the same Destination (out a different Interface) have the E flag. An asterisk () denotes the preferred path for the ECMP group.*

LLDP

Palo Alto Networks firewalls support Link Layer Discovery Protocol (LLDP), which functions at the link layer to discover neighboring devices and their capabilities. LLDP allows the firewall and other network devices to send and receive LLDP data units (LLDPDUs) to and from neighbors. The receiving device stores the information in a MIB, which the Simple Network Management Protocol (SNMP) can access. LLDP makes troubleshooting easier, especially for virtual wire deployments where the firewall would typically go undetected by a ping or traceroute.

- [LLDP Overview](#)
- [Supported TLVs in LLDP](#)
- [LLDP Syslog Messages and SNMP Traps](#)
- [Configure LLDP](#)
- [View LLDP Settings and Status](#)
- [Clear LLDP Statistics](#)

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.

Panorama and the WildFire appliance are 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 models 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 0-65,535). The value is a multiple of the LLDP Hold Time Multiplier. When the TTL value is 0, the information associated with the device is no longer valid and the firewall removes that entry from the MIB.
End of LLDPDU TLV	0	Indicates the end of the TLVs in the LLDP Ethernet frame.

The following table lists the optional TLVs that the Palo Alto Networks firewall supports:

Optional TLVs	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).

Optional TLVs	TLV Type	Purpose and Notes Regarding Firewall Implementation
		- 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.

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**.

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 **Add a Name** for the LLDP profile.
2. For **Mode**, select **transmit-receive** (default), **transmit-only**, or **receive-only**.
3. 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**.
4. For Optional TLVs, select the TLVs you want transmitted:
 - **Port Description**
 - **System Name**
 - **System Description**
 - **System Capabilities**
5. (Optional) Select **Management Address** to add one or more management addresses and **Add a Name**.
6. 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.
7. Select **IPv4** or **IPv6**, and in the adjacent field, select an IP address from the list (which lists the addresses configured on the selected interface), or enter an address.
8. Click **OK**.
9. 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.
10. 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 | Commit your changes.

Click **Commit**.

View LLDP Settings and Status

Perform the following procedure to view LLDP settings and status.

STEP 1 | View LLDP global settings.

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 second step in [Configure LLDP](#).

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.

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 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) ([RFC 5880](#)), 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](#), <http://www.rfc-editor.org/rfc/rfc5881.txt>. 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](#), <https://www.rfc-editor.org/rfc/rfc5883.txt>. However, PAN-OS does not support authentication.

- [BFD Model, Interface, and Client Support](#)
- [Non-Supported RFC Components of BFD](#)
- [BFD for Static Routes](#)
- [BFD for Dynamic Routing Protocols](#)

BFD Model, Interface, and Client Support

The following firewall models do not support BFD: PA-800 Series, PA-220, and VM-50 firewalls. The models that do support BFD support 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 and IPv6 (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.*

Configure BFD

After you read the [BFD Overview](#), which includes firewall models and interfaces supported, perform the following before configuring BFD:

- Configure one or more [Virtual Routers](#).
- Configure one or more [Static Routes](#) if you are applying BFD to static routes.
- Configure 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.

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.

STEP 2 | Configure BFD intervals.

1. 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-5200 Series firewalls is 50; minimum on VM-Series firewall is 200. Maximum is 2,000; default is 1,000.



The recommendation is to set the Desired Minimum Tx Interval on a PA-7000 Series firewall to 100 or greater; a value less than 100 is at risk of causing BFD flaps.



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.

2. 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-5200 Series firewalls is 50; minimum on VM-Series firewall is 200. Maximum is 2,000; default is 1,000.



The recommendation is to set the Required Minimum Rx Interval on a PA-7000 Series firewall to 100 or greater; a value less than 100 is at risk of causing BFD flaps.

STEP 3 | Configure the BFD Detection Time Multiplier.

Enter the **Detection Time Multiplier**. The local system calculates the detection time as the **Detection Time Multiplier** received from the remote system multiplied by the agreed transmit interval of the remote system (the greater of the **Required Minimum Rx Interval** and the last received **Desired Minimum Tx Interval**). If BFD does not receive a BFD control packet from its peer before the detection time expires, a failure has occurred. Range is 2 to 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.

STEP 4 | Configure the BFD hold time.

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.

STEP 5 | (Optional—For a BGP IPv4 implementation only) Configure hop-related settings for the BFD profile.

1. Select **Multihop** to enable BFD over BGP multihop.
2. 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.

STEP 6 | Save the BFD profile.

Click **OK**.

STEP 7 | (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.

STEP 8 | (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** list, 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:

1. Select the **Peer Group** tab.
2. Select a peer group.
3. Select a peer.
4. In the **BFD** list, 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.

5. Click **OK**.
6. Click **OK**.

A BFD column on the BGP - Peer Group/Peer list indicates the BFD profile configured for the interface.

STEP 9 | (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** list, 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:

1. Select the **Areas** tab and select an area.
2. On the **Interface** tab, select an interface.
3. In the **BFD** list, 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.

4. Click **OK**.
5. Click **OK**.

A BFD column on the OSPF **Interface** tab indicates the BFD profile configured for the interface.

STEP 10 | (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** list, 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:
 1. Select the **Interfaces** tab and select an interface.
 2. In the **BFD** list, 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.

3. Click **OK**.
5. Click **OK**.

The BFD column on the **Interface** tab indicates the BFD profile configured for the interface.

STEP 11 | Commit the configuration.

Click **Commit**.

STEP 12 | 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 13 | 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 14 | (Optional) Clear BFD transmit, receive, and drop counters.

```
clear routing bfd counters session-id all | <1-1024>
```

STEP 15 | (Optional) Clear BFD sessions for debugging.

```
clear routing bfd session-state session-id all | <1-1024>
```

Reference: BFD Details

To see the following [BFD](#) information for a virtual router, you can refer Step [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.

Name	Value (Example)	Description
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
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).

Name	Value (Example)	Description
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.
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.

Name	Value (Example)	Description
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.

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](#)
- [Control Specific ICMP or ICMPv6 Types and Codes](#)
- [Configure Session Timeouts](#)
- [Session Distribution Policies](#)
- [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 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 \(MSS\)](#)


You can use [Zone Protection Profiles](#) on the firewall to configure packet-based attack protection and thereby drop IP, TCP, and IPv6 packets with undesirable characteristics or strip undesirable options from packets before allowing them into the zone. You can also configure flood protection, specifying the rate of SYN connections per second (not matching an existing session) that trigger an alarm, cause the firewall to randomly drop SYN packets or use SYN cookies, and cause the firewall to drop SYN packets that exceed the maximum rate.

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.
- 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.



*(PAN-OS 9.0.9 and later 9.0 releases) You can configure the firewall globally to fragment IPv4 packets that exceed the egress interface MTU, even when the DF bit is set in the packet. Enable this for Layer 3 physical interfaces and IPsec tunnel interfaces using the CLI command **debug dataplane set ip4-df-ignore yes**. Restore the firewall to the default behavior by using the CLI command **debug dataplane set ipv4-df-ignore no**.*

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
- 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, in [Configure Session Settings](#), see [10](#)

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.

A UDP datagram is encapsulated in an IP packet. 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).

You can use [Zone Protection Profiles](#) on the firewall to configure flood protection and thereby specify the rate of UDP connections per second (not matching an existing session) that trigger an alarm, trigger the firewall to randomly drop UDP packets, and cause the firewall to drop UDP packets that exceed the maximum rate. (Although UDP is connectionless, the firewall tracks UDP datagrams in IP packets on a session basis; therefore if the UDP packet doesn't match an existing session, it is considered a new session and it counts as a connection toward the thresholds.)

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. You can control ICMPv4 and ICMPv6 packets in several ways:

- Create [Security Policy Rules Based on ICMP and ICMPv6 Packets](#) and select the **icmp** or **ipv6-icmp** application in the rule.
- Control [ICMPv6 Rate Limiting](#) when you [Configure Session Settings](#).
- Use [Zone Protection Profiles](#) to configure flood protection, specifying the rate of ICMP or ICMPv6 connections per second (not matching an existing session) that trigger an alarm, trigger the firewall to randomly drop ICMP or ICMPv6 packets, and cause the firewall to drop ICMP or ICMPv6 packets that exceed the maximum rate.
- Use [Zone Protection Profiles](#) to configure packet based attack protection:
 - For ICMP, you can drop certain types of packets or suppress the sending of certain packets.
 - For ICMPv6 packets (Types 1, 2, 3, 4, and 137), you can specify that the firewall use the ICMP session key to match a security policy rule, which determines whether the ICMPv6 packet is allowed or not. (The firewall uses the security policy rule, overriding the default behavior of using the embedded packet to determine a session match.) When the firewall drops ICMPv6 packets that match a security policy rule, the firewall logs the details in Traffic logs.

Security Policy Rules Based on ICMP and ICMPv6 Packets

The firewall forwards ICMP or ICMPv6 packets only if a security policy rule allows the session (as the firewall does for other packet types). The firewall determines a session match in one of two ways, depending on whether the packet is an ICMP or ICMPv6 error packet or redirect packet as opposed to an ICMP or ICMPv6 informational packet:

- **ICMP Types 3, 5, 11, and 12 and ICMPv6 Types 1, 2, 3, 4, and 137**—The firewall by default looks up the embedded IP packet bytes of information from the original datagram that caused the error (the invoking packet). If the embedded packet matches an existing session, the firewall forwards or drops the ICMP or ICMPv6 packet according to the action specified in the security policy rule that matches that same session. (You can use [Zone Protection Profiles](#) with packet based attack protection to override this default behavior for the ICMPv6 types.)
- **Remaining ICMP or ICMPv6 Packet Types**—The firewall treats the ICMP or ICMPv6 packet as if it belongs to a new session. If a security policy rule matches the packet (which the firewall recognizes as an **icmp** or **ipv6-icmp** session), the firewall forwards or drops the packet based on the security policy rule action. Security policy counters and traffic logs reflect the actions.

If no security policy rule matches the packet, the firewall applies its default security policy rules, which allow intrazone traffic and block interzone traffic (logging is disabled by default for these rules).



Although you can override the default rules to enable logging or change the default action, we don't recommend you change the default behavior for a specific case because it will impact all traffic that those default rules affect. Instead, create security policy rules to control and log ICMP or ICMPv6 packets explicitly.

There are two ways to create explicit security policy rules to handle ICMP or ICMPv6 packets that are not error or redirect packets:

- **Create a security policy rule to allow (or deny) all ICMP or ICMPv6 packets**—In the security policy rule, specify the application **icmp** or **ipv6-icmp**; the firewall allows (or denies) all IP packets matching the ICMP protocol number (1) or ICMPv6 protocol number (58), respectively, through the firewall.
- **Create a custom application and a security policy rule to allow (or deny) packets from or to that application**—This more granular approach allows you to [Control Specific ICMP or ICMPv6 Types and Codes](#).

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 don't flood the network segments protected by the firewall.

First the global **ICMPv6 Error Packet Rate (per sec)** controls the rate at which ICMPv6 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 ICMPv6 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 ICMPv6 message that can be sent. The token count is decremented each time an ICMPv6 message is sent; when the bucket reaches zero tokens, no more ICMPv6 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](#).

Control Specific ICMP or ICMPv6 Types and Codes

Use this task to create a custom ICMP or ICMPv6 application and then create a security policy rule to allow or deny that application.

STEP 1 | Create a custom application for ICMP or ICMPv6 message types and codes.

1. Select **Object > Applications** and **Add** a custom application.
2. On the **Configuration** tab, enter a **Name** for the custom application and a **Description**. For example, enter the name ping6.
3. For **Category**, select **networking**.
4. For **Subcategory**, select **ip-protocol**.
5. For **Technology**, select **network-protocol**.
6. Click **OK**.
7. On the **Advanced** tab, select **ICMP Type** or **ICMPv6 Type**.
8. For **Type**, enter the number (range is 0-255) that designates the ICMP or ICMPv6 message type you want to allow or deny. For example, Echo Request message (ping) is 128.
9. If the Type includes codes, enter the **Code** number (range is 0-255) that applies to the **Type** value you want to allow or deny. Some **Type** values have Code 0 only.

10. Click **OK**.

STEP 2 | Create a Security policy rule that allows or denies the custom application you created.

[Create a Security Policy Rule](#). On the **Application** tab, specify the name of the custom application you just created.

STEP 3 | Commit your changes.

Click **Commit**.

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. 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. The timeouts are global, meaning they apply to all of the sessions of that type on the firewall.

You can also configure a global ARP cache timeout setting, which controls how long the firewall keeps ARP entries (IP address-to-hardware addresses mappings) in its cache.

In addition to the global settings, you can 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.



If you change the TCP or UDP timers at the application level, these timers for predefined applications and shared custom applications will be implemented across all virtual systems. If you need an application's timers to be different for a virtual system, you must create a custom application, assign it unique timers, and then assign the custom application to a unique virtual system.

Perform the following task 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.

STEP 1 | Access the session timeouts.

Select **Device** > **Setup** > **Session** and edit the Session Timeouts.

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 to 15,999,999; 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 to 15,999,999; 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 to 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 to 15,999,999; 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](#).

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 to 15,999,999.
- **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: 3,600. Range: 1 to 15,999,999.
- **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 to 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 to 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 to 604,800.
- **TCP Time Wait**—Maximum length of time after receiving the second FIN or a RST. Default: 15. Range: 1 to 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 to 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 to 15,999,999.
- **UDP**—Maximum length of time that a UDP session remains open without a UDP response. Default: 30. Range: 1 to 15,999,999.
- See also the **Scan** timeout in the section [\(Optional\) Change miscellaneous timeouts](#).

STEP 5 | (Optional) Change ICMP timeouts.

- **ICMP**—Maximum length of time that an ICMP session can be open without an ICMP response. Default: 6. Range: 1 to 15,999,999.
- See also the **Discard Default** and **Scan** timeout in the section [\(Optional\) Change miscellaneous timeouts](#).

STEP 6 | Click **OK** and **Commit**.

STEP 7 | (Optional) Change the ARP cache timeout.

1. Access the CLI and specify how many seconds the firewall keeps ARP entries in its cache. Use the operational command **set system setting arp-cache-timeout <value>**, where the range is 60 to 65,535; default is 1,800.

If you decrease the timeout and existing entries in the cache have a TTL greater than the new timeout, the firewall removes those entries and refreshes the ARP cache. If you increase the timeout and existing entries have a TTL less than the new timeout, they expire according to the TTL and the firewall caches new entries with the larger timeout value.

2. View the ARP cache timeout setting with the operational CLI command **show system setting arp-cache-timeout**.

Configure Session Settings

This topic describes various settings for sessions other than timeout values. Perform these tasks if you need to change the default 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 9,216 bytes and are available on certain models.
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 1,500 bytes; the range is 576 to 1,500 bytes.
 - If you enabled jumbo frames, the **Global MTU** defaults to 9,192 bytes; the range is 9,192 to 9,216 bytes.



Jumbo Frames can take up to five times more memory compared to normal packets and can reduce the number of available packet-buffers by 20%. This reduces the queue sizes dedicated for out of order, application identification, and other such packet processing tasks. As of PAN-OS 8.1, if you enable the jumbo frame global MTU configuration and reboot your firewall, packet buffers are then redistributed to process jumbo frames more efficiently.

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.

STEP 5 | Tune NAT session settings.

- **NAT64 IPv6 Minimum Network MTU**—Sets the global MTU for IPv6 translated traffic. The default of 1,280 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 model](#).

- 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 model 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.

STEP 7 | Enable packet buffer protection.

1. Select **Packet Buffer Protection** to enable the firewall to take action against sessions that can overwhelm the its packet buffer and causes legitimate traffic to be dropped.
2. If you enable packet buffer protection, you can tune the thresholds and timers that dictate how the firewall responds to packet buffer abuse.
 - **Alert (%)**: When packet buffer utilization exceeds this threshold, the firewall creates a log event. The threshold is set to 50% by default and the range is 0% to 99%. If the value is set to 0%, the firewall does not create a log event.
 - **Activate (%)**: When a packet buffer utilization exceeds this threshold, the firewall applies random early drop (RED) to abusive sessions. The threshold is set to 50% by default and the range is 0% to 99%. If the value is set to 0%, the firewall does not apply RED.



Alert events are recorded in the system log. Events for dropped traffic, discarded sessions, and blocked IP address are recorded in the threat log.

- **Block Hold Time (sec)**: The amount of time a RED-mitigated session is allowed to continue before it is discarded. By default, the block hold time is 60 seconds. The range is 0 to 65,535 seconds. If the value is set to 0, the firewall does not discard sessions based on packet buffer protection.
- **Block Duration (sec)**: This setting defines how long a session is discarded or an IP address is blocked. The default is 3,600 seconds with a range of 0 seconds to 15,999,999 seconds. If this value is set to 0, the firewall does not discard sessions or block IP addresses based on packet buffer protection.

STEP 8 | 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 9 | Save the session settings.

Click **OK**.

STEP 10 | 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 > Other Info**.
3. Select **Adjust TCP MSS** and enter a value for one or both of the following:
 - **IPv4 MSS Adjustment Size** (range is 40 to 300 bytes; default is 40 bytes).
 - **IPv6 MSS Adjustment Size** (range is 60 to 300 bytes; default is 60 bytes).
4. Click **OK**.

STEP 11 | Commit your changes.

Click **Commit**.

STEP 12 | Reboot the firewall after changing the jumbo frame configuration.

1. Select **Device > Setup > Operations**.
2. Click **Reboot Device**.

Session Distribution Policies

Session distribution policies define how PA-5200 and PA-7000 Series firewalls distribute security processing (App-ID, Content-ID, URL filtering, SSL decryption, and IPSec) among dataplane processors (DPs) on the firewall. Each policy is specifically designed for a certain type of network environment and firewall configuration to ensure that the firewall distributes sessions with maximum efficiency. For example, the Hash session distribution policy is best fit for environments that use large scale source NAT.

The number of DPs on a firewall varies based on the firewall model:

Firewall Model	Dataplane Processor(s)
PA-7000 Series	Depends on the number of installed Network Processing Cards (NPCs). Each NPC has multiple dataplane processors (DPs) and you can install multiple NPCs in the firewall.
PA-5220 firewall	1  <i>The PA-5220 firewall has only one DP so sessions distribution policies do not have an effect. Leave the policy set to the default (round-robin).</i>

Firewall Model	Dataplane Processor(s)
PA-5250 firewall	2
PA-5260 and PA-5280 firewalls	3

The following topics provide information about the available session distribution policies, how to change an active policy, and how to view session distribution statistics.

- [Session Distribution Policy Descriptions](#)
- [Change the Session Distribution Policy and View Statistics](#)

Session Distribution Policy Descriptions

The following table provides information about [Session Distribution Policies](#) to help you decide which policy best fits your environment and firewall configuration.

Session Distribution Policy	Description
Fixed	Allows you to specify the dataplane processor (DP) that the firewall will use for security processing. Use this policy for debugging purposes.
Hash	The firewall distributes sessions based on a hash of the source address or destination address. Hash based distribution improves the efficiency of NAT address resource management and reduces latency for NAT session setup by avoiding potential IP address or port conflicts. Use this policy in environments that use large scale source NAT with dynamic IP translation or Dynamic IP and Port translation or both. When using dynamic IP translation, select the source address option. When using dynamic IP and port translation, select the destination address option.
Ingress-slot (default on PA-7000 Series firewalls)	(PA-7000 Series firewalls only) New sessions are assigned to a DP on the same NPC on which the first packet of the session arrived. The selection of the DP is based on the session-load algorithm but, in this case, sessions are limited to the DPs on the ingress NPC. Depending on the traffic and network topology, this policy generally decreases the odds that traffic will need to traverse the switch fabric. Use this policy to reduce latency if both ingress and egress are on the same NPC. If the firewall has a mix of NPCs (PA-7000 20G and PA-7000 20GXM for example), this policy can isolate the increased capacity to the corresponding NPCs and help to isolate the impact of NPC failures.
Random	The firewall randomly selects a DP for session processing.

Session Distribution Policy	Description
Round-robin (default on PA-5200 Series firewalls)	The firewall selects the dataplane processor based on a round-robin algorithm between active dataplanes so that input, output, and security processing functions are shared among all dataplanes. Use this policy in low to medium demand environments where a simple and predictable load balancing algorithm will suffice. In high demand environments, we recommend that you use the session-load algorithm.
Session-load	This policy is similar to the round-robin policy but uses a weight-based algorithm to determine how to distribute sessions to achieve balance among the DPs. Because of the variability in the lifetime of a session, the DPs may not always experience an equal load. For example, if the firewall has three DPs and DP0 is at 25% of capacity, DP1 is at 25%, and DP2 is at 50%, new session assignment will be weighted towards the DP with the lower capacities. This helps improve load balancing over time. Use this policy in environments where sessions are distributed across multiple NPC slots, such as in an inter-slot aggregate interface group or environments with asymmetric forwarding. You can also use this policy or the ingress-slot policy if the firewall has a combination of NPCs with different session capacities (such as a combination of PA-7000 20G and PA-7000 20GXM NPCs).
Symmetric-hash	(PA-5200 Series and PA-7000 Series firewalls running PAN-OS 8.0 or later) The firewall selects the DP by a hash of sorted source and destination IP addresses. This policy provides the same results for server-to-client (s2c) and client-to-server (c2s) traffic (assuming the firewall does not use NAT). Use this policy in high-demand IPSec or GTP deployments. With these protocols, each direction is treated as a unidirectional flow where the flow tuples cannot be derived from each other. This policy improves performance and reduces latency by ensuring that both directions are assigned to the same DP, which removes the need for inter-DP communication.

Change the Session Distribution Policy and View Statistics

The following table describes how to view and change the active [Session Distribution Policies](#) and describes how to view session statistics for each dataplane processor (DP) in the firewall.

Task	Command
Show the active session distribution policy.	Use the show session distribution policy command to view the active session distribution policy.

Task	Command
	The following output is from a PA-7080 firewall with four NPCs installed in slots 2, 10, 11, and 12 with the ingress-slot distribution policy enabled: <pre>> show session distribution policy</pre> Ownership Distribution Policy: ingress-slot Flow Enabled Line Cards: [2, 10, 11, 12]Packet Processing Enabled Line Cards: [2, 10, 11, 12]
Change the active session distribution policy.	Use the set session distribution-policy <policy> command to change the active session distribution policy. For example, to select the session-load policy, enter the following command: <pre>> set session distribution-policy session-load</pre>
View session distribution statistics.	Use the show session distribution statistics command to view the dataplane processors (DPs) on the firewall and the number of sessions on each active DP. The following output is from a PA-7080 firewall: <pre>> show session distribution statistics DP Active Dispatched Dispatched/sec ----- s1dp0 78698 7829818 1473 s1dp1 78775 7831384 1535 s3dp0 7796 736639 1488 s3dp1 7707 737026 1442</pre> The DP Active column lists each dataplane on the installed NPCs. The first two characters indicate the slot number and the last three characters indicate the dataplane number. For example, s1dp0 indicates dataplane 0 on the NPC in slot 1 and s1dp1 indicates dataplane 1 on the NPC in slot1. The Dispatched column shows the total number of sessions that the dataplane processed since the last time the firewall rebooted. The Dispatched/sec column indicates the dispatch rate. If you add the numbers in the Dispatched column, the total equals the number of active sessions on the firewall. You can also view the total number of active sessions by running the show session info CLI command.  <i>The PA-5200 Series firewall output will look similar, except that the number of DPs depends on the model and there is only one NPC slot (s1).</i>

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.

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 **Add** 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** list, 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 | Commit your changes.

Click **OK** and **Commit**.

Tunnel Content Inspection

The firewall can inspect the traffic content of cleartext tunnel protocols without terminating the tunnel:

- [Generic Routing Encapsulation \(GRE\) \(RFC 2784\)](#)
- Non-encrypted IPSec traffic [[NULL Encryption Algorithm for IPSec \(RFC 2410\)](#) and transport mode AH IPSec]
- General Packet Radio Service (GPRS) Tunneling Protocol for User Data ([GTP-U](#))
- Virtual Extensible Local Area Network (VXLAN) ([RFC 7348](#))



Tunnel content inspection is for cleartext tunnels, not for VPN or LSVPN tunnels, which carry encrypted traffic.

You can use tunnel content inspection to enforce Security, DoS Protection, and QoS policies on traffic in these types of tunnels and traffic nested within another cleartext tunnel (for example, a Null Encrypted IPSec tunnel inside a GRE tunnel). You can view tunnel inspection logs and tunnel activity in the ACC to verify that tunneled traffic complies with your corporate security and usage policies.

All firewall models support tunnel content inspection for GRE, non-encrypted IPSec, and VXLAN protocols. Only [firewalls that support GTP security](#) support GTP-U tunnel content inspection—see the PAN-OS Releases by Model that Support GTP and SCTP Security in the [Compatibility Matrix](#).

- [Tunnel Content Inspection Overview](#)
- [Configure Tunnel Content Inspection](#)
- [View Inspected Tunnel Activity](#)
- [View Tunnel Information in Logs](#)
- [Create a Custom Report Based on Tagged Tunnel Traffic](#)

Tunnel Content Inspection Overview

Your firewall can inspect tunnel content anywhere on the network where you do not have the opportunity to terminate the tunnel first. As long as the firewall is in the path of a GRE, non-encrypted IPSec, GTP-U, or [VXLAN](#) tunnel, the firewall can inspect the tunnel content.

- Enterprise customers who want tunnel content inspection can have some or all of the traffic on the firewall tunneled using GRE, VXLAN, or non-encrypted IPSec. For security, QoS, and reporting reasons, you want to inspect the traffic inside the tunnel.
- Service Provider customers use GTP-U to tunnel data traffic from mobile devices. You want to inspect the inner content without terminating the tunnel protocol, and you want to record user data from your users.


The firewall supports tunnel content inspection on Ethernet interfaces, subinterfaces, AE interfaces, VLAN interfaces, and VPN and LSVPN tunnel interfaces. (The cleartext tunnel that the firewall inspects can be inside a VPN or LSVPN tunnel that terminates at the firewall, hence a VPN or LSVPN tunnel interface. In other words, when the firewall is a VPN or LSVPN endpoint, the firewall can inspect the traffic of any non-encrypted tunnel protocols that tunnel content inspection supports.)

Tunnel content inspection is supported in Layer 3, Layer 2, virtual wire, and tap deployments. Tunnel content inspection works on shared gateways and on virtual system-to-virtual system communications.

Single Tunnel



Tunnel-in-Tunnel



The preceding figure illustrates the two levels of tunnel inspection the firewall can perform. When a firewall configured with Tunnel Inspection policy rules receives a packet:

- The firewall first performs a Security policy check to determine whether the tunnel protocol (Application) in the packet is permitted or denied. (IPv4 and IPv6 packets are supported protocols inside the tunnel.)
- If the Security policy allows the packet, the firewall matches the packet to a Tunnel Inspection policy rule based on source zone, source address, source user, destination zone, and destination address. The Tunnel Inspection policy rule determines the tunnel protocols that the firewall inspects, the maximum level of encapsulation allowed (a single tunnel or a tunnel within a tunnel), whether to allow packets containing a tunnel protocol that doesn't pass strict header inspection per [RFC 2780](#), and whether to allow packets containing unknown protocols.
- If the packet passes the Tunnel Inspection policy rule's match criteria, the firewall inspects the inner content, which is subject to your Security policy (**required**) and optional policies you can specify. (The supported policy types for the original session are listed in the following table).
- If the firewall instead finds another tunnel, the firewall recursively parses the packet for the second header and is now at level two of encapsulation, so the second tunnel inspection policy rule, which matches a tunnel zone, must allow a maximum tunnel inspection level of two levels for the firewall to continue processing the packet.
 - If your rule allows two levels of inspection, the firewall performs a Security policy check on this inner tunnel and then the Tunnel Inspection policy check. The tunnel protocol you use in an inner tunnel can differ from the tunnel protocol you use in the outer tunnel.
 - If your rule doesn't allow two levels of inspection, the firewall bases its action on whether you configured it to drop packets that have more levels of encapsulation than the maximum tunnel inspection level you configured.

By default, the content encapsulated in a tunnel belongs to the same security zone as the tunnel, and is subject to the Security policy rules that protect that zone. However, you can configure a *tunnel zone*, which gives you the flexibility to configure Security policy rules for inside content that differ from the Security

policy rules for the tunnel. If you use a different tunnel inspection policy for the tunnel zone, it must always have a maximum tunnel inspection level of two levels because by definition the firewall is looking at the second level of encapsulation.

The firewall doesn't support a Tunnel Inspection policy rule that matches traffic for a tunnel that terminates on the firewall; the firewall discards packets that match the inner tunnel session. For example, when an IPSec tunnel terminates on the firewall, don't create a Tunnel Inspection policy rule that matches the tunnel you terminate. The firewall already inspects the inner tunnel traffic so no Tunnel Inspection policy rule is needed.



Although tunnel content inspection works on shared gateways and on virtual system-to-virtual system communications, you can't assign tunnel zones to shared gateways or virtual system-to-virtual system communications; they are subject to the same Security policy rules as the zones to which they belong.

Both the inner tunnel sessions and the outer tunnel sessions count toward the maximum session capacity for the firewall model.

The following table indicates with a check mark which types of policy you can apply to an outer tunnel session, an inner tunnel session, and the inside, original session:

Policy Type	Outer Tunnel Session	Inner Tunnel Session	Inside, Original Session
App-Override	✓ VXLAN Only	—	✓
DoS Protection	✓	✓	✓
NAT	✓	—	—
Policy-Based Forwarding (PBF) and Symmetric Return	✓	—	—
QoS	—	—	✓
Security (required)	✓	✓	✓
User-ID	✓	✓	✓
Zone Protection	✓	✓	✓

VXLAN is different than other protocols. The firewall can use either of two different sets of session keys to create outer tunnel sessions for VXLAN.

- VXLAN UDP Session—A six-tuple key (zone, source IP, destination IP, protocol, source port, and destination port) creates a VXLAN UDP Session.
- VNI Session—A five-tuple key that incorporates the tunnel ID (the VXLAN Network Identifier, or VNI) and uses zone, source IP, destination IP, protocol, and tunnel ID (VNI) to create a VNI Session.


You can [View Inspected Tunnel Activity](#) on the ACC or [View Tunnel Information in Logs](#). To facilitate quick viewing, configure a Monitor tag so you can monitor tunnel activity and filter log results by that tag.

The ACC tunnel activity provides data in various views. For the Tunnel ID Usage, Tunnel Monitor Tag, and Tunnel Application Usage, the data for **bytes**, **sessions**, **threats**, **content**, and **URLs** come from the Traffic Summary database. For the Tunnel User, Tunneled Source IP and Tunneled Destination IP Activity, data for **bytes** and **sessions** come from Traffic Summary database, data for **threats** come from the Threat Summary, data for **URLs** come from the URL Summary, and data for **contents** come from the Data database, which is a subset of the Threat logs.

If you enable NetFlow on the interface, NetFlow will capture statistics for the outer tunnel only, to avoid double-counting (counting bytes of both outer and inner flows).

For the Tunnel Inspection policy rule and tunnel zone capacities for your firewall model, see the [Product Selection tool](#).

The following figure illustrates a corporation that runs multiple divisions and uses different Security policies and a Tunnel Inspection policy. A Central IT team provides connectivity between regions. A tunnel connects Site A to Site C; another tunnel connects Site A to Site D. Central IT places a firewall in the path of each tunnel; the firewall in the tunnel between Sites A and C performs tunnel inspection; the firewall in the tunnel between Sites A and D has no tunnel inspection policy because the traffic is very sensitive.



Configure Tunnel Content Inspection

Perform this task to configure tunnel content inspection for a tunnel protocol that you allow through a tunnel.

STEP 1 | Create a Security policy rule to allow packets that use a specific application (such as the GRE application) through the tunnel from the source zone to the destination zone.

[Create a Security policy rule](#)



The firewall can create tunnel inspection logs at the start of a session, at the end of a session, or both. When you specify Actions for the Security policy rule, select Log at Session Start for long-lived tunnel sessions, such as GRE sessions.

STEP 2 | Create a tunnel inspection policy rule.

1. Select **Policies > Tunnel Inspection** and **Add** a policy rule.
2. On the **General** tab, enter a tunnel inspection policy rule **Name**, beginning with an alphanumeric character and containing zero or more alphanumeric, underscore, hyphen, period, and space characters.
3. (Optional) Enter a **Description**.

-
4. (Optional) For reporting and logging purposes, specify a **Tag** that identifies the packets that are subject to the Tunnel Inspection policy rule.

STEP 3 | Specify the criteria that determine the source of packets to which the tunnel inspection policy rule applies.

1. Select the **Source** tab.
2. **Add a Source Zone** from the list of zones (default is **Any**).
3. (Optional) **Add a Source Address**. You can enter an IPv4 or IPv6 address, an address group, or a Geo Region address object (**Any**).
4. (Optional) Select **Negate** to choose any addresses except those you specify.
5. (Optional) **Add a Source User** (default is **any**). **Known-user** is a user who has authenticated; an **Unknown** user has not authenticated.

STEP 4 | Specify the criteria that determine the destination of packets to which the tunnel inspection policy rule applies.

1. Select the **Destination** tab.
2. **Add a Destination Zone** from the list of zones (default is **Any**).
3. (Optional) **Add a Destination Address**. You can enter an IPv4 or IPv6 address, an address group, or a Geo Region address object (default is **Any**).

You can also configure a new address or address group.
4. (Optional) Select **Negate** to choose any addresses except those you specify.

STEP 5 | Specify the tunnel protocols that the firewall will inspect for this rule.

1. Select the **Inspection** tab.
2. **Add** one or more tunnel **Protocols** that you want the firewall to inspect:
 - **GRE**—Firewall inspects packets that use Generic Route Encapsulation (GRE) in the tunnel.
 - **GTP-U**—Firewall inspects packets that use General Packet Radio Service (GPRS) Tunneling Protocol for User Data (GTP-U) in the tunnel.
 - **Non-encrypted IPSec**—Firewall inspects packets that use non-encrypted IPSec (Null EncryVpted IPSec or transport mode AH IPSec) in the tunnel.
 - **VXLAN**—Firewall inspects packets that use the Virtual Extensible Local Area Network (VXLAN) tunneling protocol in the tunnel.

STEP 6 | Specify how many levels of encapsulation the firewall inspects and the conditions under which the firewall drops a packet.

1. Select **Inspect Options**.
2. Select the **Maximum Tunnel Inspection Levels** that the firewall will inspect:
 - **One Level** (default)—Firewall inspects content that is in the outer tunnel only.
For VXLAN, the firewall inspects a VXLAN payload to find the encapsulated content or applications within the tunnel. You must select **One Level** because VXLAN inspection only occurs on the outer tunnel.
 - **Two Levels (Tunnel In Tunnel)**—Firewall inspects content that is in the outer tunnel and content that is in the inner tunnel.
3. Select any, all, or none of the following to specify whether the firewall drops a packet under each condition:
 - **Drop packet if over maximum tunnel inspection level**—Firewall drops a packet that contains more levels of encapsulation than are configured for **Maximum Tunnel Inspection Levels**.
 - **Drop packet if tunnel protocol fails strict header check**—Firewall drops a packet that contains a tunnel protocol that uses a header that is non-compliant with the RFC for the protocol. Non-

compliant headers can indicate suspicious packets. This option causes the firewall to verify GRE headers against RFC 2890.



If your firewall is tunneling GRE with a device that implements a version of GRE older than [RFC 2890](#), you should not enable the option to Drop packet if tunnel protocol fails strict header check.

- **Drop packet if unknown protocol inside tunnel**—Firewall drops a packet that contains a protocol inside the tunnel that the firewall can't identify.

For example, if this option is selected, the firewall drops encrypted IPSec packets that match the tunnel inspection policy rule because the firewall can't read them. Thus, you can allow IPSec packets and the firewall will allow only null-encrypted IPSec and AH IPSec packets.

- **Return scanned VXLAN tunnel to source**—When traffic is redirected (steered) to the firewall, VXLAN encapsulates the packet. Traffic steering is most common in public cloud environments. Enable **Return scanned VXLAN tunnel to source** to return the encapsulated packet to the originating VXLAN tunnel endpoint (VTEP). This option is only supported on Layer 3, Layer 3 subinterface, aggregate interface Layer 3, and VLAN.

4. Click **OK**.

STEP 7 | Manage tunnel inspection policy rules.

Use the following to manage tunnel inspection policy rules:

- (Filter field)—Displays only the tunnel policy rules named in the filter field.
- **Delete**—Removes selected tunnel policy rules.
- **Clone**—An alternative to the **Add** button; duplicates the selected rule with a new name, which you can then revise.
- **Enable**—Enables the selected tunnel policy rules.
- **Disable**—Disables the selected tunnel policy rules.
- **Move**—Moves the selected tunnel policy rules up or down in the list; packets are evaluated against the rules in order from the top down.
- **Highlight Unused Rules**—Highlights tunnel policy rules that no packets have matched since the last time the firewall was restarted.

STEP 8 | (Optional) Create a tunnel source zone and tunnel destination zone for tunnel content and configure a Security policy rule for each zone.



The best practice is to create tunnel zones for your tunnel traffic. Thus, the firewall creates separate sessions for tunneled and non-tunneled packets that have the same five-tuple (source IP address and port, destination IP address and port, and protocol).



Assigning tunnel zones to tunnel traffic on a PA-5200 Series firewall causes the firewall to do tunnel inspection in software; tunnel inspection is not offloaded to hardware.

1. If you want tunnel content to be subject to Security policy rules that are different from the Security policy rules for the zone of the outer tunnel (configured earlier), select **Network > Zones** and **Add a Name** for the Tunnel Source Zone.
2. For **Location**, select the virtual system.
3. For **Type**, select **Tunnel**.
4. Click **OK**.
5. Repeat these substeps to create the Tunnel Destination Zone.
6. [Configure a Security policy rule](#) for the Tunnel Source Zone.



Because you might not know the originator of the tunnel traffic or the direction of the traffic flow and you don't want to inadvertently prohibit traffic for an application through the tunnel, specify both tunnel zones as the Source Zone and both tunnel zones as the Destination Zone in your Security policy rule, or select Any for both the source and destination zones; then specify the Applications.

7. **Configure a Security policy rule** for the Tunnel Destination Zone. The tip in the previous step for configuring a Security policy rule for the Tunnel Source Zone applies to the Tunnel Destination Zone, as well.

STEP 9 | (Optional) Specify the Tunnel Source Zone and Tunnel Destination Zone for the inner content.

1. Specify the Tunnel Source Zone and Tunnel Destination Zone (that you just added) for the inner content. Select **Policies > Tunnel Inspection** and on the **General** tab, select the **Name** of the tunnel inspection policy rule you created.
2. Select **Inspection**.
3. Select **Security Options**.
4. **Enable Security Options** (disabled by default) to cause the inner content source to belong to the **Tunnel Source Zone** you specify and to cause the inner content destination to belong to the **Tunnel Destination Zone** you specify.

If you don't **Enable Security Options**, the inner content source belongs to the same source zone as the outer tunnel source and the inner content destination belongs to the same destination zone as the outer tunnel destination, which means they are subject to the same Security policy rules that apply to those outer zones.

5. For **Tunnel Source Zone**, select the appropriate tunnel zone you created in the previous step so that the policies associated with that zone apply to the tunnel source zone. Otherwise, by default, the inner content will use the same source zone that is used in the outer tunnel and the policies of the outer tunnel source zone apply to the inner content source zone, as well.
6. For **Tunnel Destination Zone**, select the appropriate tunnel zone you created in the previous step so that the policies associated with that zone apply to the tunnel destination zone. Otherwise, by default, the inner content will use the same destination zone that is used in the outer tunnel and the policies of the outer tunnel destination zone apply to the inner content destination zone, as well.



If you configure a Tunnel Source Zone and Tunnel Destination Zone for the tunnel inspection policy rule, you should configure a specific Source Zone (in Step 3) and a specific Destination Zone (in Step 4) in the match criteria of the tunnel inspection policy rule, instead of specifying a Source Zone of Any and a Destination Zone of Any. This tip ensures the direction of zone reassignment corresponds appropriately to the parent zones.



On a PA-5200 Series or PA-7080 firewall, if you use multicast underlay while inspecting VXLAN, the inner session would be duplicated on multiple dataplanes and a race condition could happen. To avoid the drop of some packets, the following requirements apply:

- You must configure a separate tunnel content inspection rule to match outer VXLAN packets going to each VXLAN tunnel endpoint (VTEP).
- In the separate rule, you assign a tunnel zone. Using a different tunnel zone would make the inner session different for each endpoint. The race condition would not happen, and no packet drop would be seen.

7. Click **OK**.

STEP 10 | Set monitoring options for traffic that matches a tunnel inspection policy rule.

1. Select **Policies > Tunnel Inspection** and select the tunnel inspection policy rule you created.

2. Select **Inspection > Monitor Options**.
3. Enter a **Monitor Name** to group similar traffic together for purposes of logging and reporting.
4. Enter a **Monitor Tag (number)** to group similar traffic together for logging and reporting (range is 1 to 16,777,215). The tag number is globally defined.



This field does not apply to the VXLAN protocol. VXLAN logs automatically use the VNI ID from the VXLAN header.



If you tag tunnel traffic, you can later filter on the Monitor Tag in the tunnel inspection log and use the ACC to view tunnel activity based on Monitor Tag.

5. **Override Security Rule Log Setting** to enable logging and log forwarding options for sessions that meet the selected tunnel inspection policy rule. If you don't select this setting, tunnel log generation and log forwarding are determined by the log settings for the Security policy rule that applies to the tunnel traffic. You can 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, VXLAN, or GTP-U) sessions and the traffic logs store the inner traffic flows.
 6. Select **Log at Session Start** to log traffic at the start of a session.
-  *The best practice for Tunnel logs is to log both at session start and session end because tunnels can stay up for long periods of time. For example, GRE tunnels can come up when the router boots and never terminate until the router is rebooted. If you don't log at session start, you will never see in the ACC that there is an active GRE tunnel.*
7. Select **Log at Session End** to log traffic at the end of a session.
 8. Select a **Log Forwarding** profile that determines where the firewall forwards tunnel logs for sessions that meet the tunnel inspection rule. Alternatively, you can create a new Log Forwarding profile if you [Configure Log Forwarding](#).
 9. Click **OK**.

STEP 11 | (Optional, VXLAN Only) Configure a VXLAN ID (VNI). By default, all VXLAN network interfaces (VNIs) are inspected. If you configure one or more VXLAN IDs, the policy inspects only those VNIs.



Only the VXLAN protocol uses the Tunnel ID tab to specify the VNI.

1. Select the **Tunnel ID** tab and click **Add**.
2. Assign a **Name**. The name is a convenience, and is not a factor in logging, monitoring, or reporting.
3. In the **VXLAN ID (VNI)** field, enter a single VNI, a comma-separated list of VNIs, a range of VNIs (with a hyphen as the separator), or a combination of these. For example, you can specify:

1677002, 1677003, 1677011–1677038, 1024

STEP 12 | (Optional) If you enabled **Rematch Sessions (Device > Setup > Session)**, ensure the firewall doesn't drop existing sessions when you create or revise a tunnel inspection policy by disabling **Reject Non-SYN TCP** for the zones that control your tunnel Security policy rules.

The firewall displays the following warning when you:

- Create a tunnel inspection policy rule.
- Edit a tunnel inspection policy rule by adding a **Protocol** or by increasing the **Maximum Tunnel Inspection Levels** from **One Level** to **Two Levels**.

- **Enable Security Options** in the **Security Options** tab by either adding new zones or changing one zone to another zone.



*Warning: Enabling tunnel inspection policies on existing tunnel sessions will cause existing TCP sessions inside the tunnel to be treated as non-syn-tcp flows. To ensure existing sessions are not dropped when the tunnel inspection policy is enabled, set the **Reject Non-SYN TCP** setting for the zone(s) to **no** using a **Zone Protection** profile and apply it to the zones that control the tunnel's security policies. Once the existing sessions have been recognized by the firewall, you can re-enable the **Reject Non-SYN TCP** setting by setting it to **yes** or **global**.*

1. Select **Network > Network Profiles > Zone Protection** and **Add** a profile.
2. Enter a **Name** for the profile.
3. Select **Packet Based Attack Protection > TCP Drop**.
4. For **Reject Non-SYN TCP**, select **no**.
5. Click **OK**.
6. Select **Network > Zones** and select the zone that controls your tunnel Security policy rules.
7. For **Zone Protection Profile**, select the Zone Protection profile you just created.
8. Click **OK**.
9. Repeat the previous three substeps (12.f, 12.g, and 12.h) to apply the Zone Protection profile to additional zones that control your tunnel Security policy rules.
10. After the firewall has recognized the existing sessions, you can re-enable **Reject Non-SYN TCP** by setting it to **yes** or **global**.

STEP 13 | (Optional) Limit fragmentation of traffic in a tunnel.

1. Select **Network > Network Profiles > Zone Protection** and **Add** a profile by **Name**.
2. Enter a **Description**.
3. Select **Packet Based Attack Protection > IP Drop > Fragmented traffic**.
4. Click **OK**.
5. Select **Network > Zones** and select the tunnel zone where you want to limit fragmentation.
6. For **Zone Protection Profile**, select the profile you just created to apply the Zone Protection profile to the tunnel zone.
7. Click **OK**.

STEP 14 | **Commit** your changes.

View Inspected Tunnel Activity

Perform the following task to view activity of inspected tunnels.

STEP 1 | Select **ACC** and select a **Virtual System** or **All** virtual systems.

STEP 2 | Select **Tunnel Activity**.

STEP 3 | Select a Time period to view, such as **Last 24 Hrs** or **Last 30 Days**.

STEP 4 | For **Global Filters**, click the **+** or **-** buttons to use **ACC Filters** on tunnel activity.

STEP 5 | View inspected tunnel activity; you can display and sort data in each window by **bytes**, **sessions**, **threats**, **content**, or **URLs**. Each window displays a different aspect of tunnel data in graph and table format:

- **Tunnel ID Usage**—Each tunnel protocol lists the Tunnel IDs of tunnels using that protocol. Tables provide totals of Bytes, Sessions, Threats, Content, and URLs for the protocol. Hover over the tunnel ID to get a breakdown per tunnel ID.
- **Tunnel Monitor Tag**—Each tunnel protocol lists tunnel monitor tags of tunnels using that tag. Tables provide totals of Bytes, Sessions, Threats, Content, and URLs for the tag and for the protocol. Hover over the tunnel monitor tag to get a breakdown per tag.
- **Tunneled Application Usage**—Application categories graphically display types of applications grouped into media, general interest, collaboration, and networking, and color-coded by their risk. The Application tables also include a count of users per application.
- **Tunneled User Activity**—Displays a graph of bytes sent and bytes received, for example, along an x-axis of date and time. Hover over a point on the graph to view data at that point. The Source User and Destination User table provides data per user.
- **Tunneled Source IP Activity**—Displays graphs and tables of bytes, sessions, and threats, for example, from an Attacker at an IP address. Hover over a point on the graph to view data at that point.
- **Tunneled Destination IP Activity**—Displays graphs and tables based on destination IP addresses. View threats per Victim at an IP address, for example. Hover over a point on the graph to view data at that point.

View Tunnel Information in Logs

You can view Tunnel Inspection logs themselves or view tunnel inspection information in other types of logs.

GRE, Non-Encrypted IPSec, and GTP-U Protocols

- When there is a TCI traffic rule match, GRE, IPSec, and GTP-U protocols are logged in the Tunnel Inspection log with the Tunnel log type, the matched protocol, and the configured Monitor name and Monitor tag (number).
- When there is no TCI rule match, all protocols are logged under Traffic logs.

VXLAN Protocol

- When there is a TCI traffic rule match, VXLAN protocol is logged in the Tunnel Inspection log with the Tunnel (VXLAN) log type, the configured Monitor name, and the Tunnel ID (VNI).

In the Traffic log for the inner session, the Tunnel Inspected flag indicates a VNI session. The Parent Session is the session that was active when the inner session was created so the ID might not match the current Session ID.

- When there is no TCI rule match, VNI sessions are logged in Traffic logs with the UDP protocol, source port 0, and destination port 4789 (the default).
- View Tunnel inspection logs.
 1. Select **Monitor > Logs > Tunnel Inspection** and view the log data to identify the tunnel **Applications** used in your traffic and any concerns, such as high counts for packets failing Strict Checking of headers.
 2. Click the Detailed Log View  to see details about a log.
- View other logs for tunnel inspection information.
 1. Select **Monitor > Logs**.
 2. Select **Traffic, Threat, URL Filtering, WildFire Submissions, Data Filtering, or Unified**.
 3. For a log entry, click the Detailed Log View .
 4. In the Flags window, see if the **Tunnel Inspected** flag is checked. A Tunnel Inspected flag indicates the firewall used a Tunnel Inspection policy rule to inspect the inside content or inner tunnel. Parent

Session information refers to an outer tunnel (relative to an inner tunnel) or an inner tunnel (relative to inside content).

On the **Traffic, Threat, URL Filtering, WildFire Submissions, Data Filtering** logs, only direct parent information appears in the Detailed Log View of the inner session log, no tunnel log information. If you configured two levels of tunnel inspection, you can select the parent session of this direct parent to view the second parent log. (You must monitor the **Tunnel Inspection** log as shown in the prior step to view tunnel log information.)

5. If you are viewing the log for an inside session that is Tunnel Inspected, click the **View Parent Session** link in the General section to see the outside session information.

Create a Custom Report Based on Tagged Tunnel Traffic

You can create a report to gather information based on the tag you applied to tunnel traffic.

STEP 1 | Select **Monitor > Manage Custom Reports** and click **Add**.

STEP 2 | For Database, select the Traffic, Threat, URL, Data Filtering, or WildFire Submissions log.

STEP 3 | For Available Columns, select Flags and Monitor Tag, along with other data you want in the report.

You can also [Generate Custom Reports](#).

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, Authentication, 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
- > Track Rules Within a Rulebase
- > Enforce Policy Rule Description, Tag, and Audit Comment
- > 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
- > Test Policy Rules

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.

For all policy types, when you [Enforce Policy Rule Description, Tag, and Audit Comment](#), you can use the audit comment archive to view how a policy rule changed over time. The archive, which includes the audit comment history and the configuration logs, enables you to compare configuration versions and review who created or modified and why.

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 more 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 .
Authentication	Identify traffic that requires users to authenticate. For more details, see Authentication Policy .
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 .

Security Policy

Security policy protects network assets from threats and disruptions and helps to optimally allocate network resources for enhancing productivity and efficiency in business processes. On a 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, the user, and the service.



To ensure that end users authenticate when they try to access your network resources, the firewall evaluates [Authentication Policy](#) before Security policy.

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 bidirectional traffic in that session (client to server and server to client). 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 traffic (within a zone) and deny all interzone traffic (between zones). Although these rules are part of the predefined 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 and, after a match is triggered, 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 you enable logging 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.

After an administrator configures a rule, you can [View Policy Rule Usage](#) to determine when and how many times traffic matches the Security policy rule to determine its effectiveness. As your rulebase evolves, change and audit information get lost over time unless you archived this information at the time the rule is created or modified. You can [Enforce Policy Rule Description, Tag, and Audit Comment](#) to ensure that all administrators enter audit comments so that you can view the audit comment archive and review comments and configuration log history and can compare rule configuration versions for a selected rule. Together, you now have more visibility into and control over the rulebase.

- [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 table:

Required/ Optional	Field	Description
Required	Name	A label (up to 63 characters) that identifies the rule.
	UUID	The Universally Unique Identifier (UUID) is a distinct 32-character string that permanently identifies rules so that you can track a rule regardless of any changes to it, such as the name.

Required/ Optional	Field	Description
	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 that 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 <i>Allow</i> or <i>Deny</i> action for the traffic based on the criteria you define in the rule. When you configure the firewall to deny traffic, it either resets the connection or silently drops packets. To provide a better user experience, you can configure granular options to deny 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	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 <i>IT-sanctioned applications</i> or <i>High-risk applications</i> .
	Description	A text field, up to 1024 characters, used to describe the rule.
	Source Address	Define host IP addresses, subnets, address objects (of type IP netmask, IP range, FQDN, or IP wildcard mask), address 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).

Required/ Optional	Field	Description
	Destination Address	The location or destination for the packet. Define IP addresses, subnets, address objects (of type IP netmask, IP range, FQDN, or IP wildcard mask), address 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).
	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 Palo Alto Networks firewalls, you must purchase a URL filtering license.  <i>To block or allow traffic based on URL category, you must apply a URL Filtering profile to the security policy rules. Define the URL Category as Any and attach a URL Filtering profile to the security policy. See Set Up a Basic Security Policy for information on using the default profiles in your security policy.</i>
	Service	Allows you to select a Layer 4 (TCP or UDP) port for the application. You can choose <i>any</i>, specify a port, or use <i>application-default</i> to permit use of the standards-based port for the application. For example, for applications with well-known port numbers such as DNS, the <i>application-default</i> 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.  <i>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.</i>
	Security Profiles	Provide additional protection from threats, vulnerabilities, and data leaks. Security profiles are evaluated only for rules that have an <i>allow</i> action.

Required/ Optional	Field	Description
	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)	Allows the traffic.
Deny	Blocks traffic and enforces the default <i>Deny Action</i> 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 <i>communication with the destination is administratively prohibited</i> —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

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 **Add** a new rule.
2. In the **General** tab, enter a descriptive **Name** for the rule.
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**.

STEP 4 | Define the matching criteria for the destination fields in the packet.

1. In the **Destination** tab, set the **Destination Zone**.
2. Specify a **Destination IP Address** or leave the value set to **any**.



As a best practice, use address objects as the Destination Address to enable access to only specific servers or specific groups of servers especially for commonly exploited services, such as DNS and SMTP. 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 that 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** you want to safely enable. You can select multiple applications or you can 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 that the rule allows are allowed only 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 destined for that 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 disable this setting if you don't want any logs generated when traffic matches this rule or you can select **Log at Session Start** for more detailed logging.
- Select a **Log Forwarding** profile.



As a best practice, do not select the check box to Disable Server Response Inspection (DSRI). Selecting this option prevents the firewall from inspecting packets from the server

to the client. For the best security posture, the firewall must inspect both the client-to-server flows and the server-to-client flows to detect and prevent threats.

STEP 9 | Attach security profiles to enable the firewall to scan all allowed traffic for threats.



Make sure you [create best practice security profiles](#) that help 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 | Click **Commit** to save the policy rule to the running configuration on the firewall.

STEP 11 | To verify that you have set up your basic security policies effectively, test whether your security policy rules are being evaluated and determine which security policy rule applies to a traffic flow.

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:

1. Select **Device > Troubleshooting**, and select **Security Policy Match** from the Select Test drop-down.
2. Enter the Source and Destination IP addresses.
3. Enter the Protocol.
4. **Execute** the security policy match test.



STEP 12 | After waiting long enough to allow traffic to pass through the firewall, [View Policy Rule Usage](#) to monitor the policy rule usage status and determine the effectiveness of the policy rule.

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.



If you need to export specific parts of the configuration for internal review or audit, you can [Export Configuration Table Data](#) as a PDF or CSV file.

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), an IP wildcard address (IPv4 address/wildcard mask) 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.  <i>The predefined External Dynamic Lists (EDLs) on the firewall count toward the maximum number of address objects that a firewall model supports.</i>
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/

Policy Object	Description
	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).  <i>To view the standard ports used by an application, in Objects > Applications search for the application and click the link. A succinct description displays.</i>

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, see [Create Best Practice Security Profiles for the Internet Gateway](#) to learn how you can create custom profiles.



For recommendations on the best-practice settings for security profiles, see [Create Best Practice Security Profiles for the Internet Gateway](#).

You can add security profiles that are commonly applied together to [Create 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).

Profile Type	Description
Antivirus Profiles	Antivirus profiles protect against viruses, worms, and trojans as well as spyware downloads. Using a stream-based malware prevention engine, which inspects traffic the moment the first packet is received, the Palo Alto Networks antivirus solution can provide protection for clients without significantly impacting the performance of the firewall. This profile scans for a wide variety of malware in executables, PDF files, HTML and JavaScript viruses, including support for scanning inside compressed files and data encoding schemes. If you have enabled Decryption on the firewall, the profile also enables scanning of decrypted content. The default profile inspects all of the listed protocol decoders for viruses, and generates alerts for SMTP, IMAP, and POP3 protocols while blocking for FTP, HTTP, and SMB protocols. You can configure the action for a decoder or Antivirus signature and specify how the firewall responds to a threat event: • Default—For each threat signature and Antivirus signature that is defined by Palo Alto Networks, a default action is specified internally. Typically, the default action is an alert or a reset-both. The default action is displayed in parenthesis, for example default (alert) in the threat or Antivirus signature. • Allow—Permits the application traffic.  <i>The Allow action does not generate logs related to the signatures or profiles.</i> • Alert—Generates an alert for each application traffic flow. The alert is saved in the threat log. • Drop—Drops the application traffic.

Profile Type	Description
	• 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 low 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  <i>The Allow action does not generate logs related to the signatures or profiles.</i> • 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.

Profile Type	Description
	• 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.  <i>In some cases, when the profile action is set to reset-both, the associated threat log might display the action as reset-server. This occurs when the firewall detects a threat at the beginning of a session and presents the client with a 503 block page. Because the block page disallows the connection, the client-side does not need to be reset and only the server-side connection is reset.</i> • 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. 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  <i>The Allow action does not generate logs related to the signatures or profiles.</i> • 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.

Profile Type	Description
	• Reset Both—For TCP, resets the connection on both client and server ends. For UDP, drops the connection.  <i>In some cases, when the profile action is set to reset-both, the associated threat log might display the action as reset-server. This occurs when the firewall detects a threat at the beginning of a session and presents the client with a 503 block page. Because the block page disallows the connection, the client-side does not need to be reset and only the server-side connection is reset.</i> • Block IP— This action blocks traffic from either a source or a source-destination pair. It is configurable for a specified period of time.
URL Filtering Profiles	URL Filtering profiles enable you to monitor and control how users access the web over HTTP and HTTPS. The firewall comes with a default profile that is configured to block websites such as known malware sites, phishing sites, and adult content sites. You can use the default profile in a security policy, clone it to be used as a starting point for new URL filtering profiles, or add a new URL profile that will have all categories set to allow for visibility into the traffic on your network. You can then customize the newly added URL profiles and add lists of specific websites that should always be blocked or allowed, which provides more granular control over URL categories.
Data Filtering Profiles	Data filtering profiles prevent sensitive information such as credit card or social security numbers from leaving a protected network. The data filtering profile also allows you to filter on key words, such as a sensitive project name or the word confidential. It is important to focus your profile on the desired file types to reduce false positives. For example, you may only want to search Word documents or Excel spreadsheets. You may also only want to scan web-browsing traffic, or FTP. You can create custom data pattern objects and attach them to a Data Filtering profile to define the type of information on which you want to filter. Create data pattern objects based on: • Predefined Patterns—Filter for credit card and social security numbers (with or without dashes) using predefined patterns. • Regular Expressions—Filter for a string of characters. • File Properties—Filter for file properties and values based on file type.  <i>If you're using a third-party, endpoint data loss prevention (DLP) solutions to populate file properties to indicate sensitive content, this option enables the firewall to enforce your DLP policy.</i> To get started, Set Up Data Filtering.
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

Profile Type	Description
	when a user attempts to download the specified file type. This allows the user to take a moment to consider whether or not they want to download a file. You can define your own custom File Blocking profiles, or choose one of the following predefined profiles when applying file blocking to a Security policy rule. The predefined profiles, which are available with content release version 653 and later, allow you to quickly enable best practice file blocking settings: • basic file blocking—Attach this profile to the Security policy rules that allow traffic to and from less sensitive applications to block files that are commonly included in malware attack campaigns or that have no real use case for upload/download. This profile blocks upload and download of PE files (.scr, .cpl, .dll, .ocx, .pif, .exe) , Java files (.class, .jar), Help files (.chm, .hlp) and other potentially malicious file types, including .vbe, .hta, .wsf, .torrent, .7z, .rar, .bat. Additionally, it prompts users to acknowledge when they attempt to download encrypted-rar or encrypted-zip files. This rule alerts on all other file types to give you complete visibility into all file types coming in and out of your network. • strict file blocking—Use this stricter profile on the Security policy rules that allow access to your most sensitive applications. This profile blocks the same file types as the other profile, and additionally blocks flash, .tar, multi-level encoding, .cab, .msi, encrypted-rar, and encrypted-zip files. 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. To get started, Set Up File Blocking.
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. If a profile rule is set to forward files to the WildFire public cloud, the firewall also forwards files that match existing antivirus signatures, in addition to unknown files. 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,

Profile Type	Description
	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 prevent session exhaustion attacks. In this type of attack, a large number of hosts (bots) are used to establish as many fully established sessions as possible to consume all of a system's resources. You can enable both types of protection mechanisms in a single DoS protection profile. The DoS profile is used to specify the type of action to take and details on matching criteria for the DoS policy. The DoS profile defines settings for SYN, UDP, and ICMP floods, can enable resource protect and defines the maximum number of concurrent connections. After you configure the DoS protection profile, you then attach it to a DoS policy. When configuring DoS protection, it is important to analyze your environment in order to set the correct thresholds and due to some of the complexities of defining DoS protection policies, this guide will not go into detailed examples.
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.
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. See Create a Security Profile Group and Set Up or Override a Default Security Profile Group.

Profile Type	Description
	 For recommendations on the best-practice settings for security profiles, see Create Best Practice Security Profiles for the Internet Gateway .

Create a Security Profile Group

Use the following steps to create a security profile group and add it to a security policy.

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**.

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.*

- 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 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:
 1. Select **Policies > Security** and **Add** a new security policy.
 2. 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**).

Track Rules Within a Rulebase

To keep track of rules within a rulebase, you can refer to the *rule number*, which changes depending on the order of a rule in the rulebase. The rule number determines the order in which the firewall applies the rule.

The *universally unique identifier (UUID)* for a rule never changes even if you modify the rule, such as when you change the rule name. The UUID allows you to track the rule across rulebases even after you deleted the rule.


Rule Numbers

The firewall automatically numbers each rule within a rulebase; when you move or reorder rules, the numbers change based on the new order. When you filter the list of rules to find rules that match specific criteria, the firewall displays each rule with its number in the context of the complete set of rules in the rulebase and its place in the evaluation order.

Panorama independently numbers pre-rules, post-rules, and default rules. 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. You can **Preview Rules** in Panorama to display an ordered list of the total number of rules on a firewall.

- 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.



Rule Number	Name	Type	Action	...
1	Default Rule	Deny	Deny	...
2	...	...	...	...
3	...	...	...	...
4	...	...	...	...
5	...	...	...	...
6	...	...	...	...
7	...	...	...	...
8	...	...	...	...
9	...	...	...	...
10	...	...	...	...
11	...	...	...	...
12	...	...	...	...
13	...	...	...	...
14	...	...	...	...
15	...	...	...	...

- View the numbered list of rules on Panorama.


Select **Policies** and any rulebase under it. For example, **Policies > Security > Pre-rules**.



Rule Number	Name	Type	Action	...
1	...	...	...	...
2	...	...	...	...
3	...	...	...	...
4	...	...	...	...
5	...	...	...	...
6	...	...	...	...
7	...	...	...	...
8	...	...	...	...
9	...	...	...	...
10	...	...	...	...
11	...	...	...	...
12	...	...	...	...
13	...	...	...	...
14	...	...	...	...
15	...	...	...	...

- After you push the rules from Panorama, view the complete list of rules with numbers on the firewall.

From the web interface on 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.



Name	Type	Action	Source	Destination	Service	Rule UUID
1. Default Deny Rule	Deny	Deny	any	any	any	00000000-0000-0000-0000-000000000000
2. Permitted Traffic	Allow	Allow	any	any	any	00000000-0000-0000-0000-000000000000
3. Permitted Traffic	Allow	Allow	any	any	any	00000000-0000-0000-0000-000000000000
4. Permitted Traffic	Allow	Allow	any	any	any	00000000-0000-0000-0000-000000000000
5. Permitted Traffic	Allow	Allow	any	any	any	00000000-0000-0000-0000-000000000000
6. Permitted Traffic	Allow	Allow	any	any	any	00000000-0000-0000-0000-000000000000
7. Permitted Traffic	Allow	Allow	any	any	any	00000000-0000-0000-0000-000000000000
8. Permitted Traffic	Allow	Allow	any	any	any	00000000-0000-0000-0000-000000000000

Rule UUIDs

The universally unique identifier (UUID) for a rule is a 32-character string (based on data such as the network address and the timestamp of creation) that the firewall or Panorama assigns to the rule. The UUID uses the format 8-4-4-4-12 (where 8, 4, and 12 represent the number of unique characters separated by hyphens). UUIDs identify rules for all policy rulebases. You can also use UUIDs to identify applicable rules in the following log types: Traffic, Threat, URL Filtering, WildFire Submission, Data Filtering, GTP, SCTP, Tunnel Inspection, Configuration, and Unified.

Using the UUID to search for a rule enables you to locate a specific rule you want to find among thousands of rules that may have similar or identical names. UUIDs also simplify automation and integration for rules in third-party systems (such as ticketing or orchestration) that do not support names.

In some cases, you may need to generate new UUIDs for existing rulebases. For example, if you want to export a configuration to another firewall, you need to *regenerate the UUIDs* for the rules as you import the configuration to ensure there are no duplicate UUIDs. If you regenerate UUIDs, you are no longer able to track those rules using their previous UUIDs and the hit data and app usage data for those rules are reset.

The firewall or Panorama assigns UUIDs when you:

- Create new rules
- Clone existing rules
- Override the default security rules
- Load a named configuration and regenerate UUIDs
- Load a named configuration containing new rules that are not in the running configuration
- Upgrade the firewall or Panorama to a PAN-OS 9.0 release

When you load a configuration that contains rules with UUIDs, the firewall considers rules to be the same if the rule name, rulebase, and virtual system all match. Panorama considers rules to be the same if the rule name, rulebase, and the device group all match.

Keep in mind the following important points for UUIDs:

- If you manage firewall policy from Panorama, UUIDs are generated on Panorama and therefore must be pushed from Panorama. If you do not push the configuration from Panorama prior to upgrading the firewalls to PAN-OS 9.0, the firewall upgrade will not succeed because it will not have the UUIDs.
- In addition, if you are upgrading an HA pair, upon upgrade to PAN-OS 9.0, each peer independently assigns UUIDs for each policy rule. Because of this, the peers will show as out of sync until you sync the configuration (**Dashboard > Widgets > System > High Availability > Sync to peer**).
- If you remove an existing high availability (HA) configuration after upgrading to PAN-OS 9.0, you must regenerate the UUIDs on one of the peers (**Device > Setup > Operations > Load named configuration snapshot > Regenerate UUIDs for the selected named configuration**) and commit the changes to prevent UUID duplication.
- All rules pushed from Panorama will share the same UUID; all rules local to a firewall will have different UUIDs. If you create a rule locally on the firewall after you push the rules from Panorama to the firewalls, the rule you created locally has its own UUID.
- To replace an RMA Panorama, make sure you **Retain Rule UUIDs** when you load the named Panorama configuration snapshot. If you do not select this option, Panorama removes all previous rule UUIDs from the configuration snapshot and assigns new UUIDs to the rules on Panorama, which means it does not retain information associated with the previous UUIDs, such as the policy rule hit count.
- Display the Rule UUID column for logs and the UUID column for policy rules.


To view the UUIDs, you must display the column, which does not display by default.

- To display the UUID in logs:
 1. Select **Monitor** and then expand the column header (▼).
 2. Select **Columns**.
 3. Enable **Rule UUID**.



- To display UUIDs on the policy rulebase:
 1. Select **Policies** and then expand the column header (▼).
 2. Select **Columns**.
 3. Enable **UUID**.

UUIDs are available for all policy rulebases.



- Copy the UUID for a log or policy rule.

Copying the UUID allows you to paste the UUID in to searches, the ACC, custom reports, filters, and anywhere else you want to locate a rule identified by that UUID.


1. Select the ellipses that display when you move your cursor over the entry in the Rule UUID column.

	Receive Time	Type	Rule UUID
	06/06 15:07:36	end	98abbb07-b04e-47bf-bb02-d02e951ca7f7 ...
	06/06 15:07:26	end	98abbb07-b04e-47bf-bb02-d02e951ca7f7
	06/06 15:07:18	end	98abbb07-b04e-47bf-bb02-d02e951ca7f7
	06/06 15:05:48	end	98abbb07-b04e-47bf-bb02-d02e951ca7f7

2. Copy the UUID from the pop-up.

	Receive Time	Type	Rule UUID
	06/06 15:07:36	end	98abbb07-b04e-47bf-bb02-d02e951ca7f7
	06/06 15:07:26	end	98abbb07-b04e-47bf-bb02-d02e951ca7f7
	06/06 15:07:18	end	98abbb07-b04e-47bf-bb02-d02e951ca7f7
	06/06 15:05:48	end	98abbb07-b04e-47bf-bb02-d02e951ca7f7

You can also go to the **Policies** tab, click the arrow to the right of the rule name, and **Copy UUID**.



- Check the Configuration Logs to view UUIDs for deleted rules.
To view the UUID for a deleted rule, select **Monitor** > **Logs** > **Configuration**.

Enforce Policy Rule Description, Tag, and Audit Comment

When creating or modifying rules, you can require a rule description, tag, and audit comment to ensure your policy rulebase is correctly organized and grouped, and to preserve important rule history for auditing purposes. By requiring a rule description, tag, and audit comment, you can simplify your policy rulebase review by ensuring that rules are appropriately grouped, and that the rule change history is tracked when creating or modifying a rule. For uniformity, you can set specific requirements for what the audit comment can include.

By default, enforcement of a description, tag, and audit comment is not enabled. You can specify whether a description, tag, audit comment, or any combination of these three is required to successfully add or modify a rule. The audit comment archive allows you to view the audit comments entered for a selected rule, review the configuration log history, and compare rule configuration versions.

STEP 1 | [Launch the Web Interface.](#)

STEP 2 | Select **Device > Setup > Management** and edit the Policy Rulebase Settings.

STEP 3 | Configure the settings you want to enforce. In this example, tags and audit comments are required for all policies.



Enforce audit comments for policy rules to capture the reason an administrator creates or modifies a rule. Requiring audit comments on policy rules helps maintain an accurate rule history for auditing purposes.

STEP 4 | Configure the Audit Comment Regular Expression to specify the audit comment format.

When administrators create or modify a rule, you can require they enter a comment those audit comments adhere to a specific format that fits your business and auditing needs by specifying letter and number expressions. For example, you can use this setting to specify regular expressions that match your ticketing number formats:

- **[0-9]{<Number of digits>}**—Requires the audit comment to contain a minimum number of digits that range from 0 to 9. For example, **[0-9]{6}** requires a minimum of six digit in a numerical expression with numbers 0 to 9.
- **<Letter Expression>**—Requires the audit comment to contain a letter expression. For example, **Reason for Change-** requires that the administrator begin the audit comment with this letter expression.
- **<Letter Expression>-[0-9]{<Number of digits>}**—Requires the audit comment to contain a predetermined character followed by a minimum number of digits that range from 0 to 9. For example, **SB-[0-9]{6}** requires the audit comment format to begin with **SB-**, followed by a minimum six digits in a numerical expression with values from 0 to 9. For example, **SB-012345**.
- **(<Letter Expression>)|(<Letter Expression>)|(<Letter Expression>)-[0-9]{<Number of digits>}**—Requires the audit comment to contain a prefix using any one of the predetermined letter expressions with a minimum number of digits that range from 0 to 9. For example, **(SB|XY|PN)-[0-9]{6}** requires the audit comment format to begin with **SB-**, **XY-**, or **PN-** followed by a minimum of six digits in a numerical expression with values from 0 to 9. For example, **SB-012345**, **XY-654321**, or **PN-012543**.

STEP 5 | Click **OK** to apply the new policy rulebase settings.



Policy Rulebase Settings

- ☒ Require Tag on policies
- ☐ Require description on policies
- ☒ Fail commit if policies have no tags or description
- ☒ Require audit comment on policies

Audit Comment Regular Expression:


- ☒ Policy Rule Hit Count
- ☒ Policy Application Usage

OK Cancel

STEP 6 | Commit the changes.



After you commit the policy rulebase settings changes, modify the existing policy rule based on the rulebase settings you decided to enforce.



Commit Status

Operation: Commit

Status: Completed

Result: Successful

Details: Configuration committed successfully

Warnings

- rulebase -> security -> rules -> Network Infrastructure is invalid. Description and Tag missing for rule entry
- rulebase -> security -> rules -> Internet Access is invalid. Description and Tag missing for rule entry
- rulebase -> security -> rules -> Data Center Applications is invalid. Description missing for rule entry
- rulebase -> security -> rules -> test-rule2 is invalid. Description and Tag missing for rule entry
- rulebase -> security -> rules -> test-rule5 is invalid. Description and Tag missing for rule entry
- rulebase -> security -> rules -> test-rule6 is invalid. Description and Tag missing for rule entry
- rulebase -> security -> rules -> test-rule7 is invalid. Description and Tag missing for rule entry
- rulebase -> security -> rules -> deny-worldwidecraft is invalid. Description missing for rule entry
- rulebase -> security -> rules -> test-rule4 is invalid. Description missing for rule entry

Warning: No Valid Threat License

Warning: Rule 'Network Infrastructure': No valid URL filtering license.


Warning: Rulebase 'security'

Rule 'Internet Access': No zone has 'enable-user-identification' turned on

Close

STEP 7 | Verify that the firewall is enforcing the new policy rulebase settings.

1. Select **Policies** and **Add** a new rule.
2. Confirm that you must add a tag and enter an audit comment click **OK**.



Security Policy Rule

General Source User Destination Application Service/URL Category Actions

Name:

Rule Type: universal (default)

Description:

Tags:

Group Rules By Tag: None

Audit Comment:

Audit Comment Action:

OK Cancel

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.



When cloning multiple policy rules, the order by which you select the rules will determine the order they are copied to the device group. For example, if you have rules 1-4 and your selection order is 2-1-4-3, the device group where these rules will be cloned will display the rules in the same order you selected. However, you can reorganize the rules as you see fit once they have been successfully copied.

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**.

Use an Address Object to Represent IP Addresses

Create an address object on the firewall to group IP addresses or to specify an FQDN, and then reference the address object in a firewall policy rule, filter, or other function to avoid having to individually specify multiple IP addresses in the rule, filter, or other function.

Furthermore, you can reference the same address object in multiple policy rules, filters, or other functions without needing to specify the same individual addresses in each use. For example, you can create an address object that specifies an IPv4 address range and then reference the address object in a Security policy rule, a NAT policy rule, and a custom report log filter.

- [Address Objects](#)
- [Create an Address Object](#)

Address Objects

An address object is a set of IP addresses that you can manage in one place and then use in multiple firewall policy rules, filters, and other functions. There are four types of address objects: **IP Netmask**, **IP Range**, **IP Wildcard Mask**, and **FQDN**.

An address object of type **IP Netmask**, **IP Range**, or **FQDN** can specify IPv4 or IPv6 addresses. An address object of type **IP Wildcard Mask** can specify only IPv4 addresses.

An address object of type **IP Netmask** requires you to enter the IP address or network using slash notation to indicate the IPv4 network or the IPv6 prefix length. For example, 192.168.18.0/24 or 2001:db8:123:1::/64.

An address object of type **IP Range** requires you to enter the IPv4 or IPv6 range of addresses separated by a hyphen.

An address object of type **FQDN** (for example, paloaltonetworks.com) provides further ease of use because DNS provides the FQDN resolution to the IP addresses instead of you needing to know the IP addresses and manually updating them every time the FQDN resolves to new IP addresses.

An address object of type **IP Wildcard Mask** is useful if you define private IPv4 addresses to internal devices and your addressing structure assigns meaning to certain bits in the address. For example, the IP address of cash register 156 in the northeastern U.S. could be 10.132.1.156 based on these bit assignments:



An address object of type **IP Wildcard Mask** specifies which source or destination addresses are subject to a Security policy rule. For example, 10.132.1.1/0.0.2.255. A zero (0) bit in the mask indicates that the bit being compared must match the bit in the IP address that is covered by the zero. A one (1) bit in the mask (a wildcard bit) indicates that the bit being compared need not match the bit in the IP address. The following snippets of an IP address and wildcard mask illustrate how they yield four matches:



After you [Create an Address Object](#):

- You can reference an address object of type **IP Netmask**, **IP Range**, or **FQDN** in a policy rule for Security, Authentication, NAT, NAT64, Decryption, DoS Protection, Policy-Based Forwarding (PBF), QoS, Application Override, or Tunnel Inspection; or in a NAT address pool, VPN tunnel, path monitoring, External Dynamic List, Reconnaissance Protection, ACC global filter, log filter, or custom report log filter.
- You can reference an address object of type **IP Wildcard Mask** only in a Security policy rule.

Create an Address Object

Create [Address Objects](#) to represent one or more IP addresses and then reference the address objects in one or more policy rules, filters, or other firewall functions. If you want to change the set of addresses, you change an address object once rather than change multiple policy rules or filters, which reduces your operational overhead.

STEP 1 | Create an address object.

1. Select **Objects > Addresses** and **Add** an address object by **Name**. The name is case-sensitive, must be unique, and can be up to 63 characters (letters, numbers, spaces, hyphens, and underscores).
2. Select the **Type** of address object:
 - **IP Netmask**—Specify a single IPv4 or IPv6 address, an IPv4 network with slash notation, or an IPv6 address and prefix. For example, 192.168.80.0/24 or 2001:db8:123:1::/64. Optionally, click **Resolve** to see the associated FQDN (based on the DNS configuration of the firewall or Panorama). To change the address object type from **IP Netmask** to **FQDN**, select the FQDN and click **Use this FQDN**. The **Type** changes to **FQDN** and the FQDN you select appears in the text field.
 - **IP Range**—Specify a range of IPv4 addresses or IPv6 addresses separated by a hyphen. For example, 192.168.40.1-192.168.40.255 or 2001:db8:123:1::1-2001:db8:123:1::22.
 - **IP Wildcard Mask**—Specify an IP wildcard address (IPv4 address followed by a slash and a mask, which must begin with a 0). For example, 10.5.1.1/0.127.248.2. A zero (0) in the mask indicates the bit being compared must match the bit in the IP address that is covered by the zero. A one (1) in the mask (wildcard bit) indicates the bit being compared need not match the bit in the IP address covered by the one.
 - **FQDN**—Specify the domain name. The FQDN initially resolves at commit time. The firewall subsequently refreshes the FQDN based on the time-to-live (TTL) of the FQDN in DNS, as long as the TTL is greater than or equal to the **Minimum FQDN Refresh Time** you configure (or the default of 30 seconds). The FQDN is resolved by the system DNS server or a DNS proxy object, if a proxy is configured. Click **Resolve** to see the associated IP address (based on the DNS configuration of the firewall or Panorama). To change the address object type from FQDN to IP Netmask, select an IP Netmask and click **Use this address**. The **Type** changes to **IP Netmask** and the IP address you select appears in the text field.
3. (Optional) Enter one or more [Use Tags to Group and Visually Distinguish Objects](#) to apply to the address object.
4. Click **OK**.

STEP 2 | Commit your changes.

STEP 3 | View logs filtered by address object, address group, or wildcard address.

1. For example, select **Monitor > Logs > Traffic** to view traffic logs.
2. Select **+** to add a log filter.
3. Select the **Address** attribute, the **in** Operator, and enter the name of the address object for which you want to view logs. Alternatively, enter an address group name or a wildcard address, such as 10.155.3.4/0.0.240.255.
4. Click **Apply**.

STEP 4 | View a custom report based on an address object.

1. Select **Monitor > Manage Custom Reports** and select a report that uses a Database such as Traffic Log.
2. Select **Filter Builder**.
3. Select an Attribute such as **Address**, **Destination Address** or **Source Address**, select an Operator, and enter the name of the address object for which you want to view the report.

STEP 5 | Use a filter in the ACC to view network activity based on a source IP address or destination IP address that uses an address object.

1. Select **ACC > Network Activity**.
2. View the **Source IP Activity—For Global Filters**, click **+** to add a filter and select one of the following: **Address** or **Source > Source Address** or **Destination > Destination Address** and select an address object.
3. View the **Destination IP Activity—For Global Filters**, click the **+** to add a filter and select one of the following: **Address** or **Source > Source Address** or **Destination > Destination Address** and select an address object.

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](#)
- [View Rules by Tag Group](#)

Create and Apply Tags

Use tags to identify the purpose of a rule or configuration object and to help you better organize your rulebase. To ensure that policy rules are properly tagged, see how to [Enforce Policy Rule Description, Tag, and Audit Comment](#). Additionally, you can [View Rules by Tag Group](#) by first creating and then setting the tag as the Group tag.

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 make the tag available.
3. **Add** a tag and enter a **Name** to identify the tag or select a zone **Name** 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 a **Color** from the 17 predefined colors. By default, **Color** is **None**.



6. Click **OK** and **Commit** to save your changes.

STEP 2 | Apply tags to policy.

1. Select **Policies** and any rulebase under it.
2. **Add** a policy rule and use the tagged objects you created in Step 1.
3. Verify that the tags are in use.



Name	Tag	Type	Zone	Address	User	Zone	Address
General Business Rules	Business Rules	Universal	All Users	any	any	All Internet	any

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 (**Tags**) 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

- Select **Objects > Tags** to perform any of the following operations with tags:
 - Click the **Name** to edit the properties of a tag.
 - Select a tag in the table and **Delete** the tag from the firewall.
 - **Clone** a tag to duplicate it 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 [View Rules by Tag Group](#).

View Rules by Tag Group

View your policy rulebase as tag groups to visually group rules based on the tagging structure you created. In this view, you can perform operational procedures such as adding, deleting, and moving the rules in the selected tag group more easily. Viewing the rulebase as tag groups maintains the rule evaluation order and a single tag may appear multiple times throughout the rulebase to visually preserve the rule hierarchy.

You must create the tag before you can assign it as a group tag on a rule. Policy rules that are already tagged on upgrade to PAN-OS 9.0 have the first tag automatically assigned as the Group tag. Before you upgrade to PAN-OS 9.0, review the tagged rules in your rulebase to ensure rules are correctly grouped. You must manually edit each tag rule and configure the correct Group tag if your rules are grouped incorrectly after you upgrade to PAN-OS 9.0.




Name	Tags	Zone	Address	User	Zone	Address
GroupTag (1)	GroupTag	any	any	any	any	any
GroupTag2 (1)						
GroupTag2 (1)						
GroupTag (1)						
GroupTag (1)						

STEP 1 | Launch the Web Interface.

STEP 2 | [Create and Apply Tags](#) you want to use for grouping rules.

STEP 3 | Assign a policy rule to a tag group.

1. Create a policy rule. Refer to [Policy](#) for more information on creating policy rules.
2. In the **Group Rules by Tag** field, select the tag from the drop-down and click **OK**.



3. **Commit** your changes.

STEP 4 | View your policy rulebase as groups.

1. (**Panorama only**) From the **Device Group**, select the device group rulebase to view or view all Shared rules.
2. Click **Policies** and select the rulebase where you created the rules in Step 2.
3. Select the **View Rulebase as Groups** option (at the bottom).




Rules not assigned a tag group display as None.



STEP 5 | Perform Group operations as needed.

1. Click **Group** to perform group operations for rules in the selected tag group.
 - (**Panorama only**) **Move rules in group to a different rulebase or device group**—Move all policy rules in the selected tag group to the Pre-Rulebase or Post-Rulebase or move them to a different device group.
 - **Change group of all rules**—Move all rules in the selected tag group to a different tag group.
 - **Move all rules in group**—Move all rules in the selected tag group to change the rule priority order.
 - **Delete all rules in group**—Delete all rules in the selected tag group.

- **Clone all rules in group**—Clone all rules in the selected tag group.



2. **Commit your changes.**

Use an External Dynamic List in Policy

An external dynamic list (formerly called dynamic block list) is a text file that you or another source hosts 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 the list is updated, 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](#)
- [Built-in External Dynamic Lists](#)
- [Configure the Firewall to Access an External Dynamic List](#)
- [Retrieve an External Dynamic List from the Web Server](#)
- [View External Dynamic List Entries](#)
- [Exclude Entries from an External Dynamic List](#)
- [Enforce Policy on an External Dynamic List](#)
- [Find External Dynamic Lists That Failed Authentication](#)
- [Disable Authentication for an External Dynamic List](#)

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. When multiple lists are referenced, you can prioritize the order of evaluation to make sure the most important EDLs are committed before capacity limits are reached. 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 uses the last successfully retrieved list for enforcing policy until the connection is restored with the web server. In cases where authentication to the EDL fails, the security policy stops enforcing the EDL. To retrieve the external dynamic list, the firewall uses the interface configured with the **Palo Alto Networks Services** service route.

The firewall supports four types of external dynamic lists:

- **Predefined IP Address**—A predefined IP address list is a type of IP address list that refers to the built-in, dynamic IP lists with fixed or “predefined” contents. These [Built-in External Dynamic Lists](#)—for bulletproof hosting providers, known malicious, and high-risk IP addresses—are automatically added to your firewall if you have an active Threat Prevention license. A predefined IP address list can also refer to an EDL that uses one of the built-in lists as a source. Because you can’t modify the contents of a predefined list, you might use one as a source for a different EDL if you want to add or exclude list entries.
- **IP Address**—The firewall typically enforces policy for a source or destination IP address that is defined as a static object on the firewall (see [Enforce Policy on an External Dynamic List](#)). 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 (see [Use an External Dynamic List in a URL Filtering Profile](#)).
- **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 feeds 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>**. You can also specify the firewall to include the subdomains of a specified domain. For example, if your domain list includes paloaltonetworks.com, all lower level components of the domain name (e.g., *.paloaltonetworks.com) will also be included as part of the list. When this setting is enabled, each domain in a given list requires an additional entry, effectively doubling the number of entries used by the list. For details on configuring domain lists, see [Configure DNS Sinkholing for a List of Custom Domains](#).

On each firewall model, you can add a maximum of 30 custom EDLs with unique sources to enforce policy. The external dynamic list limit is 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-5200 Series and the PA-7000 Series firewalls support a maximum of 150,000 total IP addresses; all other models 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. The IP addresses in predefined IP address lists do not count toward the limit.
- **URL and domain**—The maximum number of URLs and domains supported varies by model. No limits are enforced for the number of URL or domain entries per list. Refer to the following table for specifics on your model:

Model	URL List Entry Limits	Domain List Entry Limits
PA-5200 Series, PA-7000 Series (upgraded with the PA-7000 20GXM NPC, PA-7000 20GQXM NPC, or the PA-7000 100G NPC).  <i>PA-7000 appliances with mixed NPCs only support the standard capacities.</i>	250,000	4,000,000

Model	URL List Entry Limits	Domain List Entry Limits
VM-500, VM-700	100,000	2,000,000
PA-850, PA-820, PA-3200 Series	100,000	1,000,000
PA-7000 Series (and appliances upgraded with the PA-7000 20GQ NPC or the PA-7000 20G NPC), VM-300	100,000	500,000
PA-220, VM-50, VM-50 (Lite), VM-100, VM-1000-HV	50,000	50,000

List entries only count toward the firewall limits if they belong to an external dynamic list that is referenced in policy.



- *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 model. To ensure that the entries do not exceed the limit, check the number of entries currently used in policy. Select Objects > External Dynamic Lists and click List Capacities.*
- *An external dynamic list should not be empty. If you want to stop using the list, remove the reference from the policy rule or profile rather than leave the list empty. If the list is empty the firewall fails to refresh the list and continues to use the last information it retrieved.*
- *As a best practice, Palo Alto Networks recommends using shared EDLs when multiple virtual systems are used. Using individual EDLs with duplicate entries for each virtual system uses more memory, which might over-utilize firewall resources.*
- *EDL entry counts on firewalls operating multi-virtual systems take additional factors into account (such as DAGs, number of vsys, rules bases) to generate a more accurate capacity consumption listing. This might result in a discrepancy in capacity usage after upgrading from PAN-OS 8.x releases.*
- *Depending on the features enabled on the firewall, memory usage limits might be exceeded before EDL capacity limits are met due to memory allocation updates. As a best practice, Palo Alto Networks recommends reviewing EDL capacities and, when necessary, removing or consolidating EDLs into shared lists to minimize memory usage.*

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. The entries in a predefined IP address list comply with the formatting guidelines for IP address lists.

- [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
```



For an IP address that is blocked, you can display a notification page only if the protocol is HTTP.

Domain List

You can use placeholder characters in domain lists to configure a single entry to match against multiple website subdomains, pages, including entire top-level domains, as well as matches to specific web pages.

Follow these guidelines when creating domain list entries:

- 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://.
- You can use an asterisk (*) to indicate a wildcard value.
- You can use a caret (^) to indicate an exact match value.
- The following characters are considered token separators: . / ? & = ; +

Every string separated by one or two of these characters is a token. Use wildcard characters as token placeholders, indicating that a specific token can contain any value.

- Wildcard characters must be the only character within a token; however, an entry can contain multiple wildcards.
- Each domain entry can be up to 255 characters in length.

When to use the asterisk (*) wildcard:

Use an asterisk (*) wildcard to indicate one or multiple variable subdomains. For example, to specify enforcement for Palo Alto Network's website regardless of the domain extension used, which might be one or two subdomains depending on location, you would add the entry: *.paloaltonetworks.com. This entry would match to both docs.paloaltonetworks.com and support.paloaltonetworks.com.

You can also use this wildcard to indicate entire top-level domains. For example, to specify enforcement of a TLD named .work, you would add the entry *.work. This matches all websites ending with .work.



The (*) wildcard can only be prepended in domain entries.

Asterisk (*) examples

EDL Domain List Entry	Matching Sites
*.company.com	eng.tools.company.com

EDL Domain List Entry	Matching Sites
	support.tools.company.com tools.company.com docs.company.com
*.click	all websites ending with a top-level domain of .click.

When to use a caret (^) character:

Use carets (^) to indicate an exact match of a subdomain. For example, **^paloaltonetworks.com** matches only paloaltonetworks.com. This entry does not match to any other site.

Caret (^) examples

EDL Domain List Entry	Matching Site
^company.com	company.com
^eng.company.com	eng.company.com

URL List

See [URL Category Exceptions](#).

Built-in External Dynamic Lists

With an active Threat Prevention license, Palo Alto Networks provides built-in IP address EDLs that you can use to protect against malicious hosts.

- **Palo Alto Networks Bulletproof IP Addresses**—Contains IP addresses provided by bulletproof hosting providers. Because bulletproof hosting providers place few, if any, restrictions on content, attackers frequently use these services to host and distribute malicious, illegal, and unethical material.
- **Palo Alto Networks High-Risk IP Addresses**—Contains malicious IP addresses from threat advisories issued by trusted third-party organizations. Palo Alto Networks compiles the list of threat advisories, but does not have direct evidence of the maliciousness of the IP addresses.
- **Palo Alto Networks Known Malicious IP Addresses**—Contains IP addresses that are verified malicious based on WildFire analysis, Unit 42 research, and data gathered from telemetry ([Share Threat Intelligence with Palo Alto Networks](#)). Attackers use these IP addresses almost exclusively to distribute malware, initiate command-and-control activity, and launch attacks.

The firewall receives updates for these feeds in content updates, allowing the firewall to automatically enforce policy based on the latest threat intelligence from Palo Alto Networks. You cannot modify the contents of the built-in lists. Use them as-is (see [Enforce Policy on an External Dynamic List](#)), or create a custom external dynamic list that uses one of the lists as a source (see [Configure the Firewall to Access an External Dynamic List](#)) and [exclude entries](#) from the list as needed.



Configure the Firewall to Access an External Dynamic List

You must establish the connection between the firewall and the source that hosts the external dynamic list before you can [Enforce Policy on an External Dynamic List](#).

STEP 1 | (Optional) Customize the service route that the firewall uses to retrieve external dynamic lists.

Select **Device > Setup > Services > Service Route Configuration > Customize** and modify the **External Dynamic Lists** service route.



*The firewall does not use the External Dynamic Lists service route to retrieve **Built-in External Dynamic Lists**; content updates modify or update the contents of those lists (active Threat Prevention license required).*

STEP 2 | Find an external dynamic list to use with the firewall.

- Create an external dynamic list and host it on a web server. Enter IP addresses, domains, or URLs in a blank text file. Each list entry 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

See the [Formatting Guidelines for an External Dynamic List](#) to ensure that the firewall does not skip list entries. To prevent commit errors and invalid entries, do not prefix `http://` or `https://` to any of the entries.

- Use an external dynamic list hosted by another source and verify that it follows the [Formatting Guidelines for an External Dynamic List](#).

STEP 3 | Select **Objects > External Dynamic Lists**.

STEP 4 | Click **Add** and enter a descriptive **Name** for the list.

STEP 5 | (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.



As a best practice, Palo Alto Networks recommends using shared EDLs when multiple virtual systems are used. Using individual EDLs with duplicate entries for each vsys uses more memory, which might over-utilize firewall resources.

STEP 6 | (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.

STEP 7 | 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](#).

If you are using a Domain List, you can optionally enable **Automatically expand to include subdomains** to also include the subdomains of a specified domain. For example, if your domain list includes paloaltonetworks.com, all lower level components of the domain name (e.g., *.paloaltonetworks.com) will also be included as part of the list. Keep in mind, when this setting is enabled, each domain in a given list requires an additional entry, effectively doubling the number of entries that are consumed.

STEP 8 | 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**.

If you are creating a list of type Predefined IP, select a Palo Alto Networks malicious IP address feed to use as a source.

STEP 9 | If the list source is secured with SSL (i.e. lists with an HTTPS URL), enable server authentication. Select a **Certificate Profile** or create a **New Certificate Profile** for authenticating the server that hosts the list. The certificate profile you select must have root CA (certificate authority) and intermediate CA certificates that match the certificates installed on the server you are authenticating.



Maximize the number of external dynamic lists that you can use to enforce policy. Use the same certificate profile to authenticate external dynamic lists from the same source URL. If you assign different certificate profiles to external dynamic lists from the same source URL, the firewall counts each list as a unique external dynamic list.

STEP 10 | Enable client authentication if the list source has an HTTPS URL and requires basic HTTP authentication for list access.

1. Select **Client Authentication**.
2. Enter a valid **Username** to access the list.
3. Enter the **Password** and **Confirm Password**.

STEP 11 | (Not available on Panorama) Click **Test Source URL** to verify that the firewall can connect to the web server.



The Test Source URL function is not available when authentication is used for EDL access.

STEP 12 | (Optional) Specify the **Check for updates** 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](#).

STEP 13 | Click **OK** and **Commit**.

STEP 14 | (Optional) EDLs are shown top to bottom, in order of evaluation. Use the directional controls at the bottom of the page to change the list order. This allows you to or order the lists to make sure the most important EDLs are committed before capacity limits are reached.



You can only change the EDL order when Group By Type is deselected.

STEP 15 | **Enforce Policy on an External Dynamic List.**



If the server or client authentication fails, the firewall ceases to enforce policy based on the last successfully retrieved external dynamic list. [Find External Dynamic Lists That Failed Authentication](#) and view the reasons for authentication failure.

Retrieve an External Dynamic List from the Web Server

When you [Configure the Firewall to Access an External Dynamic List](#), you can configure the firewall to retrieve the list from the web server on an hourly (default) five minute, daily, weekly, or monthly basis. If you

have added or deleted IP addresses from the list and need to trigger an immediate refresh, use the following process to fetch the updated 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 is queued.

STEP 3 | To view the status of the job in the Task Manager, see [Manage and Monitor Administrative Tasks](#).

STEP 4 | (Optional) After the firewall retrieves the list, [View External Dynamic List Entries](#).


View External Dynamic List Entries

Before you [Enforce Policy on an External Dynamic List](#), you can view the contents of an external dynamic list directly on the firewall to check if it contains certain IP addresses, domains, or URLs. The entries displayed are based on the version of the external dynamic list that the firewall most recently retrieved.

STEP 1 | Select **Objects > External Dynamic Lists**.

STEP 2 | Click the external dynamic list you want to view.

STEP 3 | Click **List Entries and Exceptions** and view the objects that the firewall retrieved from the list.



The list might be empty if:

- The firewall has not yet retrieved the external dynamic list. To force the firewall to retrieve an external dynamic list immediately, [Retrieve an External Dynamic List from the Web Server](#).
- The firewall is unable to access the server that hosts the external dynamic list. Click **Test Source URL** to verify that the firewall can connect to the server.

STEP 4 | Enter an IP address, domain, or URL (depending on the type of list) in the filter field and Apply Filter () to check if it's in the list. [Exclude Entries from an External Dynamic List](#) based on which IP addresses, domains, and URLs you need to block or allow.

STEP 5 | (Optional) View the [AutoFocus Intelligence Summary](#) for a list entry. Hover over an entry to open the drop-down and then click **AutoFocus**.

Exclude Entries from an External Dynamic List

As you view the entries of an external dynamic list, you can exclude up to 100 entries from the list. The ability to exclude entries from an external dynamic list gives you the option to enforce policy on some (but not all) of the entries in a list. This is helpful if you cannot edit the contents of an external dynamic list (such as the Palo Alto Networks High-Risk IP Addresses feed) because it comes from a third-party source.

STEP 1 | [View External Dynamic List Entries](#).

STEP 2 | Select up to 100 entries to exclude from the list and click Submit () or manually **Add** a list exception.

- You cannot save your changes to the external dynamic list if you have duplicate entries in the Manual Exceptions list. To identify duplicate entries, look for entries with a red underline.
- A manual exception must match a list entry exactly. For example, if an IP address range is included as a list entry and you manually enter a single IP address within the range as a list exception, the firewall will continue to enforce policy on all the IP addresses in the range. So, to exclude that single IP address, you must first make sure that it's a standalone external dynamic list entry, and then manually add the same IP address to the list of exceptions.

STEP 3 | Click **OK** and **Commit** to save your changes.

STEP 4 | (Optional) [Enforce Policy on an External Dynamic List](#).

Enforce Policy on an External Dynamic List

Block or allow traffic based on IP addresses or URLs in an external dynamic list, or use an dynamic domain list with a DNS sinkhole to prevent access to malicious domains. Refer to the table below for the ways you can use external dynamic lists to enforce policy on the firewall.

- [Configure DNS Sinkholing for a List of Custom Domains](#).
- [Use an External Dynamic List in a URL Filtering Profile](#).
- Use an External Dynamic List of Type URL as Match Criteria in a Security Policy Rule.
 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.

1. [View External Dynamic List Entries](#) for the URL list, and attempt to access a URL from the list.
2. Verify that the action you defined is enforced.
3. To monitor the activity on the firewall:

- Select **ACC** and add a URL Domain as a global filter to view the Network Activity and Blocked Activity for the URL you accessed.
- Select **Monitor > Logs > URL Filtering** to access the detailed log view.

- Use an External Dynamic List of Type IP or Predefined 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**.
3. In the **Source/Destination** tabs, set the external dynamic list to be used as the **Source/Destination Address(es)**.
4. In the Service/ URL Category tab, make sure the **Service** is set to **application-default**.
5. 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.

6. Leave all the other options at the default values.
7. Click **OK** to save the changes.
8. **Commit** the changes.
9. Test that the policy action is enforced.
 1. [View External Dynamic List Entries](#) for the external dynamic list, and attempt to access an IP address from the list.
 2. Verify that the action you defined is enforced.
 3. Select **Monitor > Logs > Traffic** and view the log entry for the session.
 4. To verify the policy rule that matches a flow, select **Device > Troubleshooting**, and execute a Security Policy Match test:



Tips for enforcing policy on the firewall with external dynamic lists:

- When viewing external dynamic lists on the firewall (Objects > External Dynamic Lists), click List Capacities to compare how many IP addresses, domains, and URLs are currently used in policy with the total number of entries that the firewall supports for each list type.
- Use Global Find to Search the Firewall or Panorama Management Server for a domain, IP address, or URL that belongs to one or more external dynamic lists is used in policy. This is useful for determining which external dynamic list (referenced in a Security policy rule) is causing the firewall to block or allow a certain domain, IP address, or URL.
- Use the directional controls at the bottom of the page to change the evaluation order of EDLs. This allows you to order the lists to make sure the most important entries in an EDL are committed before capacity limits are reached.



You can only change the EDL order when Group By Type is deselected.

Find External Dynamic Lists That Failed Authentication

When an external dynamic list that requires SSL fails client or server authentication, the firewall generates a system log of critical severity. The log is critical because the firewall ceases to enforce policy based on the external dynamic list after it fails authentication. Use the following process to view critical system log messages notifying you of authentication failure related to external dynamic lists.

STEP 1 | Select **Monitor > Logs > System**.

STEP 2 | Construct the following filters to view all messages related to authentication failure, and apply the filters. For more information, review the complete workflow to [Filter Logs](#).

- Server authentication failure—(**eventid eq tls-edl-auth-failure**)

- Client authentication failure—(**eventid eq edl-cli-auth-failure**)

Generate Time	Type	Severity	Event	Object	Description	Generate Time	Log Type
11/09 22:08:21	sys	critical	edl-auth-failure		EDL server certificate authentication failed. The associated external dynamic list has been removed, which might impact your policy. EDL Name: [redacted] EDL Source URL: [redacted] Reason: self signed certificate in certificate chain	11/09 22:08:21	
11/09 22:08:22	auth	critical	edl-cli-auth-failure		EDL client basic authentication failed. The associated external dynamic list has been removed, which might impact your policy. EDL Name: [redacted] EDL Source URL: [redacted]	11/09 22:08:22	

STEP 3 | Review the system log messages. The message description includes the name of the external dynamic list, the source URL for the list, and the reason for the authentication failure.

The server that hosts the external dynamic list fails authentication if the certificate is expired. If you have configured the certificate profile to check certificate revocation status via Certificate Revocation List (CRL) or Online Certificate Status Protocol (OCSP), the server may also fail authentication if:

- The certificate is revoked.
- The revocation status of the certificate is unknown.
- The connection times out as the firewall is attempting to connect to the CRL/OCSP service.

For more information on certificate profile settings, refer to the steps to [Configure a Certificate Profile](#).



Verify that you added the root CA and intermediate CA of the server to the certificate profile configured with the external dynamic list. Otherwise, the firewall will not authenticate the list properly.

Client authentication fails if you have entered the incorrect username and password combination for the external dynamic list.

STEP 4 | (Optional) [Disable Authentication for an External Dynamic List](#) that failed authentication as a stop-gap measure until the list owner renews the certificate(s) of the server that hosts the list.

Disable Authentication for an External Dynamic List

Palo Alto Networks recommends that you enable authentication for the servers that host the external dynamic lists configured on your firewall. However, if you [Find External Dynamic Lists That Failed Authentication](#) and prefer to disable server authentication for those lists, you can do so through the CLI. The procedure below only applies to external dynamic lists secured with SSL (i.e., lists with an HTTPS URL); the firewall does not enforce server authentication on lists with an HTTP URL.



Disabling server authentication for an external dynamic list also disables client authentication. With client authentication disabled, the firewall will not be able to connect to an external dynamic list that requires a username and password for access.

STEP 1 | [Launch the CLI](#) and switch to configuration mode as follows:

```
username@hostname> configure
Entering configuration mode
[edit]
username@hostname#
```

The change from the > to the # symbol indicates that you are now in configuration mode.

STEP 2 | Enter the appropriate CLI command for the list type:

- IP Address

```
set external-list <external dynamic list name> type ip certificate-profile  
None
```

- Domain

```
set external-list <external dynamic list name> type domain certificate-  
profile None
```

- URL

```
set external-list <external dynamic list name> type url certificate-  
profile None
```

STEP 3 | Verify that authentication is disabled for the external dynamic list.

Trigger a refresh for the list (see [Retrieve an External Dynamic List from the Web Server](#)). If the firewall retrieves the list successfully, server authentication is disabled.

Register IP Addresses and Tags Dynamically

To mitigate the challenges of scale, lack of flexibility, and performance, network architectures today allow for virtual machines (VMs) and applications to be provisioned, changed, and deleted on demand. This agility, though, poses a challenge for security administrators because they have limited visibility in to the IP addresses of the dynamically provisioned VMs and the plethora of applications that can be enabled on these virtual resources.

Firewalls (hardware-based and VM-Series models) support the ability to register IP addresses and tags dynamically. The IP addresses and tags can be registered on the firewall directly or from Panorama. You can also automatically remove tags on the source and destination IP addresses included in a firewall log.

You can enable the dynamic registration process 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 servers, vCenter Servers, or a combination of the two. 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**—Enables you to monitor VMware ESXi, vCenter Server, AWS-VPCs, and Google Compute Engines natively on the firewall and to retrieve IP address changes when you provision or modify virtual machines on these sources. For monitoring virtual machines in your Microsoft Azure deployment, you can deploy the [VM Monitoring script](#) that runs on a virtual machine within the Azure public cloud. This script collects the IP address-to-tag mapping for all your Azure assets and uses the API to push the VM information to your Palo Alto Networks firewalls. VM Information Sources option 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](#).
- **Panorama Plugin**—You can enable a Panorama™ M-Series or virtual appliance to connect to your Azure or AWS public cloud environment and retrieve information on the virtual machines deployed within your subscription or VPC. Panorama then registers the VM information to the managed Palo Alto Networks firewalls that you configured for notification and then you can use these attributes to define dynamic address groups and attach them to Security policy rules to allow or deny traffic to and from these VMs.
- **VMware Service Manager (Integrated NSX solutions only)**—The integrated NSX solution is designed for automated provisioning and distribution of the Palo Alto Networks Next-Generation Security Operating Platform® 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. You can make API calls directly from command-line utilities such, as cURL, or by using any scripting or application framework that supports REST-based services. Refer to the [PAN-OS XML API Usage Guide](#) for details.
- **Auto-Tag**—Tag the source or destination IP address automatically when a log is generated on the firewall and register the IP address and tag mapping to a User-ID agent on the firewall or on Panorama, or to a remote User-ID agent using an HTTP server profile. For example, whenever the firewall generates a threat log, you can configure the firewall to tag the source IP address in the threat log with a specific tag name.

Additionally, you can configure the firewall to dynamically unregister a tag after a configured amount of time using a timeout. For example, you can configure the timeout to be the same duration as the DHCP lease timeout for the IP address. This allows the IP address-to-tag mapping to expire at the same time as the DHCP lease so that you don't unintentionally apply policy when the IP address is reassigned.

See [Forward Logs to an HTTP\(S\) Destination](#).

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](#).

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 on Virtual Machines in Cloud Platforms](#)
- [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, vCenter Server, AWS-VPC, Microsoft Azure VNet, and Google Cloud. 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.

You can directly configure the firewall or use Panorama templates to monitor up to 10 VM information sources. **VM Information Sources** offers easy configuration and enables you to monitor a predefined set of 16 metadata elements or attributes. See [Attributes Monitored on Virtual Machines in Cloud Platforms](#) for the list. By default, the traffic between the firewall and the monitored sources uses the management (MGT) port on the firewall.



- *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.*
- *For monitoring the virtual machines within your Azure deployment, instead of VM Monitoring Sources, you need to deploy the [VM Monitoring script](#) that runs on a virtual machine within the Azure public cloud. This script collects the IP address-to-tag mapping information for your Azure assets and publishes it to the firewalls and corresponding virtual systems you specify in the script.*
- *For Panorama version 8.1.3 and later, you can also use the Panorama plugin for AWS or Azure to retrieve VM Information and register it to the managed firewalls. See [Attributes Monitored on Virtual Machines in Cloud Platforms](#) for details.*

STEP 1 | Enable VM Monitoring.



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**. This example shows you how to add VMware ESX(i) or vCenter Server.
 2. Click **Add** and enter the following information:
 - A **Name** to identify the source 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 an **AWS VPC**, a **Google Compute Engine** instance, a **VMware ESX(i)** server, or a **VMware vCenter** Server.



The type chosen determines the fields displayed.

- Add the credentials (**Username** and **Password**) to authenticate to the server specified above.
- Use the credentials of an administrative user to enable access.
- Define the **Source**.
- (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. (range is 2-10 hours; default is 2).

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.

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 on Virtual Machines in Cloud Platforms

As you provision or remove virtual machines in the private or public cloud, you can use a Panorama plugin, a VM Monitoring script, or the VM Information Source on the next-gen firewall to monitor changes on virtual machines (VMs) deployed in the virtual environments.

VM Information Sources—On a hardware or a VM-Series firewall you can monitor virtual machine instances and retrieves changes as you provision or modify the guests configured on the monitored sources—AWS, ESXi or vCenter Server, or AWS. For each firewall (and/or virtual system if your firewall has multiple virtual system capability), you can configure up to 10 sources. For information on how VM Information Sources and Dynamic Address Groups work synchronously and enable you to monitor changes in the virtual environment, refer to the [VM-Series Deployment Guide](#). If your firewalls are configured in a high availability configuration:

- In an active/passive setup, only the active firewall monitors the VM information sources.
- In an active/active setup, only the primary firewall monitors the VM information sources.

Panorama Plugin—On a Panorama —hardware appliance or virtual appliance running version 8.1.3—you can install the plugin for Microsoft Azure and AWS. The plugin allows you to connect Panorama to your Azure public cloud subscriptions or AWS VPCs and retrieve the IP address-to-tag mapping for your virtual machines. Panorama then registers the VM information to the managed Palo Alto Networks® firewall(s) that you have configured for notification.

Use the following sections to review the options supported on each cloud vendor and the virtual machine attributes that you can monitor to create Dynamic Address Groups:

- [VMware ESXi](#)
- [Amazon Web Services \(AWS\)](#)
- [Microsoft Azure](#)
- [Google Cloud Platform](#)

VMware ESXi

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.



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 32 tags (from vCenter server and NSX Manager) can be registered to an IP address.

To collect the values assigned to the monitored VMs, use the VM Information Sources on the firewall to monitor the following predefined set of ESXi attributes:

Attributes Monitored on a VMware Source

UUID

Name

Guest OS

VM State – the power state can be poweredOff, poweredOn, standBy, and unknown.

Annotation

Version

Network – Virtual Switch Name, Port Group Name, and VLAN ID

Container Name –vCenter Name, Data Center Object Name, Resource Pool Name, Cluster Name, Host, Host IP address.

Amazon Web Services (AWS)

As you provision or modify virtual machines in your AWS VPCs, you have two ways of monitoring these instances and retrieving the tags for use as match criteria in dynamic address groups.

- **VM Information Source**—On a next-gen firewall, you can monitor up to a total of 32 tags—14 pre-defined and 18 user-defined key-value pairs (tags). The following attributes (or tag names) are available as match criteria for dynamic address groups.
- **AWS Plugin on Panorama**—The [Panorama plugin for AWS](#) allows you to connect Panorama to your AWS VPCs and retrieve the IP address-to-tag mapping for your AWS virtual machines. Panorama then registers the VM information to the managed Palo Alto Networks® firewall(s) that you have configured for notification. With the plugin, Panorama can retrieve a total of 32 tags for each virtual machine, 11 predefined tags and up to 21 user-defined tags.

Attributes Monitored on the AWS-VPC	VM Information Source on the Firewall	AWS Plugin on Panorama
Architecture	Yes	No
Guest OS	Yes	No
AMI ID	Yes	Yes
IAM Instance Profile	No	Yes
Instance ID	Yes	No
Instance State	Yes	No

Attributes Monitored on the AWS-VPC	VM Information Source on the Firewall	AWS Plugin on Panorama
Instance Type	Yes	No
Key Name	Yes	Yes
Owner ID	No	Yes
Placement—Tenancy	Yes	Yes
Placement—Group Name	Yes	Yes
Placement—Availability Zone	Yes	Yes
Private DNS Name	Yes	No
Public DNS Name	Yes	Yes
Subnet ID	Yes	Yes
Security Group ID	No	Yes
Security Group Name	No	Yes
VPC ID	Yes	Yes
Tag (key, value)	Yes; Up to a maximum of 18 user defined tags are supported. The user-defined tags are sorted alphabetically, and the first 18 tags are available for use on the firewalls.	Yes; Up to a maximum of 21 user defined tags are supported. The user-defined tags are sorted alphabetically, and the first 21 tags are available for use on Panorama and the firewalls.

Microsoft Azure

For [VM Monitoring on Azure](#) you need to retrieve the IP address-to-tag mapping for your Azure VMs and make it available as match criteria in dynamic address groups. The [Panorama plugin for Microsoft Azure](#) allows you to connect Panorama to your Azure public cloud subscriptions and retrieve the IP address-to-tag mapping for your Azure virtual machines. Panorama can retrieve a total of 26 tags for each virtual machine, 11 predefined tags and up to 15 user-defined tags and registers the VM information to the managed Palo Alto Networks® firewall(s) that you have configured for notification.

With the Panorama plugin for Azure, you can monitor the following set of virtual machine attributes within your Microsoft Azure deployment.

Attributes Monitored on Microsoft Azure	Azure Plugin on Panorama
VM Name	Yes

Attributes Monitored on Microsoft Azure	Azure Plugin on Panorama
VM Size	No
Network Security Group Name	Yes
OS Type	Yes
OS Publisher	Yes
OS Offer	Yes
OS SKU	Yes
Subnet	Yes
VNet	Yes
Azure Region	Yes
Resource Group Name	Yes
Subscription ID	Yes
User Defined Tags	Yes Up to a maximum of 15 user defined tags are supported. The user-defined tags are sorted alphabetically, and the first 15 tags are available for use on Panorama and the firewalls.

Google Cloud Platform

Using VM Information Sources on the next-gen firewall, you can monitor the following predefined set of Google Compute Engine (GCE) attributes. These Predefined attributes are created when you create certain assets in GCP.



High Availability is not supported on the firewalls.

Attributes Monitored on GCE VMs	Comment
Hostname of the VM	
Machine type	
Project ID	
Source (OS type)	
Status	

Attributes Monitored on GCE VMs	Comment
Subnetwork	
VPC Network	

Predefined Attributes

- **Project ID**—For example: google.project-id.myProjectId.
To find your project information in the Google console, select your project, then select **IAM & Admin > Settings**.
- **Service account**—Your [service account](#) in the form of an email address. For example: google.svc-accnt.sa-name@project-id.iam.gserviceaccount.com
To find your [Service account](#) view the VM instance details.
- **VPC name**—The name of a [VPC](#) network that is used for a VM-Series firewall interface. For example: google.vpc-name.myvnet.
- **Subnet name**—The name of a [subnet](#) you created for the VM-Series firewall. For example: google.untrust.web.
- **OS SKU**—The operating system you chose when you deployed a VM. For example: google.os-sku.centos-7.



This attribute is not supported if the VM uses a custom image.

- **Google zone**—The [zone](#) you selected when you deployed the VM-Series firewall. For example: google.zone.asia-east1-a.
- **Google region**—The [region](#) containing the zone you selected. For example: google.zone.asia-east1.
- **Instance group name**—[Instance group](#) name is for GKE deployments. For example: google.instance-group.<myInstanceGroup>.

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. After receiving 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 model is different. Use the following table for specifics on your model:

Model	Maximum number of dynamically registered IP addresses
M-Series and Panorama Virtual Appliances	500,000
PA-5200 Series, VM-7000 SMC-B Series	500,000
VM-500, VM-700	300,000
PA-3200 Series, VM-300	200,000
PA-7000 Series, VM-1000-HV	100,000
PA-850, VM-100	2,500
PA-820, PA-220, VM-50	1,000

The following example shows how dynamic address groups can simplify network security enforcement. The example workflow shows how to:

- Enable the VM Monitoring agent on the firewall, to monitor the VMware ESX(i) host or vCenter Server and register VM IP addresses and the associated tags.
- Create dynamic address groups and define the tags to filter. In this example, two address groups are created. One that only filters for dynamic tags and another that filters for both static and dynamic tags to populate the members of the group.
- Validate that the members of the dynamic address group are populated on the firewall.
- Use dynamic address groups in policy. This example uses two different security policies:
 - A security policy for all Linux servers that are deployed as FTP servers; this rule matches on dynamically registered tags.
 - A security policy for all Linux servers that are deployed as web servers; this rule matches on a dynamic address group that uses static and dynamic tags.
- Validate that the members of the dynamic address groups are updated as new FTP or web servers are deployed. This ensure that the security rules are enforced on these new virtual machines too.

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.



View 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**.

STEP 3 | 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')

Name	Location	Members Count	Addresses
ftp_servers		dynamic	more...
web_servers		dynamic	more...

Click to see Members/registered IP addresses

STEP 4 | Use dynamic address groups in policy.



View the [tutorial](#).

1. Select **Policies > Security**.
2. Click **Add** and enter a **Name** and a **Description** for the policy.
3. Add the **Source Zone** to specify the zone from which the traffic originates.
4. Add the **Destination Zone** at which the traffic is terminating.
5. For the **Destination Address**, select the Dynamic address group you 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**.

STEP 5 | This example shows how to create two policies: one for all access to FTP servers and the other for access to web servers.

		Source					Destination						
Name	Tags	Zone	Address	User	HQ Profile	Zone	Address	Application	Service	Action	Profile	Options	
1. Access to web servers		any	any	any	any	untrust	Web_servers	any	application-d...	✓			
2. Access to FTP servers		any	any	any	any	untrust	Ftp_servers	ftp	application-d...	✓			

STEP 6 | 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.

Address Groups - Ftp_servers	
Address	Type
10.5.124.45	registered-ip
15.0.0.45	registered-ip
fe80::250:56ff:fe6d:b6aa	registered-ip
fe80::250:56ff:fe6d:cecd	registered-ip

Policy will be enforced for all IP addresses that belong to this address group, and are displayed here.



*If you want to delete all registered IP addresses, use the CLI command **debug object registered-ip clear all** and then reboot the firewall after clearing the tags.*

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, <code>state.poweredOn</code> or that are not tagged as <code>vSwitch0</code> .	<pre>show log iptag tag_name equal state.poweredOn show log iptag tag_name not-equal switch.vSwitch0</pre>
View all dynamically registered IP addresses that were sourced by VM Information Source with name <code>vmware1</code> and tagged as <code>poweredOn</code> .	<pre>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"</pre>
Clear all IP addresses and tags learned from a specific VM Monitoring source without disconnecting the source.	<pre>debug vm-monitor clear source-name <name></pre>
Display IP addresses registered from all sources.	<pre>show object registered-ip all</pre>
Display the count for IP addresses registered from all sources.	<pre>show object registered-ip all option count</pre>
Clear IP addresses registered from all sources	<pre>debug object registered-ip clear all</pre>
Add or delete tags for a given IP address that was registered using the XML API.	<pre>debug object registered-ip test [<register/unregister>] <ip/netmask><tag></pre>
View all tags registered from a specific information source.	<pre>show vm-monitor source source-name vmware1 tag all</pre>

Example	CLI Command
	<pre> 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 </pre>
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: <pre>show log iptag datasource_type equal unknown</pre> - To view tags registered from the XML API: <pre>show log iptag datasource_type equal xml-api</pre> - To view tags registered from VM Information sources: <pre>show log iptag datasource_type equal vm-monitor</pre> - To view tags registered from the Windows User-ID agent: <pre>show log iptag datasource_type equal xml-api datasource_subtype equal user-id-agent</pre>
View all tags that are registered for a specific IP address (across all sources).	<pre>debug object registered-ip show tag-source ip ip_address tag all</pre>

Identify Users Connected through a Proxy Server

If you have a proxy server deployed between the users on your network and the firewall, the firewall might see the proxy server IP address as the source IP address in HTTP/HTTPS 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 HTTP requests that include 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 extract the end user IP address from the XFF so that User-ID can map that IP address to the username of the end user. This enables you enforce user-based policy to safely enable access to web-based applications for your users behind a proxy server. In addition, if User-ID is able to map the XFF IP address to a username, the firewall displays that username as the Source user in Traffic, Threat, WildFire Submissions, and URL Filtering logs for visibility into the web activity of users behind the proxy.

To use the XFF header for user mapping:

- ❑ The XFF header your proxy server adds must contain the source IP address of the end user who originated the request. If the header contains multiple IP addresses, the firewall uses the first IP address only. If the header contains information other than an IP address, the firewall will not be able to perform user mapping.
- ❑ You must [Enable User-ID](#).

With this option enabled, the firewall uses the IP address in the XFF header for user mapping purposes only. The source IP address the firewall logs is still that of the proxy server, not that of the source user. When you see a log event attributed to a user that the firewall mapped using an IP address extracted from an XFF header, it can be difficult to track down the specific device associated with the event. To simplify debugging and troubleshooting of events attributed to users behind the proxy server, you must also configure the firewall to populate the X-Forwarded-For column in the URL Filtering log with the IP address in the XFF header so that you can track down the specific user and device associated with an log event that is correlated with the URL Filtering log entry.

- [Use XFF Values for Policies and Logging Source Users](#)
- [Use the IP Address in the XFF Header to Troubleshoot Events](#)

Use XFF Values for Policies and Logging Source Users

You can configure the firewall map the IP address in the XFF header to a username using User-ID so that you can have visibility into and user-based policy control over the web traffic of users behind a proxy server who cannot otherwise be identified. In order to map the IP addresses from the XFF headers to usernames, you must first [Enable User-ID](#).



Enabling the firewall to use the X-Forwarded-For headers to perform user mapping does not enable the firewall to use the client IP address in the XFF header as the source address in the logs; the logs still display the proxy server IP address as the source address. However, to simplify the debugging and troubleshooting process you can configure the firewall to [Add XFF Values to URL Filtering Logs](#) to display the client IP address from the XFF header in the URL Filtering 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 policy enforcement and logging.

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**.

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 web applications.

Use the IP Address in the XFF Header to Troubleshoot Events

By default, the firewall does not log the source address of a client behind a proxy server, even if you are using this address from the X-Forwarded-For (XFF) header for user mapping. Therefore, while you can identify the specific user associated with a log event, you will not be able to easily identify the source device that originated the log event. To simplify the debugging and troubleshooting of events for users behind a proxy server, you must enable the X-Forwarded-For option within HTTP Header Logging in the URL Filtering profile that you attach to security policy rules that allow access to web-based applications. With this option enabled, the firewall logs the IP address from the XFF header as the Source address for all traffic that matches the rule.



Enabling the firewall to use the XFF header as the Source address in URL Filtering logs does not enable user mapping of the source address. To populate the source user fields, see [Use XFF Values for Policies and Logging Source Users](#).

STEP 1 | Enable the X-Forwarded-For option within HTTP Header Logging in the URL Filtering profile.

1. Select **Objects > Security Profiles > URL Filtering** and select the URL Filtering profile you want to configure, or [add](#) a new one.



You can't enable XFF logging in the default URL Filtering profile.

2. Select the **Settings** tab and select **X-Forwarded-For**.
3. Click **OK** to save the profile.

STEP 2 | Attach the URL Filtering profile to the security policy rule(s) that enable access to web applications.

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 configured for X-Forwarded-For HTTP Header Logging.
3. Click **OK** and **Commit**.

STEP 3 | Verify the firewall is logging XFF values.

1. Select **Monitor > Logs > URL Filtering**.
2. View the XFF values in one of the following ways:

-
- To display the XFF value for a single URL Filtering log—Click the spyglass icon for the log to display its details. The HTTP Headers section displays the X-Forwarded-For value.
 - To display the XFF values for all URL Filtering 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.

STEP 4 | Use the XFF field in the URL Filtering log to troubleshoot a log event in another log type.

Although only the URL Filtering logs display the IP address of the source user in the X-Forwarded-For column of the logs, if you notice an event associated with HTTP/HTTPS traffic but that you cannot identify the source IP address because it is that of the proxy server, you can use the X-Forwarded-For value in a correlated URL Filtering log to help you identify the source address associated with the log event. To do this:

1. Find an event you want investigate in a Traffic, Threat, or WildFire Submissions logs that is showing the IP address of the proxy server as the source address.
2. Click the spyglass icon for the log to display its details and look for an associated URL Filtering log at the bottom of the Detailed Log Viewer window.
3. Select the header row and then select **X-Forwarded-For** from the **Columns** drop-down to display this value. The IP address in this column of the X-Forwarded-For column represents the IP address of the source user behind the proxy server. Use this IP address to track down the device that triggered the event you are investigating.

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 silently dropped.



*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	If rule is disabled when the monitored IP address is unreachable
For an established session	wait-recover —Continue to use egress interface specified in the PBF rule	wait-recover —Continue to use egress interface specified in the PBF rule
	fail-over —Use path determined by routing table (no PBF)	fail-over —Use path determined by routing table (no PBF)
For a new session	wait-recover —Use path determined by routing table (no PBF)	wait-recover —Check the remaining PBF rules. If no match, use the routing table
	fail-over —Use path determined by routing table (no PBF)	fail-over —Check the remaining PBF rules. If no match, use the routing table

Service Versus Applications in PBF

PBF rules are applied either on the first packet (SYN) or the first response to the first packet (SYN/ACK). This means that a PBF rule may be applied before the firewall has enough information to determine the application. Therefore, application-specific rules are not recommended for use with PBF. Whenever possible, use a service object, which is the Layer 4 port (TCP or UDP) used by the protocol or application.

However, if you specify an application in a PBF rule, the firewall performs *App-ID caching*. When an application passes through the firewall for the first time, the firewall does not have enough information to identify the application and therefore cannot enforce the PBF rule. As more packets arrive, the firewall determines the application and creates an entry in the App-ID cache and retains this App-ID for the session. When a new session is created with the same destination IP address, destination port, and protocol ID, the firewall could identify the application as the same from the initial session (based on the App-ID cache) and apply the PBF rule. Therefore, a session that is not an exact match and is not the same application, can be forwarded based on the PBF rule.

Further, applications have dependencies and the identity of the application can change as the firewall receives more packets. Because PBF makes a routing decision at the start of a session, the firewall cannot enforce a change in application identity. YouTube, for example, starts as web-browsing but changes to Flash, RTSP, or YouTube based on the different links and videos included on the page. However with PBF, because the firewall identifies the application as web-browsing at the start of the session, the change in application is not recognized thereafter.



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.

STEP 1 | Create a Policy-Based Forwarding (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.



You can specify the source and destination addresses using an IP address, an address object, or an FQDN.

1. Select **Policies > Policy Based Forwarding** and **Add** a PBF policy rule.
2. Give the rule a descriptive name (**General**).
3. Select **Source** and configure the following:
 1. Select the **Type (Zone or Interface)** to which you will apply the forwarding policy and specify the relevant zone or interface. If you want to enforce symmetric return, you must select a source interface.



Only Layer 3 interfaces support PBF; loopback interfaces do not support PBF.

2. (**Optional**) Specify the **Source Address** to which the PBF rule applies. 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.



*Click **Negate** to exclude one or more Source 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, subsequent rules are not evaluated.

3. (**Optional**) **Add** and select the **Source User** or groups of users to whom the policy applies.
4. Select **Destination/Application/Service** and configure the following:
 1. **Destination Address**—By default, the rule applies to **Any** IP address. Click **Negate** to exclude one or more destination IP addresses from the PBF rule.
 2. **Add** any **Application** and **Service** that you want to control using PBF.



We do not recommend application-specific rules for use with PBF because PBF rules may be applied before the firewall has enough information to determine the application. 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](#).

STEP 2 | Specify how to forward packets that match 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. Select **Forwarding**.
2. Set the **Action** to take when matching a packet:

- **Forward**—Directs the packet to the specified **Egress Interface**.
 - **Forward to VSYS** (On a firewall enabled for multiple virtual systems)—Select the virtual system to which to forward the packet.
 - **Discard**—Drops the packet.
 - **No PBF**—Excludes packets that match the criteria for source, destination, application, or 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.
3. To trigger the specified **Action** at a daily, weekly, or non-recurring frequency, create and attach a **Schedule**.
 4. For **Next Hop**, select one of the following:
 - **IP Address**—Enter an IP address or select an address object of type IP Netmask to which the firewall forwards matching packets. An IPv4 address object must have a /32 netmask and an IPv6 address object must have a /128 netmask.
 - **FQDN**—Enter an FQDN (or select or create an address object of type FQDN) to which the firewall forwards matching packets. The FQDN can resolve to an IPv4 address, an IPv6 address, or both. If the FQDN resolves to both IPv4 and IPv6 addresses, then the PBF rule has two next hops: one IPv4 address and one IPv6 address. You can use the same PBF rule for both IPv4 and IPv6 traffic. IPv4 traffic is forwarded to the IPv4 next hop; IPv6 traffic is forwarded to the IPv6 next hop.
-  *This FQDN must resolve to an IP address that belongs to the same subnet as the interface you configured for PBF; otherwise, the firewall rejects the resolution and the FQDN remains unresolved.*
-  *The firewall uses only one IP address (from each IPv4 or IPv6 family type) from the DNS resolution of the FQDN. If the DNS resolution returns more than one address, the firewall uses the preferred IP address that matches the IP family type (IPv4 or IPv6) configured for the next hop. The preferred IP address is the first address the DNS server returns in its initial response. The firewall retains this address as preferred as long as the address appears in subsequent responses, regardless of order.*
- **None**—No next hop means the destination IP address of the packet is used as the next hop. Forwarding fails if the destination IP address is not in the same subnet as the egress interface.
5. (Optional) Enable monitoring to verify connectivity to a target IP address or to the **Next Hop IP** address if no IP address is specified. Select **Monitor** and attach a monitoring **Profile** (default or custom) that specifies the action when the monitored address is unreachable.
- You can **Disable this rule if nexthop/monitor ip is unreachable**.
 - Enter a target **IP Address** to monitor.

The **Egress Interface** can have both IPv4 and IPv6 addresses and the **Next Hop FQDN** can resolve to both IPv4 and IPv6 addresses. In this case:

1. If the egress interface has both IPv4 and IPv6 addresses and the next hop FQDN resolves to only one address family type, the firewall monitors the resolved IP address. If the FQDN resolves to both IPv4 and IPv6 addresses but the egress interface has only one address family type address, the firewall monitors the resolved next hop address that matches the address family of the egress interface.
 2. If both the egress interface and next hop FQDN have both IPv4 and IPv6 addresses, the firewall monitors the IPv4 next hop address.
 3. If the egress interface has one address family address and the next hop FQDN resolves to a different address family address, the firewall does not monitor anything.
6. (Required for asymmetric routing environments; otherwise, optional) **Enforce Symmetric Return** and **Add** one or more IP addresses in the **Next Hop Address List**. 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 (such as 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 | Commit your changes. The PBF rule is in effect.

Name	Tags	Zone/Interface	Address	User	Address	Application	Service	Action	Egress IP	Next Hop	Enforce Symmetric Return	Profile	Target	Default IP	Schedule
HTTP to ISP	none	TwoISP	any	any	HQ-Subnet	any	service:http	Forward	ethernet1/3	10.0.0.0	False	none	none	False	none

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.



STEP 1 | Configure the ingress and the egress interfaces on the firewall.

Egress interfaces can be in the same zone.

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: TwoISP
 - IP Address: 1.1.1.2/30
 - Virtual Router: Default
- Ethernet 1/3 connected to the backup ISP:
 - Zone: TwoISP
 - 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: Corporate
 - 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.

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 Corporate.
4. In the **Destination/Application/Service** tab, set the following:
 1. 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.
 2. 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.



STEP 4 | Specify where to forward traffic.

1. In the **Forwarding** tab, specify the interface to which you want to forward traffic and enable path monitoring.
2. 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).



3. 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.
4. (Required if you have asymmetric routes) Select **Enforce Symmetric Return** to ensure that return traffic from the Corporate zone to the internet is forwarded out on the same interface through which traffic ingress from the internet.
5. NAT ensures that the traffic from the internet is returned to the correct interface/IP address on the firewall.
6. Click **OK** to save the changes.

Name	Zone/Interface	Address	Address	Service	Action	Egress IP	Next Hop	Enforce Symmetric Return
Use ISP Primary	Corporate	any	Internet-external	Service-HTTP	Forward	ethernet1/1	1.1.1.1	Yes

STEP 5 | 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: Corporate

Destination Zone: TwoISP

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: Corporate

Destination Zone: TwoISP

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



Name	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation
NAT for Backup ISP	Corporate	TwoISP	eth1	any	any	any	dynamic-ip-and-port	none
NAT for Primary ISP	Corporate	TwoISP	eth2	any	any	any	dynamic-ip-and-port	none

STEP 6 | 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 Corporate.
4. In the **Destination** tab, Set the **Destination Zone** to TwoISP.
5. In the **Service/ URL Category** tab, leave the default **application-default**.
6. In the **Actions** tab, complete these tasks:
 1. Set the **Action Setting** to **Allow**.
 2. 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.



Name	Zone	Address	Zone	Address	Application	Service	Action	Profile
SecPol1	Corporate	any	TwoISP	any	any	application-d	allow	none

STEP 7 | Save the policies to the running configuration on the firewall.

Click **Commit**.

STEP 8 | 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.

- 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 198.51.100.6
Pinging 198.51.100.6 with 32 bytes of data:
Reply from 198.51.100.6: bytes=32 time=34ms TTL=117
Reply from 198.51.100.6: bytes=32 time=13ms TTL=117
Reply from 198.51.100.6: bytes=32 time=25ms TTL=117
Reply from 198.51.100.6: bytes=32 time=3ms TTL=117
Ping statistics for 198.51.100.6:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 3ms, Maximum = 34ms, Average = 18ms
```



As defined by the PBF rule, only traffic on ports 80 or 443 use the Primary ISP, hence ping is sent through the interface attached to the backup ISP.

The security policy that allows the traffic.

Session Time	Type	From Zone	To Zone	Source	Destination	To Port	Application	Action	Rule
11/05/2019 03:02	out	Corporate	Untrust	192.168.24.26	198.51.100.6	80	ping	allow	Untrust

- 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 9 | Verify that the failover to the backup ISP occurs and that the Source NAT is correctly applied.

- Unplug the connection to the primary ISP.
- 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
```

- Access a web server, and check the traffic log to verify that traffic is being forwarded through the backup ISP.



Traffic is sent through the interface attached to the backup ISP.

The security policy that allows the traffic.

Session Time	Type	From Zone	To Zone	Source	Destination	To Port	Application	Action	Rule
11/05/2019 03:04	out	Corporate	Untrust	192.168.24.26	204.79.197.200	443	ssl	allow	Untrust
11/05/2019 03:04	out	Corporate	Untrust	192.168.24.26	204.79.197.200	80	web-forwarding	allow	Untrust

- 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[Source]/Zone/Proto (translated
IP[Port]) Vsys Dst[Destination]/Zone (translated IP[Port])
-----
87212 ssl ACTIVE    FLOW  NS   192.168.54.56[53236]/Corporate/6
(2.2.2.2[12896]) vsys1 204.79.197.200[443]/TwoISP (204.79.197.200[443])
```

- Obtain the session identification number from the output and view the session details.



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 [Corporate]
    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 [TwoISP]
    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                   : Corp2ISP
  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)

```

Test Policy Rules

Test the policy rules in your running configuration to ensure that your policies appropriately allow and deny traffic and access to applications and websites in compliance with your business needs and requirements. You can test and verify that your policy rules are allowing and denying the correct traffic by executing policy match tests for your firewalls directly from the web interface.

STEP 1 | Launch the Web Interface.

STEP 2 | Select **Device > Troubleshooting to perform a policy match or connectivity test.**

STEP 3 | Enter the required information to perform the policy match test. In this example, we run a NAT policy match test.

1. **Select Test**—Select **NAT Policy Match**.
2. **From**—Select the zone traffic is originating from.
3. **To**—Select the target zone of the traffic.
4. **Source**—Enter the IP address from which traffic originated.
5. **Destination**—Enter the IP address of the target device for the traffic.
6. **Destination Port**—Enter the port used for the traffic. This port varies depending on the IP protocol used in the following step.
7. **Protocol**—Enter the IP protocol used for the traffic.
8. If necessary, enter any additional information relevant for your NAT policy rule testing.

STEP 4 | Execute the NAT policy match test.

STEP 5 | Review the **NAT Policy Match Result to see the policy rules that match the test criteria.**

Test Configuration	Test Result	Result Detail				
Select Test: NAT Policy Match From: D-vlan-trust To: D-untrust Source: Destination: Source Port: [] - [] Destination Port: 445 Protocol: 6 To Interface: None Ha Device ID: [] - [] Execute Reset	NAT Policy Match Result	<table border="1"><thead><tr><th>Name</th><th>Value</th></tr></thead><tbody><tr><td>Result</td><td>access-corp</td></tr></tbody></table>	Name	Value	Result	access-corp
Name	Value					
Result	access-corp					

Virtual Systems

This topic describes virtual systems, their benefits, typical use cases, and how to configure them. It also provides links to other topics where virtual systems are documented as they function with other features.

- > [Virtual Systems Overview](#)
- > [Communication Between Virtual Systems](#)
- > [Shared Gateway](#)
- > [Configure Virtual Systems](#)
- > [Configure Inter-Virtual System Communication within the Firewall](#)
- > [Configure a Shared Gateway](#)
- > [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.

- [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, Tunnel Inspection, Authentication, and DoS protection)
- All objects (such as address objects, application groups and filters, external dynamic lists, security profiles, decryption profiles, custom objects, etc.)
- User-ID

-
- Certificate management
 - Server profiles
 - Logging, reporting, and visibility functions

Virtual systems affect the security functions of the firewall, but virtual systems alone do not affect networking functions such as static and dynamic routing. You can segment routing for each virtual system by creating one or more virtual routers for each virtual system, as in the following use cases:

- If you have virtual systems for departments of one organization, and the network traffic for all of the departments is within a common network, you can create a single virtual router for multiple virtual systems.
- If you want routing segmentation and each virtual system's traffic must be isolated from other virtual systems, you can create one or more virtual routers for each virtual system.
- If you want to segment the user mappings so that not all mappings are shared across virtual systems, you can configure the User-ID sources on a virtual system that is not a User-ID hub. See [Share User-ID Mappings Across Virtual Systems](#).

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.
- **Ability to share IP-address-to-username mappings**—By assigning a virtual system as a User-ID hub, you can share the IP-address-to-username mappings across virtual systems to leverage the full User-ID capacity of the firewall and reduce operational complexity.

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 PA-3200 Series, PA-5200 Series, 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 to support multiple virtual systems on PA-3200 Series firewalls, and to create more than the base number of virtual systems supported on a platform.

For license information, see [Subscriptions](#). For the base and maximum number of virtual systems supported, see [Compare Firewalls](#) tool.

Multiple virtual systems are not supported on the PA-220, PA-800 Series, or VM-Series firewalls.



The default is vsys1. You cannot delete vsys1 because it is relevant to the internal hierarchy on the firewall; vsys1 appears even on firewall models that don't support multiple virtual systems.

You can [limit the resource allocations](#) for sessions, rules and VPN tunnels allowed for a virtual system, and thereby control firewall resources. Each resource setting displays the valid range of values, which [varies per firewall model](#). The default setting is 0, which means the limit for the virtual system is the limit for the firewall model. However, the limit for a specific setting isn't replicated for each virtual system. For example, if a firewall has four virtual systems, each virtual system can't have the total number of Decryption Rules allowed per firewall. After the total number of Decryption Rules for all of the virtual systems reaches the firewall limit, you cannot add more.

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. When you create an Admin Role profile and select the role to be **Virtual System**, the role applies to specific virtual systems on the firewall. From the **Command Line** tab, the two types of virtual system administrative roles are:

- **vsysadmin**—Has access to specific virtual systems on the firewall to create and manage specific aspects of virtual systems. A vsysadmin doesn't have access to network interfaces, VLANs, virtual wires, virtual routers, IPSec tunnels, GRE tunnels, DHCP, DNS Proxy, QoS, LLDP, or network profiles. Persons with vsysadmin permission can commit configurations for only the virtual systems assigned to them.
- **vsysreader**—Has read-only access to specific virtual systems on the firewall and specific aspects of virtual systems. A vsysreader doesn't have access to network interfaces, VLANs, virtual wires, virtual routers, IPSec tunnels, GRE tunnels, DHCP, DNS Proxy, QoS, LLDP, or network profiles.

A virtual system administrator can view logs of only the virtual systems assigned to that administrator. A **Superuser** or **Device administrator** can view all of the logs, select a virtual system to view, or configure a virtual system as a User-ID hub.

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.

Communication Between Virtual Systems


There are two typical scenarios where communication between virtual systems (inter-vsyz 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.

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.

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 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 Authentication 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 vsys1 to vsys2 to vsys3, or similarly from vsys1 to vsys2 to shared gateway1. Both examples involve more than two virtual systems, which is not permitted.

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.

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 creating more than the base number of virtual systems supported on the platform. See [Platform Support and Licensing for Virtual Systems](#).

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 is vsys1. You cannot delete vsys1 because it is relevant to the internal hierarchy on the firewall; vsys1 appears even on firewall models that don't support multiple virtual systems.

2. Select **Allow forwarding of decrypted content** 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.

STEP 3 | Assign interfaces to the virtual system.

The virtual routers, virtual wires, 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. Each field displays the valid range of values, which varies per firewall model. The default setting is 0, which means the limit for the virtual system is the limit for the firewall model. However, the limit for a specific setting isn't replicated for each virtual system. For example, if a firewall has four virtual systems, each virtual system can't have the total number of Decryption Rules allowed per firewall. After the total number of Decryption Rules for all of the virtual systems reaches the firewall limit, you cannot add more.

- **Sessions Limit**



If you use the `show session meter CLI` command, it displays the Maximum number of sessions allowed per dataplane, the Current number of sessions being used by the virtual system, and the Throttled number of sessions per virtual system. On a PA-5200 or PA-7000 Series firewall, the Current number of sessions being used can be greater than the Maximum configured for Sessions Limit because there are multiple dataplanes per virtual system. The Sessions Limit you configure on a PA-5200 Series or PA-7000 Series firewall is per dataplane, and will result in a higher maximum per virtual system.

- Security Rules
- NAT Rules
- Decryption Rules
- QoS Rules
- Application Override Rules
- Policy Based Forwarding Rules
- Authentication Rules
- DoS Protection Rules
- Site to Site VPN Tunnels
- Concurrent SSL VPN Tunnels


2. Click **OK**.

STEP 5 | (Optional) Configure a virtual system as a User-ID hub to [Share User-ID Mappings Across Virtual Systems](#).



IP-address-and-port-to-username mapping information from Terminal Server agents and group mapping data is not shared between the virtual system hub and the connected virtual systems.

1. For any existing virtual systems, transfer the configuration for the User-ID sources you want to share (such as monitored servers and User-ID agents) to the virtual system you will use as a hub.
2. On the **Resource** tab, select **Make this vsys a User-ID data hub**.



3. Click **Yes** to confirm, then click **OK**.

If you want to change the User-ID hub to a different virtual system or disable it, select the virtual system currently configured as a User-ID hub, then select **Resource > Change Hub**.

Virtual System

ID 1

☐ Allow forwarding of decrypted content

Name

General

Resource

Sessions Limit [1 - 524288]

Policy Limits

Security Rules [0 - 5000]

NAT Rules [0 - 5000]

Decryption Rules [0 - 500]

QoS Rules [0 - 1000]

Application Override Rules [0 - 500]

Policy Based Forwarding Rules [0 - 500]

Authentication Rules [0 - 1000]

DoS Protection Rules [0 - 1000]

VPN Limits

Site to Site VPN Tunnels [0 - 2048]

Concurrent SSL VPN Tunnels [0 - 2048]

Inter-Vsys User-ID Data Sharing

User-ID hub is vsys1 [Change Hub](#)

OK

Cancel

Select the **New User-ID hub** from the list, or select **none** to disable the User-ID hub and stop sharing mappings across virtual systems.

Inter-Vsys User-ID Data Sharing

By changing the User-ID hub, other virtual systems will not be able to access the current User-ID hub (vsys-place-holder). This could have an affect on policy matching and getting User based visibility on other virtual systems.

New User-ID hub vsys1

▼

Proceed

Cancel

Click **Proceed** to confirm and commit your changes.

STEP 6 | Commit the configuration.

Click **Commit**. The virtual system is now an object accessible from the **Objects** tab.

STEP 7 | 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 select the interfaces that belong to the virtual router.
3. Click **OK**.

STEP 8 | Configure a security zone for each interface in the virtual system.

For at least one interface, create a Layer 3 security zone. See [Configure Interfaces and Zones](#).

STEP 9 | 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 10 | Commit the configuration.

Click **Commit**.



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 11 | (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.

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. (Optional) Select a **Zone Protection Profile** (or configure one later) that provides flood, reconnaissance, or packet-based attack protection.
6. (Optional) In **Log Setting**, select a log forwarding profile for forwarding zone protection logs to an external system.
7. (Optional) Select **Enable User Identification** 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 | Commit your changes.

Click **Commit**.

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.

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 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. (Optional) Select a **Zone Protection Profile** (or configure one later) that provides flood, reconnaissance, or packet-based attack protection.
5. (Optional) In **Log Setting**, select a log forwarding profile for forwarding zone protection logs to an external system.
6. (Optional) Select **Enable User Identification** to enable User-ID for the shared gateway.
7. Click **OK**.

STEP 3 | Commit your changes.

Click **Commit**.

Customize Service Routes for a Virtual System

When a firewall is enabled for multiple virtual systems, the virtual systems inherit the global service and service route settings. For example, the firewall can use a shared email server to originate email alerts to all virtual systems. In some scenarios, you'd want to create different service routes for each virtual system.

One use case for configuring service routes at the virtual system level is if you are an ISP who needs to support multiple individual tenants on a single Palo Alto Networks firewall. Each tenant requires custom service routes to access service such as DNS, Kerberos, LDAP, NetFlow, RADIUS, TACACS+, Multi-Factor Authentication, email, SNMP trap, syslog, HTTP, User-ID Agent, VM Monitor, and Panorama (deployment of content and software updates). 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.



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.

- [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.

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.

The service route for a service strictly follows how you configured the server profile for the service:

- If you define a server profile (**Device > Server Profiles**) for the Shared location, the firewall uses the global service route for that service.
- If you define a server profile for a specific virtual system, the firewall uses the virtual system-specific service route for that service.
- If you define a server profile for a specific virtual system but the virtual system-specific service route for that service is not configured, the firewall uses the global service route for that service.



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.

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:
 - **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 the step to customize.
 - **Customize**—Allows you to specify a 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 on a service. (Only services that are relevant to a virtual system are available.)



*To easily use the same source address for multiple services, select the checkbox for the services, click **Set Selected Routes**, and continue.*

- To limit the list for Source Address, select a **Source Interface**, then select a Source Address (from that interface) as the service route. Selecting **Any** Source Interface makes all IP addresses on all interfaces for the virtual system available in the Source Address list from which you select an address. You can select **Inherit Global Setting**.
 - **Source Address** will indicate **Inherited** if you selected **Inherit Global Setting** for the **Source Interface** or it will indicate the source address you selected. If you selected **Any** for **Source Interface**, select an IP address 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 the prior steps to configure source addresses for other external services.
 7. Click **OK**.

STEP 2 | Commit your changes.

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

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 supports using a logging card.

Depending on your firewall configuration, you might have one of the following card types:

- **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. If your firewall has an LPC installed, you need to configure a log card port.
- **Log Forwarding Card (LFC)**—Supports high-speed log forwarding of all dataplane logs to an external log collector (for example, Panorama and syslog servers). You can create and configure subinterfaces for virtual systems. If your firewall has an LFC installed, you do not need to configure a log card port.

In other Palo Alto Networks models, the dataplane sends logging service route traffic to the management plane, which sends the traffic to logging servers. In a PA-7000 Series firewall, the LPC or LFC have only one

interface, and dataplanes for multiple virtual systems send logging server traffic (types mentioned above) to the PA-7000 Series firewall logging card. The logging card 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 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. For information about the LPC or LFC, see the [PA-7000 Series Hardware Reference Guide](#).

- [Configure a PA-7000 Series LPC for Logging per Virtual System](#)
- [Configure a PA-7000 Series LFC for Logging per Virtual System](#)

Configure a PA-7000 Series LPC for Logging per Virtual System

If you have enabled multi-vsyst capability on a PA-7000 Series firewall with a Log Processing Card (LPC) installed, you can configure logging for different virtual systems as described in the following workflow.

STEP 1 | Create a Log Card subinterface.

1. Select **Network > Interfaces > Ethernet** and select the interface to be the Log Card interface.
2. Enter the **Interface Name**.
3. For **Interface Type**, select **Log Card**.
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. 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 | Commit your changes.

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](#).

Configure a PA-7000 Series LFC for Logging per Virtual System

If you have enabled multi-vsyst capability on a PA-7000 Series firewall with a Log Forwarding Card (LFC) installed, you can configure logging for different virtual systems as described in the following workflow.

STEP 1 | Create a Log Forwarding Card subinterface.

1. Select **Device** > **Log Forwarding Card** and add a subinterface.
2. For **Interface Name**, after the period, enter the subinterface assigned to the tenant's virtual system.
3. (Optional) Enter a **Comment**.
4. For **Tag**, enter a VLAN tag value.



Make the tag the same as the subinterface number for ease of use, but it can be a different number.

5. On the **Config** tab, in the **Assign Interface to Virtual System** field, select the virtual system to which the LFC subinterface is assigned. Alternatively, you can click **Virtual Systems** to add a new virtual system.
6. Click **OK**.

STEP 2 | (Optional) Enter the addresses assigned to the subinterface, and configure the default gateway.

1. Select the **Network** 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 3 | Commit your changes.

Click **OK** and **Commit**.

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.

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.

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 | Commit the configuration.

Click **Commit**.

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 [Virtual Wire Interfaces](#).
- If you have configured User-ID and multiple virtual systems, you can share user mappings across virtual systems. See [Share User-ID Mappings Across Virtual Systems](#).

Zone Protection and DoS Protection

Segmenting the network into functional and organizational zones reduces the network's attack surface—the portion of the network exposed to potential attackers. Zone protection defends network zones against flood attacks, reconnaissance attempts, packet-based attacks, and attacks that use non-IP protocols. Tailor a Zone Protection profile to protect each zone (you can apply the same profile to similar zones). Denial-of-service (DoS) protection defends specific critical systems against flood attacks, especially devices that user access from the internet such as web servers and database servers, and protects resources from session floods. Tailor DoS Protection profiles and policy rules to protect each set of critical devices. Visit the Best Practices documentation portal to get a checklist of Zone Protection and DoS Protection best practices.



Check and monitor firewall dataplane CPU consumption to ensure that each firewall is properly sized to support DoS and Zone Protection along with any other features that consume CPU cycles, such as decryption. If you use Panorama to manage your firewalls, use Device Monitor (Panorama > Managed Devices > Health) to check and monitor the CPU consumption of all managed firewalls at one time.

- > Network Segmentation Using Zones
- > How Do Zones Protect the Network?
- > Zone Defense
- > Configure Zone Protection to Increase Network Security
- > DoS Protection Against Flooding of New Sessions

Network Segmentation Using Zones

The larger the network, the more difficult it is to protect. A large, unsegmented network presents a large attack surface that can be difficult to manage and protect. Because traffic and applications have access to the entire network, once an attacker gains entry to a network, the attacker can move laterally through the network to access critical data. A large network is also more difficult to monitor and control. Segmenting the network limits an attacker's ability to move through the network by preventing lateral movement between zones.

A security zone is a group of one or more physical or virtual firewall interfaces and the network segments connected to the zone's interfaces. You control protection for each zone individually so that each zone receives the specific protections it needs. For example, a zone for the finance department may not need to allow all of the applications that a zone for IT allows.

To fully protect your network, all traffic must flow through the firewall. [Configure Interfaces and Zones](#) to create separate zones for different functional areas such as the internet gateway, sensitive data storage, and business applications, and for different organizational groups such as finance, IT, marketing, and engineering. Wherever there is a logical division of functionality, application usage, or user access privileges, you can create a separate zone to isolate and protect the area and apply the appropriate security policy rules to prevent unnecessary access to data and applications that only one or some groups need to access. The more granular the zones, the greater the visibility and control you have over network traffic. Dividing your network into zones helps to create a [Zero Trust architecture](#) that executes a security philosophy of trusting no users, devices, applications, or packets, and verifying everything. The end goal is to create a network that allows access only to the users, devices, and applications that have legitimate business needs, and to deny all other traffic.

How to appropriately restrict and permit access to zones depends on the network environment. For example, environments such as semiconductor manufacturing floors or robotic assembly plants, where the workstations control sensitive manufacturing equipment, or highly restricted access areas, may require physical segmentation that permits no access from outside devices (no mobile device access).

In environments where users can access the network with mobile devices, enabling [User-ID](#) and [App-ID](#) in conjunction with segmenting the network into zones ensures that users receive the appropriate access privileges regardless of where they access the network, because access privileges are tied to a user or a user group instead of to a device in one particular zone.

The protection requirements for different functional areas and groups may also differ. For example, a zone that handles a large amount of traffic may require different flood protection thresholds than a zone that normally handles less traffic. The ability to define the appropriate protection for each zone is another reason to segment the network. What appropriate protection is depends on your network architecture, what you want to protect, and what traffic you want to permit and deny.

How Do Zones Protect the Network?

Zones not only protect your network by segmenting it into smaller, more easily managed areas, zones also protect the network because you can control access to zones and traffic movement between zones.

Zones prevent uncontrolled traffic from flowing through the firewall interfaces into your network because firewall interfaces can't process traffic until you assign them to zones. The firewall applies zone protection on ingress interfaces, where traffic enters the firewall in the direction of flow from the originating client to the responding server (c2s), to filter traffic before it enters a zone.

The firewall interface type and the zone type (Tap, virtual wire, L2, L3, Tunnel, or External) must match, which helps to protect the network against admitting traffic that doesn't belong in a zone. For example, you can assign an L2 interface to an L2 zone or an L3 interface to an L3 zone, but you can't assign an L2 interface to an L3 zone.

In addition, a firewall interface can belong to one zone only. Traffic destined for different zones can't use the same interface, which helps to prevent inappropriate traffic from entering a zone and enables you to configure the protection appropriate for each individual zone. You can connect more than one firewall interface to a zone to increase bandwidth, but each interface can connect to only one zone.

After the firewall admits traffic to a zone, traffic flows freely within that zone and is not logged. The more [granular you make each zone](#), the greater the control you have over the traffic that accesses each zone, and the more difficult it is for malware to move laterally across the network between zones. Traffic can't flow between zones unless a security policy rule allows it and the zones are of the same zone type (Tap, virtual wire, L2, L3, Tunnel, or External). For example, a security policy rule can allow traffic between two L3 zones, but not between an L3 zone and an L2 zone. The firewall logs traffic that flows between zones when a security policy rule permits interzone traffic.

By default, security policy rules prevent lateral movement of traffic between zones, so malware can't gain access to one zone and then move freely through the network to other targets.



Tunnel zones are for non-encrypted tunnels. You can apply different security policy rules to the tunnel content and to the zone of the outer tunnel, as described in the [Tunnel Content Inspection Overview](#).

Zone Defense

Zone Protection profiles defend zones against flood, reconnaissance, packet-based, and non-IP-protocol-based attacks. DoS Protection profiles used in DoS Protection policy rules defend specific, critical devices against targeted flood and resource-based attacks. A DoS attack overloads the network or targeted critical systems with large amounts of unwanted traffic an attempt to disrupt network services.

Plan to defend your network against different types of DoS attacks:

- **Application-Based Attacks**—Target weaknesses in a particular application and try to exhaust its resources so legitimate users can't use it. An example of this is the [Slowloris](#) attack.
- **Protocol-Based Attacks**—Also known as state-exhaustion attacks, these attacks target protocol weaknesses. A common example is a [SYN flood attack](#).
- **Volumetric Attacks**—High-volume attacks that attempt to overwhelm the available network resources, especially bandwidth, and bring down the target to prevent legitimate users from accessing those resources. An example of this is a [UDP flood attack](#).

There are no default Zone Protection profiles or DoS Protection profiles and DoS Protection policy rules. Configure and apply zone protection based on each zone's traffic characteristics and configure DoS protection based on the individual critical systems you want to protect in each zone.

- [Zone Defense Tools](#)
- [How Do the Zone Defense Tools Work?](#)
- [Firewall Placement for DoS Protection](#)
- [Zone Protection Profiles](#)
- [Packet Buffer Protection](#)
- [DoS Protection Profiles and Policy Rules](#)

Zone Defense Tools

Effective defense against DoS attacks requires a layered approach. The first layer of defense should be a dedicated, high-volume DDoS protection device at the internet-facing network perimeter and a perimeter router, switch, or other hardware-based packet drop device with appropriate access control lists (ACLs) to defend against volumetric attacks that the session-based firewall isn't designed to handle. The firewall adds more granular layers of DoS attack defense and also visibility into application traffic that dedicated DDoS devices don't provide.

Palo Alto Networks firewalls provide four complementary tools to layer in DoS protection for your network zones and critical devices:

- [Zone Protection profiles](#) defend the ingress zone edge against IP flood attacks, reconnaissance port scans and host sweeps, IP packet-based attacks, and non-IP protocol attacks. The ingress zone is where traffic enters the firewall in the direction of flow from the client to the server (c2s), where the client is the originator of the flow and the server is the responder. Zone Protection profiles provide a second layer of broad defense against DoS attacks, based on the aggregate traffic entering the zone, by limiting the new connections-per-second (CPS) to the zone. Zone Protection profiles don't take individual devices (IP addresses) into account because the profiles apply to the aggregate traffic entering the zone.

Zone protection profiles defend the network as a session is formed, before the firewall performs DoS Protection policy and Security policy rule lookups, and consume fewer CPU cycles than a DoS Protection policy or Security policy rule lookup. If a Zone Protection profile denies traffic, the firewall doesn't spend CPU cycles on policy rule lookups.

Apply Zone Protection profiles to every zone, both internet-facing and internal.

- [DoS Protection profiles and policy rules](#) defend specific individual endpoints and resources against flood attacks, especially high-value targets that users access from the internet. While a Zone Protection profile

defends the zone from flood attacks, a DoS Protection policy rule with an appropriate DoS Protection profile defends critical individual systems in a zone from targeted flood attacks, providing a granular third layer of defense against DoS attacks.



Because the intent of DoS protection is to defend critical devices and because it consumes resources, DoS protection defends only the devices you specify in a DoS Protection policy rule. No other devices are protected.

DoS Protection profiles set flood protection thresholds (new CPS limits) for individual devices or groups of devices, resource protection thresholds (session limits for specified endpoints and resources), and whether the profile applies to [aggregate or classified](#) traffic. DoS Protection policy rules specify match criteria (source, destination, service ports), the action to take when traffic matches the rule, and the [aggregate and classified DoS Protection profiles](#) associated with each rule.

Aggregate DoS Protection policy rules apply the CPS thresholds defined in an aggregate DoS Protection profile to the combined traffic of all the devices that meet the DoS Protection policy rule match criteria. For example, if you configure the aggregate DoS Protection profile to limit the CPS rate to 20,000, the 20,000 CPS limit applies to the aggregate number of connections for the entire group. In this case, one device could receive the majority of the allowed connections.

Classified DoS Protection policy rules apply the CPS thresholds defined in a classified DoS Protection profile to each individual device that matches the policy rule. For example, if you configure the classified DoS Protection profile to limit the CPS rate to 4,000, then no device in the group can accept more than 4,000 CPS. A DoS Protection policy can have one aggregate profile and one classified profile.



Classified profiles can classify connections by source IP, destination IP, or both. For internet-facing zones, classify by destination IP only because the firewall can't scale to hold the internet routing table.

Apply DoS Protection only to critical devices, especially popular attack targets that users access from the internet, such as web servers and database servers.

- For existing sessions, [Packet Buffer Protection](#) protects the firewall (and therefore the zone) against single-session DoS attacks that attempt to overwhelm the firewall's packet buffer, using thresholds and timers to mitigate abusive sessions. You configure Packet Buffer Protection settings globally and apply them per zone.
- [Security Policy](#) rules affect both the ingress and egress flows of a session. To establish a session, incoming traffic must match an existing Security policy rule. If there is no match, the firewall discards the packet. A Security policy allows or denies traffic between zones (interzone) and within zones (intrazone) using criteria including zones, IP addresses, users, applications, services, and URL categories.



Apply the [best practice Vulnerability Protection profile](#) to each Security policy rule to help defend against DoS attacks.

The default Security policy rules don't permit traffic to travel between zones, so you need to configure a Security policy rule if you want to allow interzone traffic. All intrazone traffic is allowed by default. You can configure Security policy rules to match and control intrazone, interzone, or universal (intrazone and interzone) traffic.



Zone Protection profiles, DoS Protection profiles and policy rules, and Security policy rules only affect dataplane traffic on the firewall. Traffic originating on the firewall management interface does not cross the dataplane, so the firewall does not match management traffic against these profiles or policy rules.

- You can also search the [Palo Alto Networks Threat Vault](#) (requires a valid support account and login) for threats by hash, CVE, signature ID, domain name, URL, or IP address.

How Do the Zone Defense Tools Work?

When a packet arrives at the firewall, the firewall attempts to match the packet to an existing session, based on the ingress zone, egress zone, source IP address, destination IP address, protocol, and application derived from the packet header. If the firewall finds a match, then the packet uses the Security policy rules that already control the session. If the packet doesn't match an existing session, the firewall uses Zone Protection profiles, DoS Protection profiles and policy rules, and Security policy rules to determine whether to establish a session or discard the packet, and the level of access the packet receives.

After traffic passes through your dedicated DDoS device at the internet-facing network edge, the first protection the firewall applies is the broad defense of the Zone Protection profile, if one is attached to the zone. The firewall determines the zone from the interface on which the packet arrives (each interface is assigned to only one zone and all interfaces that carry traffic must belong to a zone). If the Zone Protection profile denies the packet, the firewall discards the packet and saves resources by not needing to look up the DoS Protection policy or Security policy. The firewall applies Zone Protection profiles only to new sessions (packets that do not match an existing session). After the firewall establishes a session, the firewall bypasses the Zone Protection profile lookup for succeeding packets in that session.

If the Zone Protection profile doesn't drop the packet, the second protection the firewall applies is a DoS Protection policy rule. If a Zone Protection profile allows a packet based on the total aggregate amount of traffic going to the zone, a DoS Protection policy rule may deny the packet if it is going to a particular destination or coming from a particular source that has exceeded the flood protection or resource protection settings in the rule's DoS Protection profile. If the packet matches a DoS Protection policy rule, the firewall applies the rule to the packet. If the rule denies access, the firewall discards the packet and doesn't perform a Security policy lookup. If the rule allows access, the firewall performs a Security policy lookup. Like the Zone Protection profile, the firewall enforces DoS Protection policy only on new sessions.

The third protection the firewall applies is a [Security policy](#) lookup, which happens only if the Zone Protection profile and DoS Protection policy rules allow the packet. If the firewall finds no Security policy rule match for the packet, the firewall discards the packet. If the firewall finds a matching Security policy rule, the firewall applies the rule to the packet. The firewall enforces the Security policy rule on traffic in both directions (c2s and s2c) for the life of the session. Apply the [best practice Vulnerability Protection profile](#) to all Security policy rules to help defend against DoS attacks.

The fourth protection the firewall applies is packet buffer protection, which you apply globally to protect the device and can also apply individually to zones to prevent single-session DoS attacks that attempt to overwhelm the firewall's packet buffer. For global protection, the firewall used Random Early Drop (RED) to drop packets (not sessions) when the level of traffic crosses protection thresholds. For per-zone protection, the firewall blocks the source IP address if it violates the packet buffer thresholds. Unlike zone and DoS protection, packet buffer protection applies to existing sessions.

Firewall Placement for DoS Protection

The firewall is a session-based device that isn't designed to scale to millions of connections-per-second (CPS) to defend against large volumetric DoS attacks. The firewall treats each unique flow (based on ingress and egress zone, source and destination IP, protocol, and application) as a session, spends CPU cycles on packet inspection at the port and the IP level to provide visibility into application traffic, and must count each session for the flood threshold counters, so firewall placement is critical to avoid flooding the firewall.

For the best DoS protection, *place firewalls as close to the resources you're protecting as possible*. This reduces the number of sessions the firewall needs to handle and therefore the amount of firewall resources required to provide DoS protection.

At the internet-facing perimeter, do *not* place firewalls you use for DoS protection or zone protection in front of dedicated DDoS devices and perimeter routers and switches. Make those high-volume devices your first line of DoS defense to mitigate volumetric flood attacks. For zone and DoS protection at the perimeter,

use high-capacity firewalls and place them *behind* the high-volume devices. As a rule, the closer a firewall is to the perimeter, the higher capacity it must be to handle the volume of traffic.

The way you segment your network into zones can help mitigate internal DoS attacks. Smaller zones provide greater visibility into traffic and prevent lateral movement of malware better because more traffic must cross zones, and to allow interzonal traffic requires you to create a specific Security policy rule (all intrazonal traffic is allowed by default). Consider revisiting your segmentation approach if your network is relatively unsegmented.

Baseline CPS Measurements for Setting Flood Thresholds

Flood protection thresholds determine the number of new connections-per-second (CPS) to allow for a zone (Zone Protection profile), for a group of devices in a zone (aggregate DoS Protection policy), or for individual devices in a zone (classified DoS Protection policy), when to throttle new connections to begin mitigating a potential flood attack, and when to drop all new connections. The default Zone Protection profile and DoS Protection profile flood protection thresholds aren't appropriate for most networks because each network is unique. You need to understand the aggregate normal and peak CPS for each zone to set effective Zone Protection profile thresholds, and for the individual critical systems you want to defend to set effective DoS Protection profile thresholds that don't inadvertently set thresholds too high and allow flood attacks or set thresholds too low and throttle traffic.

- [CPS Measurements to Take](#)
- [How to Measure CPS](#)

CPS Measurements to Take

Measure average and peak CPS traffic over the course of at least five business days or until you're confident that the measurements reflect the network's typical traffic patterns; the longer measurement period, the more accurate the measurements. Take into account special events, quarterly events, and annual events that may spike the number of CPS you need to support. You may need to adjust Zone Protection profiles and schedule adjusted DoS Protection policy rules to accommodate these types of events if your firewalls have the capacity to handle extra traffic. Take the following baseline measurements:

- For Zone Protection profiles, measure the average and peak CPS ingressing each zone.
- For aggregate DoS Protection profiles, measure the combined average and peak CPS for each group of devices you want to protect.
- For classified DoS Protection profiles, measure the average and peak CPS of the individual devices you want to protect.

Also understand the capacity of your firewalls and how other resource-consuming features such as decryption affect the number of connections each firewall can control. As a general rule, the closer a firewall is to the perimeter, the greater its capacity needs to be because it handles more traffic. The datasheet for each firewall model includes the total new sessions per second (CPS) the firewall supports and the [Firewall Comparison Tool](#) enables you to compare the CPS (and other metrics) of different firewall models.

How to Measure CPS

There are many ways to measure CPS:

- If you use Panorama to manage your firewalls, use [Device Monitoring](#) to measure CPS coming into a firewall (**Panorama > Managed Devices > Health > All Devices**). Device Monitoring can also show you a 90-day trend line of CPU average and peak use to help you understand the typical available capacity of each firewall.
- Run the operational CLI command **show session info**.



The operational CLI command **show counter interface** displays two times the actual CPS value. If you use this command, divide the CPS value by two to derive the real CPS value.

- For setting appropriate DoS Protection profile thresholds, work with application teams to understand the normal and peak CPS to their servers and the maximum CPS those servers can support.

In addition, you can filter firewall Traffic logs and Threat logs for the destination IP addresses of the critical devices you want to protect to obtain normal and peak session activity information.

- Use third-party tools such as Wireshark or NetFlow to collect and analyze network traffic.
- Use scripts to automate CPS information collection and continuous monitoring, and to mine information from the logs.
- Configure every Security policy rule on the firewall to **Log at Session End**. If you have no monitoring tools such as NetFlow or Wireshark, and cannot obtain or develop automated scripts, **Log at Session End** captures the number of connections at the session end. While this doesn't provide CPS information, it does give you an idea of the number of connections in a session.



To conserve resources, the firewall measures the aggregate CPS at ten-second intervals. For this reason, measurements you see on the firewall may not catch bursts within the ten-second interval. Although the average CPS measurements aren't affected, the peak CPS measurements may not be precise. For example, if the firewall logs report a 5,000 CPS average in a ten-second interval, it's possible that 4,000 CPS came in a one-second burst and the other 1,000 CPS were spread out over the remaining nine seconds.

To gather historical CPS data over time, if you use an SNMP server, you can use your own management tools to poll SNMP MIBs. However, it is important to understand that the CPS measurements in the MIBs show twice the actual CPS value (for example, if the true CPS measurement is 10,000, the MIBs show 20,000 as the value). You can still see trends from the MIBs and you can divide the CPS values by two to derive the true values. The SNMP MIB OIDs are: PanZoneActiveTcpCps, PanZoneActiveUdpCps, and PanZoneOtherIpCps. Because the firewall only takes measurements and updates the SNMP server every 10 seconds, poll every 10 seconds.

In addition, create separate [log forwarding profiles](#) for flood events so the appropriate administrator receives emails that contain only flood (potential DoS attack) events. Set Log Forwarding for both zone protection and DoS protection threshold events.



After you implement zone and DoS protection, use these methods to monitor the deployment, so as your network evolves and traffic patterns change, you adjust flood protection thresholds.

Zone Protection Profiles

Apply a Zone Protection profile to [each zone](#) to defend it based on the aggregate traffic entering the ingress zone.



In addition to configuring zone protection and DoS protection, apply the [best practice Vulnerability Protection profile](#) to each Security policy rule to help defend against DoS attacks.

- [Flood Protection](#)
- [Reconnaissance Protection](#)
- [Packet-Based Attack Protection](#)
- [Protocol Protection](#)

Flood Protection

A Zone Protection profile with flood protection configured defends an entire ingress zone against SYN, ICMP, ICMPv6, UDP, and other IP flood attacks. The firewall measures the aggregate amount of each flood type entering the zone in new connections-per-second (CPS) and compares the totals to the thresholds you configure in the Zone Protection profile. (You protect critical individual devices within a zone with [DoS Protection profiles and policy rules](#).)



Measure and monitor firewall dataplane CPU consumption to ensure that each firewall is properly sized to support DoS and Zone Protection and any other features that consume CPU cycles, such as decryption. If you use Panorama to manage your firewalls, [Device Monitoring](#) (Panorama > Managed Devices > Health > All Devices) shows you the CPU and memory consumption of each managed firewall. It can also show you a 90-day trend line of CPU average and peak use to help you understand the typical available capacity of each firewall.

For each flood type, you set three thresholds for new CPS entering the zone, and you can set a drop **Action** for SYN floods. If you know the baseline CPS rates for the zone, use these guidelines to set the initial thresholds, and then monitor and adjust the thresholds as necessary.

- **Alarm Rate**—The new CPS threshold to trigger an alarm. Target setting the **Alarm Rate** to 15-20% above the average CPS rate for the zone so that normal fluctuations don't cause alerts.
- **Activate**—The new CPS threshold to activate the flood protection mechanism and begin dropping new connections. For ICMP, ICMPv6, UDP, and other IP floods, the protection mechanism is Random Early Drop (RED, also known as Random Early Detection). For SYN floods only, you can set the drop **Action** to SYN Cookies or RED. Target setting the **Activate** rate to just above the peak CPS rate for the zone to begin mitigating potential floods.
- **Maximum**—The number of connections-per-second to drop incoming packets when RED is the protection mechanism. Target setting the **Maximum** rate to approximately 80-90% of firewall capacity, taking into account other features that consume firewall resources.

If you don't know the baseline CPS rates for the zone, start by setting the **Maximum** CPS rate to approximately 80-90% of firewall capacity and use it to derive reasonable flood mitigation alarm and activation rates. Set the **Alarm Rate** and **Activate** rate based on the Maximum rate. For example, you could set the **Alarm Rate** to half the **Maximum** rate and adjust it depending on how many alarms you receive and the firewall resources being consumed. Be careful setting the **Activate Rate** since it begins to drop connections. Because normal traffic loads experience some fluctuation, it's best not to drop connections too aggressively. . Err on the high side and adjust the rate if firewall resources are impacted.



SYN Flood Protection is the only type for which you set the drop Action. Start by setting the Action to SYN Cookies. SYN Cookies treats legitimate traffic fairly and only drops traffic that fails the SYN handshake, while using Random Early Drop drops traffic randomly, so RED may affect legitimate traffic. However, SYN Cookies is more resource-intensive because the firewall acts as a proxy for the target server and handles the three-way handshake for the server. The tradeoff is not dropping legitimate traffic (SYN Cookies) versus preserving firewall resources (RED). Monitor the firewall, and if SYN Cookies consumes too many resources, switch to RED. If you don't have a dedicated DDoS prevention device in front of the firewall, always use RED as the drop mechanism.

The default threshold values are high so that activating a Zone Protection profile doesn't unexpectedly drop legitimate traffic. Adjust the thresholds to values appropriate for your network's traffic. The best method for understanding how to set reasonable flood thresholds is to take baseline measurements of average and peak CPS for each flood type to determine the normal traffic conditions for each zone and to understand the capacity of the firewall, including the impact of other resource-consuming features such as decryption. Monitor and adjust the flood thresholds as needed and as your network evolves.



Firewalls with multiple dataplane processors (DPs) distribute connections across DPs. In general, the firewall divides the CPS threshold settings equally across its DPs. For example, if a firewall has five DPs and you set the Alarm Rate to 20,000 CPS, each DP has an Alarm Rate of 4,000 CPS ($20,000 / 5 = 4,000$), so if the new sessions on a DP exceeds 4,000, it triggers the Alarm Rate threshold for that DP.

Reconnaissance Protection

Similar to the military definition of reconnaissance, the network security definition of reconnaissance is when attackers attempt to gain information about your network's vulnerabilities by secretly probing the network to find weaknesses. Reconnaissance activities are often preludes to a network attack. *Enable Reconnaissance Protection on all zones* to defend against port scans and host sweeps:

- **Port scans** discover open ports on a network. A port scanning tool sends client requests to a range of port numbers on a host, with the goal of locating an active port to exploit in an attack. Zone Protection profiles defend against TCP and UDP port scans.
- **Host sweeps** examine multiple hosts to determine if a specific port is open and vulnerable.

You can use reconnaissance tools for legitimate purposes such as pen testing of network security or the strength of a firewall. You can specify up to 20 IP addresses or netmask address objects to exclude from Reconnaissance Protection so that your internal IT department can conduct pen tests to find and fix network vulnerabilities.

You can set the action to take when reconnaissance traffic (excluding pen testing traffic) exceeds the configured threshold when you [Configure Reconnaissance Protection](#). Retain the default **Interval** and **Threshold** to log a few packets for analysis before blocking the reconnaissance operation.

Packet-Based Attack Protection

Packet-based attacks take many forms. Zone Protection profiles check IP, TCP, ICMP, IPv6, and ICMPv6 packet headers and protect a zone by:

- Dropping packets with undesirable characteristics.
- Stripping undesirable options from packets before admitting them to the zone.

Select the drop characteristics for each packet type when you [Configure Packet Based Attack Protection](#). The best practices for each IP protocol are:

- **IP Drop**—Drop **Unknown** and **Malformed** packets. Also drop **Strict Source Routing** and **Loose Source Routing** because allowing these options allows adversaries to bypass Security policy rules that use the Destination IP address as the matching criteria. For internal zones only, check **Spoofed IP Address** so only traffic with a source address that matches the firewall routing table can access the zone.
- **TCP Drop**—Retain the default **TCP SYN with Data** and **TCP SYNACK with Data** drops, drop **Mismatched overlapping TCP segment** and **Split Handshake** packets, and strip the **TCP Timestamp** from packets.



Enabling Rematch Sessions (Device > Setup > Session > Session Settings) is a best practice that applies committed newly configured or edited Security Policy rules to existing sessions. However, if you [configure Tunnel Content Inspection](#) on a zone and Rematch Sessions is enabled, you must also disable Reject Non-SYN TCP (change the selection from Global to No), or else when you enable or edit a Tunnel Content Inspection policy, the firewall drops all existing tunnel sessions. Create a separate Zone Protection profile to disable Reject Non-SYN TCP only on zones that have Tunnel Content Inspection policies and only when you enable Rematch Sessions.

- **ICMP Drop**—There are no standard best practice settings because dropping ICMP packets depends on how you use ICMP (or if you use ICMP). For example, if you want to block ping activity, you can block **ICMP Ping ID 0**.

- **IPv6 Drop**—If compliance matters, ensure that the firewall drops packets with non-compliant routing headers, extensions, etc.
- **ICMPv6 Drop**—If compliance matters, ensure that the firewall drops certain packets if the packets don't match a Security policy rule.

Protocol Protection

In a Zone Protection profile, Protocol Protection defends against non-IP protocol based attacks. Enable Protocol Protection to block or allow non-IP protocols between security zones on a Layer 2 VLAN or on a virtual wire, or between interfaces within a single zone on a Layer 2 VLAN (Layer 3 interfaces and zones drop non-IP protocols so non-IP Protocol Protection doesn't apply). [Configure Protocol Protection](#) to reduce security risks and facilitate regulatory compliance by preventing less secure protocols from entering a zone, or an interface in a zone.



If you don't configure a Zone Protection profile that prevents non-IP protocols in the same zone from going from one Layer 2 interface to another, the firewall allows the traffic because of the default intrazone allow Security policy rule. You can create a Zone Protection profile that [blocks protocols such as LLDP](#) within a zone to prevent discovery of networks reachable through other zone interfaces.

If you need to discover which non-IP protocols are running on your network, use monitoring tools such as NetFlow, Wireshark, or other third-party tools discover non-IP protocols on your network. Examples of non-IP protocols you can block or allow are LLDP, NetBEUI, Spanning Tree, and Supervisory Control and Data Acquisition (SCADA) systems such as Generic Object Oriented Substation Event (GOOSE), among many others.

Create an **Exclude List** or an **Include List** to configure Protocol Protection for a zone. The **Exclude List** is a block list—the firewall blocks all of the protocols you place in the **Exclude List** and allows all other protocols. The **Include List** is an allow list—the firewall allows only the protocols you specify in the list and blocks all other protocols.



Use include lists for Protocol Protection instead of exclude lists. Include lists specifically sanction only the protocols you want to allow and block the protocols you don't need or didn't know were on your network, which reduces the attack surface and blocks unknown traffic.

A list supports up to 64 Ethertype entries, each identified by its [IEEE hexadecimal Ethertype](#) code. Other sources of Ethertype codes are standards.ieee.org/develop/regauth/ethertype/eth.txt and <http://www.cavebear.com/archive/cavebear/Ethernet/type.html>. When you configure zone protection for non-IP protocols on zones that have Aggregated Ethernet (AE) interfaces, you can't block or allow a non-IP protocol on only one AE interface member because AE interface members are treated as a group.



Protocol Protection doesn't allow blocking IPv4 (Ethertype 0x0800), IPv6 (0x86DD), ARP (0x0806), or VLAN-tagged frames (0x8100). The firewall always implicitly allows these four Ethernets in an Include List even if you don't explicitly list them and doesn't permit you to add them to an Exclude List.

Packet Buffer Protection

Packet Buffer Protection defends your firewall and network from single session DoS attacks that can overwhelm the firewall's packet buffer and cause legitimate traffic to drop. Although you don't configure Packet Buffer Protection in a Zone Protection profile or in a DoS Protection profile or policy rule, Packet Buffer Protection defends ingress zones. While zone and DoS protection apply to new sessions (connections) and are granular, Packet Buffer Protection applies to existing sessions and is global.



VM-Series firewalls do not support packet buffer protection.

You [Configure Packet Buffer Protection](#) globally to protect the entire firewall and also enable Packet Buffer Protection on each zone to protect zones:

- **Global Packet Buffer Protection**—The firewall monitors sessions from all zones (regardless of whether Packet Buffer Protection is enabled in a zone) and how those sessions utilize the packet buffer. You must configure Packet Buffer Protection globally (**Device** > **Setup** > **Session Settings**) to protect the firewall and to enable it on individual zones. When packet buffer consumption reaches the configured **Activate** percentage, the firewall used Random Early Drop (RED) to drop packets from the offending sessions (the firewall doesn't drop complete sessions at the global level).
- **Per-Zone Packet Buffer Protection**—Enable Packet Buffer Protection on each zone (**Network** > **Zones**) to layer in a second level of protection. When packet buffer consumption crosses the **Activate** threshold and global protection begins to apply RED to session traffic, that starts the **Block Hold Time** timer. The **Block Hold Time** is the amount of time in seconds that the offending session can continue before the firewall blocks the entire session. The offending session remains blocked until the **Block Duration** time expires.



If you don't enable Packet Buffer Protection globally, it won't be active in zones until you enable it globally.

Take baseline measurements of firewall packet buffer utilization over a period of time—at least one business week, but a longer measurement period provides a better baseline—to understand typical usage. One method of taking baseline measurements is to use a script to automate packet buffer monitoring, which also provides a good method for continuous monitoring and helps you understand when anomalous events occur. Your Palo Alto Networks account team can provide a sample script you can use to develop your own script.

If baseline measurements consistently show abnormally high packet buffer utilization, then the firewall's capacity may be undersized for typical traffic loads. In this case, consider resizing the firewall deployment. Otherwise, you need to tune the Packet Buffer Protection thresholds carefully to prevent impacted buffers from overflowing (and to prevent dropping legitimate traffic). When firewall sizing is correct for the deployment, only an attack should cause a large spike in buffer usage.



Overrunning the firewall packet buffer negatively impacts the firewall's packet forwarding capabilities. When the buffers are full, no packets can enter the firewall on any interface, not just the interface that experienced the attack.

The best practices for setting the thresholds are:

- **Alert and Activate**—Start with the default threshold values (50% in both cases), monitor packet buffer utilization, and adjust the thresholds as necessary. If the firewall is sized correctly, buffer utilization should be well below 50%. If the packet buffer utilization crosses **Alert** threshold, the firewall creates an alert entry in the System log.
- **Block Hold Time**—When packet buffer utilization triggers the **Activate** threshold, the **Block Hold Time** sets the amount of time the offending session can continue before the firewall blocks the session. During the **Block Hold Time**, the firewall continues to apply RED to the packets of offending sessions. Start with the default **Block Hold Time** threshold value (60 seconds), monitor packet buffer utilization, and adjust the threshold as necessary. If the packet buffer utilization percentage falls below the **Activate** threshold before the **Block Hold Time** expires, the timer resets and doesn't start until the **Activate** threshold is crossed again. Increasing the **Block Hold Time** imposes a greater penalty on offending sessions and reducing it imposes a lesser penalty on offending sessions.
- **Block Duration**—When the **Block Hold Time** expires, the firewall blocks the offending session for the period of time defined by the **Block Duration**. Start with the default threshold value (3600 seconds), monitor packet buffer utilization, and adjust the threshold as necessary. When you enable Packet Buffer

Protection on a zone, **Block Duration** affects every session from the IP address even if only one session from an IP address overutilizes the packet buffer. If you believe that blocking an IP address for one hour (3600 seconds) is too great a penalty, reduce the **Block Duration** to an acceptable value.

In addition to monitoring the buffer utilization of individual sessions, Packet Buffer Protection can also block an IP address if certain criteria are met. While the firewall monitors the packet buffers, if it detects a source IP address rapidly creating sessions that would not individually be seen as an attack, it blocks that IP address for the configured **Block Duration**.



Network Address Translation (NAT) (an external source that has translated its internet-bound traffic using source NAT) can give the appearance of greater packet buffer utilization because of IP address translation activity. If this occurs, adjust the thresholds in a way that penalizes individual sessions but doesn't penalize the underlying IP addresses (so other sessions from the same IP address aren't affected). To do this, reduce the **Block Hold Time** so the firewall blocks individual sessions that overutilize the buffers faster, and reduce the **Block Duration** so that the underlying IP address is not unduly penalized.

DoS Protection Profiles and Policy Rules

DoS Protection profiles and DoS Protection policy rules combine to protect specific groups of critical resources and individual critical resources against session floods. Compared to Zone Protection profiles, which protect entire zones from flood attacks, DoS protection provides granular defense for specific systems, especially critical systems that users access from the internet and are often attack targets, such as web servers and database servers. Apply both types of protection because if you only apply a Zone Protection profile, then a DoS attack that targets a particular system in the zone can succeed if the total connections-per-second (CPS) doesn't exceed the zone's **Activate** and **Maximum** rates.

DoS Protection is resource-intensive, so use it only for critical systems. Similar to Zone Protection profiles, DoS Protection profiles specify flood thresholds. DoS Protection policy rules determine the devices, users, zones, and services to which DoS Profiles apply.



*In addition to configuring DoS protection and zone protection, apply the **best practice Vulnerability Protection** profile to each Security policy rule to help defend against DoS attacks.*

- [Classified Versus Aggregate DoS Protection](#)
- [DoS Protection Profiles](#)
- [DoS Protection Policy Rules](#)

Classified Versus Aggregate DoS Protection

You can configure *aggregate* and *classified* [DoS Protection Profiles](#), and apply one profile or one of each type of profile to [DoS Protection Policy Rules](#) when you [configure DoS Protection](#).

- **Aggregate**—Sets thresholds that apply to the entire group of devices specified in a DoS Protection policy rule instead of to each individual device, so one device could receive the majority of the allowed connection traffic. For example, a **Max Rate** of 20,000 CPS means the total CPS for the group is 20,000, and an individual device can receive up to 20,000 CPS if other devices don't have connections. Aggregate DoS Protection policies provide another layer of broad protection (after your dedicated DDoS device at the internet perimeter and Zone Protection profiles) for a particular group of critical devices when you want to apply extra constraints on specific subnets, users, or services.
- **Classified**—Sets flood thresholds that apply to each individual device specified in a DoS Protection policy rule. For example, if you set an **Max Rate** of 5,000 CPS, each device specified in the rule can accept up to 5,000 CPS before it drops new connections. If you apply a classified DoS Protection policy rule to more than one device, the devices governed by the rule should be similar in terms of capacity and

how you want to control their CPS rates because classified thresholds apply to each individual device. Classified profiles protect individual critical resources.

When you configure a DoS Protection policy rule with a classified DoS Protection profile (**Option/Protection > Classified > Address**), use the **Address** field to specify whether incoming connections count toward the profile thresholds based on matching the **source-ip-only**, **destination-ip-only**, or **src-dest-ip-both** (the firewall counts both the source and the destination IP addresses matches toward the thresholds). Counters consume resources, so the way you count address matches affects firewall resource consumption. You can use classified DoS protection to:

- Protect critical individual devices, especially servers that users access from the internet and are often attack targets, such as web servers, database servers, and DNS servers. Set appropriate flood and resource protection thresholds in a classified DoS Protection profile. Create a DoS Protection policy rule that applies the profile to each server's IP address by adding the IP addresses as the rule's destination criteria, and set the **Address** to **destination-ip-only**.



Do not use source-IP-only or src-dest-ip-both classification for internet-facing zones in classified DoS Protection policy rules because the firewall doesn't have the capacity to store counters for every possible IP address on the internet. Increment the threshold counter for source IPs only for internal zone or same-zone rules. In perimeter zones, use destination-ip-only.

- Monitor the CPS rate for a suspect host or group of hosts (the zone that contains the hosts cannot be internet-facing). Set an appropriate alarm threshold in a classified DoS Protection profile to notify you if a host initiates an unusually large number of connections. Create a DoS Protection policy rule that applies the profile to the individual source or source address group and set the **Address** to **source-ip-only**. Investigate hosts that initiate enough new connections to set off the alarm.

How you configure the **Address** (**source-ip-only**, **destination-ip-only**, or **src-dest-ip-both**) for classified profiles depends on your DoS protection goals, what you are protecting, and whether the protected device(s) are in internet-facing zones.



The firewall uses more resources to track src-dest-ip-both as the Address than to track source-IP-only or destination-ip-only because the counters consume resources for both the source and destination IP addresses instead of just one of the two.

If you apply both an aggregate and a classified DoS Protection profile to the same DoS Protection policy rule, the firewall applies the aggregate profile first and then applies the classified profile if needed. For example, we protect a group of five web servers with both types of profiles in a DoS Protection policy rule. The aggregate profile configuration drops new connections when the combined total for the group reaches a **Max Rate** of 25,000 CPS. The classified profile configuration drops new connections to any individual web server in the group when it reaches a **Max Rate** of 6,000 CPS. There are three scenarios where new connection traffic crosses **Max Rate** thresholds:

- The new CPS rate exceeds the aggregate **Max Rate** but doesn't exceed the classified **Max Rate**. In this scenario, the firewall applies the aggregate profile and blocks all new connections for the configured Block Duration.
- The new CPS rate doesn't exceed the aggregate **Max Rate**, but the CPS to one of the web servers exceeds the classified **Max Rate**. In this scenario, the firewall checks the aggregate profile and finds that the rate for the group is less than 25,000 CPS, so the firewall doesn't block new connections based on that. Next, the firewall checks the classified profile and finds that the rate for a particular server exceeds 6,000 CPS. The firewall applies the classified profile and blocks new connections to that particular server for the configured Block Duration. Because the other servers in the group are within the classified profile's **Max Rate**, their traffic is not affected.
- The new CPS rate exceeds the aggregate **Max Rate** and also exceeds the classified **Max Rate** for one of the web servers. In this scenario, the firewall checks the aggregate profile and finds that the rate for the group exceeds 25,000 CPS, so the firewall blocks new connections to limit the group's total CPS. The

firewall then checks the classified profile and finds that the rate for a particular server exceeds 6,000 CPS (so the aggregate profile enforced the group's combined limit, but that wasn't enough to protect this particular server). The firewall applies the classified profile and blocks new connections to that particular server for the configured Block Duration. Because the other servers in the group are within the classified profile's **Max Rate**, their traffic is not affected.



If you want both an aggregate and a classified DoS Protection profile to apply to the same traffic, you must apply both profiles to the same DoS Protection policy rule. If you apply the aggregate profile to one rule and the classified profile to a different rule, even if they specify exactly the same traffic, the firewall can apply only one profile because when the traffic matches the first DoS Protection policy rule, the firewall executes the Action specified in that rule and doesn't compare to the traffic to any subsequent rules, so the traffic never matches the second rule and the firewall can't apply its action. (This is the same way that Security policy rules work.)

DoS Protection Profiles

DoS Protection profiles set thresholds that [protect against new session IP flood attacks](#) and provide resource protection (maximum concurrent session limits for specified endpoints and resources). DoS Protection profiles protect specific devices (classified profiles) and groups of devices (aggregate profiles) against SYN, UDP, ICMP, ICMPv6, and Other IP flood attacks. Configuring Flood Protection thresholds in a DoS Protection profile is similar to configuring [Flood Protection](#) in a Zone Protection profile, but Zone Protection profiles protect entire ingress zones, while DoS protection profiles and policy rules are granular and targeted, and can even be classified to a single device (IP address). The firewall measures the aggregate number of connections-per-second (CPS) to a group of devices (aggregate profile) or measures the CPS to individual devices (classified profile).



Measure and monitor firewall dataplane CPU consumption to ensure that each firewall is properly sized to support DoS and Zone Protection and any other features that consume CPU cycles, such as decryption. If you use Panorama to manage your firewalls, [Device Monitoring](#) (Panorama > Managed Devices > Health > All Devices) shows you the CPU and memory consumption of each managed firewall. It can also show you a 90-day trend line of CPU average and peak use to help you understand the typical available capacity of each firewall.

For each flood type, you set three thresholds for new CPS to a group of devices (aggregate) or to individual devices (classified) and a **Block Duration**, and you can set a drop **Action** for SYN floods:

- **Alarm Rate**—When new CPS exceeds this threshold, the firewall generates a DoS alarm. For classified profiles, set the rate to 15-20% above the device's average CPS rate so that normal fluctuations don't cause alerts. For aggregate profiles, set the rate to 15-20% above the group's average CPS rate.
- **Activate Rate**—When new CPS exceeds this threshold, the firewall begins to drop new connections to mitigate the flood until the CPS rate drops below the threshold. For classified profiles, the **Max Rate** should be an acceptable CPS rate for the device(s) you're protecting (the **Max Rate** won't flood the critical device(s)). You can set the **Activate Rate** to the same threshold as the **Max Rate** so that the firewall doesn't use RED or SYN Cookies to begin dropping traffic before it reaches the **Max Rate**. Set the **Activate Rate** lower than the **Max Rate** only if you want to drop traffic before it reaches the **Max Rate**. For aggregate profiles, set the threshold just above the average peak CPS rate for the group to begin mitigating floods using RED (or SYN Cookies for SYN floods).
- **Max Rate**—When new CPS exceeds this threshold, the firewall blocks (drops) all new connections from the offending IP address for the specified **Block Duration** time period. For classified profiles, base the **Max Rate** threshold on the capacity of the device(s) you're protecting so that the CPS rate can't flood them. For aggregate profiles, set to 80-90% of the group's capacity.
- **Block Duration**—When new CPS exceeds the **Max Rate**, the firewall blocks new connections from the offending IP address. The **Block Duration** specifies the amount of time the firewall continues to block

the IP address's new connections. While the firewall blocks new connections, it doesn't count incoming connections and doesn't increment the threshold counters. For classified and aggregate profiles, use the default value (300 seconds) to block the attacking session without penalizing legitimate sessions from the source for too long a period of time.



SYN Flood Protection is the only type for which you set the drop Action. Start by setting the Action to SYN Cookies. SYN Cookies treats legitimate traffic fairly and only drops traffic that fails the SYN handshake, while using Random Early Drop drops traffic randomly, so RED may affect legitimate traffic. However, SYN Cookies is more resource-intensive because the firewall acts as a proxy for the target server and handles the three-way handshake for the server. The tradeoff is not dropping legitimate traffic (SYN Cookies) versus preserving firewall resources (RED). Monitor the firewall, and if SYN Cookies consumes too many resources, switch to RED. If you don't have a dedicated DDoS prevention device in front of the firewall, always use RED as the drop mechanism.

The default threshold values are high so that DoS Protection profiles don't unexpectedly drop legitimate traffic. Monitor connection traffic and adjust the thresholds to values appropriate for your network. Start by taking baseline measurements of average and peak CPS for each flood type to determine the normal traffic conditions for the critical devices you want to protect. Because normal traffic loads experience some fluctuation, it's best not to drop connections too aggressively. Monitor and adjust the flood thresholds as needed and as your network evolves.

Another method of setting flood thresholds is to use the baseline measurements to set the maximum CPS you want to allow and work back from there to derive reasonable flood mitigation alarm and activation rates.



Firewalls with multiple dataplane processors (DPs) distribute connections across DPs. In general, the firewall divides the CPS threshold settings equally across its DPs. For example, if a firewall has five DPs and you set the Alarm Rate to 20,000 CPS, each DP has an Alarm Rate of 4,000 CPS ($20,000 / 5 = 4,000$), so if the new sessions on a DP exceeds 4,000, it triggers the Alarm Rate threshold for that DP.

In addition to setting IP flood thresholds, you can also use DoS Protection profiles to detect and prevent session exhaustion attacks in which a large number of hosts (bots) establish as many sessions as possible to consume a target's resources. On the profile's **Resources Protection** tab, you can set the maximum number of concurrent sessions that the device(s) defined in the DoS Protection policy rule to which you apply the profile can receive. When the number of concurrent sessions reaches its maximum limit, new sessions are dropped.

The maximum number of concurrent sessions to set depends on your network context. Understand the number of concurrent sessions that the resources you are protecting (defined in the DoS Protection policy rule to which you attach the profile) can handle. Set the threshold to approximately 80% of the resources' capacity, then monitor and adjust the threshold as needed.

For aggregate profiles, the **Resources Protection** threshold applies to all traffic of the devices defined in the policy rule (source and destination). For classified profiles, the **Resources Protection** threshold applies to the traffic based on whether the classified policy rule applies to the source IP only, to the destination IP only, or to both the source and destination IPs.

DoS Protection Policy Rules

DoS Protection policy rules control the systems to which the firewall applies DoS protection (the flood thresholds configured in DoS Protection profiles that you attach to DoS Protection policy rules), what action to take when traffic matches the criteria defined in the rule, and how to log DoS traffic. Because DoS protection consumes firewall resources, use it only to defend specific critical resources against session floods, especially common targets that users access from the internet, such as web servers and database servers. Use Zone Protection profiles to protect entire zones against floods and other attacks. DoS

Protection policy rules provide granular matching criteria so that you have the flexibility to define exactly what you want to protect:

- Source zone, interface, IP address (including whole regions), and user.
- Destination zone, interface, and IP address (including whole regions).
- Services (by port and protocol). DoS protection applies only to the services you specify. However, specifying services doesn't allow the services and implicitly block all other services. Specifying services limits DoS protection to those services, but doesn't block other services.



In addition to protecting service ports in use on critical servers, you can also protect against DoS attacks on the unused service ports of critical servers. For critical systems, you can do this by creating one DoS Protection policy rule and profile to protect ports with services running, and a different DoS Protection policy rule and profile to protect ports with no services running. For example, you can protect a web server's normal service ports, such as 80 and 443, with one policy/profile, and protect all of the other service ports with the other policy/profile. Be aware of the firewall's capacity so that servicing the DoS counters doesn't impact performance.

When traffic matches a DoS Protection policy rule, the firewall takes one of three actions:

- **Deny**—The firewall denies access and doesn't apply a DoS Protection profile. Traffic that matches the rule is blocked.
- **Allow**—The firewall permits access and doesn't apply a DoS Protection profile. Traffic that matches the rule is allowed.
- **Protect**—The firewall protects the devices defined in the DoS Protection policy rule by applying the specified DoS Protection profile or profiles thresholds to traffic that matches the rule. A rule can have one aggregate DoS Protection profile and one classified DoS Protection profile, and for classified profiles, you can use the source IP, destination IP, or both to increment the flood threshold counters, as described in [Classified Versus Aggregate DoS Protection](#). Incoming packets count against both DoS Protection profile thresholds if they match the rule.

The firewall applies DoS Protection profiles only if the **Action** is **Protect**. If the DoS Protection policy rule's **Action** is **Protect**, specify the appropriate aggregate and/or classified DoS Protection profiles in the rule so that the firewall applies the DoS Protection profile's thresholds to traffic that matches the rule. Most rules are **Protect** rules.

The **Allow** and **Deny** actions enable you to make exceptions within larger groups but do not apply DoS protection to the traffic. For example, you can deny the traffic from most of a group but allow a subset of that traffic. Conversely, you can allow the traffic from most of a group and deny a subset of that traffic.

You can **Schedule** when a DoS Protection policy rule is active (start and end time, recurrence period). One use case for scheduling is to apply different flood thresholds at different times of the day or week. For example, if your business experiences significantly less traffic at night than during the day, you may want to apply higher flood thresholds during the day than at night. Another use case is to schedule special thresholds for special events, providing that the firewall supports the CPS rates.

For easier management and granular reporting, configure **Log Forwarding** to separate DoS protection logs from other threat logs. Forward DoS threshold violation events directly to the administrators via email in addition to forwarding the logs to a server such as an SNMP or syslog server. Providing that the firewalls are appropriately sized, threshold breaches should not be frequent and will be strong indicators of an attack attempt.

Configure Zone Protection to Increase Network Security

The following topics provide zone protection configuration examples:

- [Configure Reconnaissance Protection](#)
- [Configure Packet Based Attack Protection](#)
- [Configure Protocol Protection](#)
- [Configure Packet Buffer Protection](#)

Configure Reconnaissance Protection

Configure one of the following [Reconnaissance Protection](#) actions for the firewall to take in response to the corresponding reconnaissance attempt:

- **Allow**—The firewall allows the port scan or host sweep reconnaissance to continue.
- **Alert**—The firewall generates an alert for each port scan or host sweep that matches the configured threshold within the specified time interval. Alert is the default action.
- **Block**—The firewall drops all subsequent packets from the source to the destination for the remainder of the specified time interval.
- **Block IP**—The firewall drops all subsequent packets for the specified **Duration**, in seconds (the range is 1-3,600). **Track By** determines whether the firewall blocks source or source-and-destination traffic.

STEP 1 | Configure Reconnaissance Protection.

1. Select **Network > Network Profiles > Zone Protection**.
2. Select a Zone Protection profile or **Add** a new profile and enter a **Name** for it.
3. On the Reconnaissance Protection tab, select the scan types to protect against.
4. Select an **Action** for each scan. If you select Block IP, you must also configure **Track By** (source or source-and-destination) and **Duration**.
5. Set the **Interval** in seconds. This options defines the time interval for port scan and host sweep detection.
6. Set the **Threshold**. The threshold defines the number of port scan events or host sweeps that occurs within the interval configured above that triggers an action.

STEP 2 | (Optional) Configure a Source Address Exclusion.

1. On the Reconnaissance Protection tab, **Add** a Source Address Exclusion.
 1. Enter a descriptive **Name** for the allowed address.
 2. Set the Address Type to **IPv4** or **IPv6** and then select an address object or enter an IP address.
 3. Click **OK**.
2. Click **OK** to save the Zone Protection profile.
3. **Commit** your changes.

Configure Packet Based Attack Protection

To enhance security for a zone, [Packet-Based Attack Protection](#) allows you to specify whether the firewall drops IP, IPv6, TCP, ICMP, or ICMPv6 packets that have certain characteristics or strips certain options from the packets.

For example, you can drop TCP SYN and SYN-ACK packets that contain data in the payload during a TCP three-way handshake. A Zone Protection profile by default is set to drop SYN and SYN-ACK packets with data (you must apply the profile to the zone).

The [TCP Fast Open](#) option ([RFC 7413](#)) preserves the speed of a connection setup by including data in the payload of SYN and SYN-ACK packets. A Zone Protection profile treats handshakes that use the TCP Fast Open option separately from other SYN and SYN-ACK packets; the profile by default is set to allow the handshake packets if they contain a valid Fast Open cookie.



If you have existing Zone Protection profiles in place when you upgrade to PAN-OS 8.0, the three default settings will apply to each profile and the firewall will act accordingly.

Beginning with PAN-OS 8.1.2 and later releases, you can use a CLI command (Step 4 in this task) to enable the firewall to generate a Threat log when the firewall receives and drops the following types of packets, so that you can more easily analyze these occurrences and also fulfill audit and compliance requirements:

- Teardrop attack
- DoS attack using ping of death

Furthermore, the same CLI command also enables the firewall to generate Threat logs for the following types of packets if you enable the corresponding Packet Based Attack Protection:

- Fragmented IP packets
- IP address spoofing
- ICMP packets larger than 1024 bytes
- Packets containing ICMP fragments
- ICMP packets embedded with an error message
- First packets for a TCP session that are not SYN packets

STEP 1 | Create a Zone Protection profile and configure Packet-Based Attack Protection settings.

1. Select **Network > Network Profiles > Zone Protection** and **Add** a new profile.
2. Enter a **Name** for the profile and an optional **Description**.
3. Select **Packet Based Attack Protection**.
4. On each tab (**IP Drop**, **TCP Drop**, **ICMP Drop**, **IPv6 Drop**, and **ICMPv6 Drop**), select the [Packet-Based Attack Protection settings](#) you want to enforce to protect a zone.
5. Click **OK**.

STEP 2 | Apply the Zone Protection profile to a security zone that is assigned to interfaces you want to protect.

1. Select **Network > Zones** and select the zone where you want to assign the Zone Protection profile.
2. **Add** the **Interfaces** belonging to the zone.
3. For **Zone Protection Profile**, select the profile you just created.
4. Click **OK**.

STEP 3 | **Commit** your changes.

STEP 4 | ([PAN-OS 8.1.2 and later releases](#)) Enable the firewall to generate Threat logs for a teardrop attack and a DoS attack using ping of death, and also generate Threat logs for the types of packets listed above if you enable the corresponding packet-based attack protection (in Step 1). For example, if you enable packet-based attack protection for **Spoofed IP address**, using the following CLI causes the firewall to generate a Threat log when the firewall receives and drops a packet with a spoofed IP address.

1. [Access the CLI](#).
2. Use the operational CLI command **set systemsetting additional-threat-log on**. Default is **off**.


Configure Protocol Protection

Protect virtual wire or Layer 2 security zones from non-IP protocol packets by using [Protocol Protection](#).

- [Use Case: Non-IP Protocol Protection Between Security Zones on Layer 2 Interfaces](#)
- [Use Case: Non-IP Protocol Protection Within a Security Zone on Layer 2 Interfaces](#)

Use Case: Non-IP Protocol Protection Between Security Zones on Layer 2 Interfaces

In this use case, the firewall is in a Layer 2 VLAN divided into two subinterfaces. VLAN 100 is 192.168.100.1/24, subinterface .6. VLAN 200 is 192.168.100.1/24, subinterface .7. Non-IP protocol protection applies to ingress zones. In this use case, if the Internet zone is the ingress zone, the firewall blocks the Generic Object Oriented Substation Event (GOOSE) protocol. If the User zone is the ingress zone, the firewall allows the GOOSE protocol. The firewall implicitly allows IPv4, IPv6, ARP, and VLAN-tagged frames in both zones.



STEP 1 | Configure two VLAN subinterfaces.

1. Select **Network > Interfaces > VLAN** and **Add** an interface.
2. **Interface Name** defaults to vlan. After the period, enter 7.
3. On the **Config** tab, **Assign Interface To** the **VLAN 200**.
4. Click **OK**.
5. Select **Network > Interfaces > VLAN** and **Add** an interface.
6. **Interface Name** defaults to vlan. After the period, enter 6.
7. On the **Config** tab, **Assign Interface To** the **VLAN 100**.
8. Click **OK**.

STEP 2 | Configure protocol protection in a Zone Protection profile to block GOOSE protocol packets.

1. Select **Network > Network Profiles > Zone Protection** and **Add** a profile.
2. Enter the **Name** Block GOOSE.
3. Select **Protocol Protection**.
4. Choose **Rule Type** of **Exclude List**.
5. Enter the **Protocol Name**, GOOSE, to easily identify the Ethertype on the list. The firewall doesn't verify that the name you enter matches the Ethertype code; it uses only the Ethertype code to filter.
6. Enter **Ethertype** code 0x88B8. The Ethertype must be preceded by 0x to indicate a hexadecimal value. Range is 0x0000 to 0xFFFF.
7. Select **Enable** to enforce the protocol protection. You can disable a protocol on the list, for example, for testing.
8. Click **OK**.

STEP 3 | Apply the Zone Protection profile to the Internet zone.

1. Select **Network > Zones** and **Add** a zone.
2. Enter the **Name** of the zone, Internet.
3. For **Location**, select the virtual system where the zone applies.
4. For **Type**, select **Layer2**.
5. **Add** the **Interface** that belongs to the zone, vlan.7.
6. For **Zone Protection Profile**, select the profile Block GOOSE.
7. Click **OK**.

STEP 4 | Configure protocol protection to allow GOOSE protocol packets.

Create another Zone protection profile named Allow GOOSE, and choose **Rule Type** of **Include List**.



When configuring an Include list, include all required non-IP protocols; an incomplete list can result in legitimate non-IP traffic being blocked.

STEP 5 | Apply the Zone Protection profile to the User zone.

1. Select **Network > Zones** and **Add** a zone.
2. Enter the **Name** of the zone, User.
3. For **Location**, select the virtual system where the zone applies.
4. For **Type**, select **Layer2**.
5. **Add** the **Interface** that belongs to the zone, vlan.6.
6. For **Zone Protection Profile**, select the profile Allow GOOSE.
7. Click **OK**.

STEP 6 | Commit.

Click **Commit**.

STEP 7 | View the number of non-IP packets the firewall has dropped based on protocol protection.

[Access the CLI.](#)


```
> show counter global name pkt_nonip_pkt_drop
> show counter global name pkt_nonip_pkt_drop delta yes
```

Use Case: Non-IP Protocol Protection Within a Security Zone on Layer 2 Interfaces

If you don't implement a Zone Protection profile with non-IP protocol protection, the firewall allows non-IP protocols in a single zone to go from one Layer 2 interface to another. In this use case, blocking LLDP packets ensures that LLDP for one network doesn't discover a network reachable through another interface in the zone.

In the following figure, the Layer 2 VLAN named Datacenter is divided into two subinterfaces: 192.168.1.1/24, subinterface .7 and 192.168.1.2/24, subinterface .8. The VLAN belongs to the User zone. By applying a Zone Protection profile that blocks LLDP to the User zone:

- Subinterface .7 blocks LLDP from its switch to the firewall at the red X on the left, preventing that traffic from reaching subinterface .8.
- Subinterface .8 blocks LLDP from its switch to the firewall at the red X on the right, preventing that traffic from reaching subinterface .7.



STEP 1 | Create a subinterface for an Ethernet interface.

1. Select **Network > Interfaces > Ethernet** and select a Layer 2 interface, in this example, ethernet1/1.
2. Select **Add Subinterfaces**.
3. The **Interface Name** defaults to the interface (ethernet 1/1). After the period, enter 7.
4. For **Tag**, enter 300.
5. For **Security Zone**, select User.
6. Click **OK**.

STEP 2 | Create a second subinterface for the Ethernet interface.

1. Select **Network > Interfaces > Ethernet** and select the Layer 2 interface: ethernet1/1.
2. Select **Add Subinterfaces**.
3. The **Interface Name** defaults to the interface (ethernet 1/1). After the period, enter 8.
4. For **Tag**, enter 400.
5. For **Security Zone**, select User.
6. Click **OK**.

STEP 3 | Create a VLAN for the Layer2 interface and two subinterfaces.

1. Select **Network > VLANs** and **Add** a VLAN.

2. Enter the **Name** of the VLAN; for this example, enter Datacenter.
3. For **VLAN Interface**, select **None**.
4. For **Interfaces**, click **Add** and select the Layer 2 interface: ethernet1/1, and two subinterfaces: ethernet1/1.7 and ethernet1/1.8.
5. Click **OK**.

STEP 4 | Block non-IP protocol packets in a Zone Protection profile.

1. Select **Network > Network Profiles > Zone Protection** and **Add** a profile.
2. Enter the **Name**, in this example, Block LLDP.
3. Enter a profile **Description**—Block LLDP packets from an LLDP network to other interfaces in the zone (intrazone).
4. Select **Protocol Protection**.
5. Choose **Rule Type** of **Exclude List**.
6. Enter **Protocol Name** LLDP.
7. Enter **Ethertype** code 0x88cc. The Ethertype must be preceded by 0x to indicate a hexadecimal value.
8. Select **Enable**.
9. Click **OK**.

STEP 5 | Apply the Zone Protection profile to the security zone to which Layer 2 VLAN belongs.

1. Select **Network > Zones**.
2. **Add** a zone.
3. Enter the **Name** of the zone, User.
4. For **Location**, select the virtual system where the zone applies.
5. For **Type**, select **Layer2**.
6. **Add** an **Interface** that belongs to the zone, ethernet1/1.7
7. **Add** an **Interface** that belongs to the zone, ethernet1/1.8.
8. For **Zone Protection Profile**, select the profile Block LLDP.
9. Click **OK**.

STEP 6 | Commit.

Click **Commit**.

STEP 7 | View the number of non-IP packets the firewall has dropped based on protocol protection.

[Access the CLI](#).

```
> show counter global name pkt_nonip_pkt_drop
> show counter global name pkt_nonip_pkt_drop delta yes
```

Configure Packet Buffer Protection

You configure [Packet Buffer Protection](#) settings globally and then apply them per ingress zone. When the firewall detects high buffer utilization, the firewall only monitors and takes action against sessions from zones with packet buffer protection enabled. Therefore, if the abusive session is from a zone without packet buffer protection, the high packet buffer utilization continues. Packet buffer protection can be applied to a zone but it is not active until global settings are configured and enabled.

STEP 1 | Configure the global session thresholds.

1. Select **Device > Setup > Session**.

-
2. **Edit** the Session Settings.
 3. Select the **Packet Buffer Protection** check box to enable and configure the packet buffer protection thresholds.
 4. Enter a value for each threshold and timer to define the packet buffer protection behavior.
 - **Alert (%)**—When packet buffer utilization exceeds this threshold for more than 10 seconds, the firewall creates a log event every minute. The firewall generates log events when packet buffer protection is enabled globally. The default threshold is 50% and the range is 0% to 99%. If the value is 0%, the firewall does not create a log event.
 - **Activate (%)**—When a packet buffer utilization exceeds this threshold, the firewall applies RED to abusive sessions. The default threshold is 50% and the range is 0% to 99%. If the value is 0%, the firewall does not apply RED.
-  *The firewall records alert events in the System log and events for dropped traffic, discarded sessions, and blocked IP address in the Threat log.*
- **Block Hold Time (sec)**—The amount of time a RED-mitigated session is allowed to continue before the firewall discards it. By default, the block hold time is 60 seconds. The range is 0 to 65,535 seconds. If the value is 0, the firewall does not discard sessions based on packet buffer protection.
 - **Block Duration (sec)**—This setting defines how long a session remains discarded or an IP address remains blocked. The default is 3,600 seconds with a range of 1 seconds to 15,999,999 seconds.
5. Click **OK**.
 6. **Commit** your changes.

STEP 2 | Enable packet buffer protection on an ingress zone.

1. Select **Network > Zones**.
2. Choose an ingress zone and click on its name.
3. Select the **Enable Packet Buffer Protection** check box in the Zone Protection section.
4. Click **OK**.
5. **Commit** your changes.

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 against DoS attacks of new sessions only, 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.


[DoS Protection Profiles and Policy Rules](#) work together to provide protection against flooding of many incoming SYN, UDP, ICMP, and ICMPv6 packets, and other types of IP packets. You determine what thresholds constitute flooding. In general, the DoS Protection profile sets the thresholds at which the firewall generates a DoS alarm, takes action such as Random Early Drop, and drops additional incoming connections. A DoS Protection policy rule that is set to protect (rather than to allow or deny packets) determines the criteria for packets to match (such as source address) in order to be counted toward the thresholds. This flexibility allows you to block certain traffic, or allow certain traffic and treat other traffic as DoS traffic. When the incoming rate exceeds your maximum threshold, the firewall blocks incoming traffic from the source address.

- [Multiple-Session DoS Attack](#)
- [Single-Session DoS Attack](#)
- [Configure DoS Protection Against Flooding of New Sessions](#)
- [End a Single Session DoS Attack](#)
- [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 Protection policy 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 Protection policy rule, new connections per second are counted toward the Alert, Activate, and Max Rate thresholds.

Sequence of Events as Firewall Quarantines an IP Address

	 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.



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 don't 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 (the address remains on the block list) until the Block Duration expires.



The 1-second checks illustrated in the preceding figure occur on firewall models that have multiple dataplane CPUs and a hardware network processor. All single dataplane systems or systems without a hardware network processor perform this mitigation in software and use a 5-second interval.

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 evaluation. You must configure a Security policy rule to allow or deny traffic 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 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 or drop rule in place. Hence, a single-session attack requires a Security policy deny or drop 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 from that source IP address, such as packets with a different application. 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. You can [Monitor Blocked IP Addresses](#) to view the block list, remove entries from it, and get additional information about an IP address on the block list.



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 Security policy enforcement. 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 [End a Single Session DoS Attack](#).

Configure DoS Protection Against Flooding of New Sessions

STEP 1 | 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. (Required for single-session attack mitigation or attacks that have not triggered the DoS Protection policy threshold; optional for multiple-session attack mitigation).



This step is one of the steps typically performed to stop an existing attack. See [End a Single Session DoS Attack](#).

- [Create a Security Policy Rule](#)

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 connections per second (cps) exceed the **Activate Rate** threshold:
 1. **Random Early Drop**—The firewall uses an algorithm to progressively start dropping that type of packet. If the attack continues, the higher the incoming cps rate (above the **Activate Rate**) gets, the more packets the firewall drops. The firewall drops packets until the incoming cps rate reaches the **Max Rate**, at which point the firewall drops all incoming connections. **Random Early Drop (RED)** is the default action for **SYN Flood**, and the only action for **UDP Flood**, **ICMP Flood**, **ICMPv6 Flood**, and **Other IP Flood**. RED is more efficient than SYN Cookies and can handle larger attacks, but doesn't discern between good and bad traffic.
 2. **SYN Cookies**—Rather than immediately sending the SYN to the server, the firewall generates a cookie (on behalf of the server) to send in the SYN-ACK to the client. The client responds with its ACK and the cookie; upon this validation the firewall then sends the SYN to the server. The **SYN Cookies** action requires more firewall resources than **Random Early Drop**; it's more discerning because it affects bad traffic.
5. (Optional) On each of the flood tabs, change the following thresholds to suit your environment:
 - **Alarm Rate (connections/s)**—Specify the threshold rate (cps) above which a DoS alarm is generated. (Range is 0-2,000,000; default is 10,000.)
 - **Activate Rate (connections/s)**—Specify the threshold rate (cps) 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. (For SYN Flood, you can select the action that occurs.)

- **Max Rate (connections/s)**—Specify the threshold rate of incoming connections per second that the firewall allows. When the threshold is exceeded, new connections 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 identify 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 aren't part of an attack.

7. Click **OK**.

STEP 3 | Configure a DoS Protection policy rule that specifies the criteria for matching the incoming traffic.



The firewall resources are finite, so you wouldn't want to classify using source address on an internet-facing zone because there can be an enormous number of unique IP addresses that match the DoS Protection policy rule. That would require many counters and the firewall would run out of tracking resources. Instead, define a DoS Protection policy rule that classifies using the destination address (of the server you are protecting).

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). Choose zone or interface depending on your deployment and what you want to protect. For example, if you have only one interface coming into the firewall, choose Interface.
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 against DoS attacks only on the server that has a specific destination address, and you also want to ensure that every source IP address won't surpass a specific cps threshold to that server.

13. Click **OK**.

STEP 4 | Commit.

Click **Commit**.

End a Single Session DoS Attack

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.

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, use the 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 [Multiple-Session DoS Attack](#).

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 model (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.

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 SZONESRC      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 model only, 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 PA-5200 Series and PA-7000 Series models only, you can limit output to a dataplane. For example:

```
admin@PA-5260> show running resource-monitor ingress-backlogs dp dp1
```

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 [End a Single Session DoS Attack](#). At a minimum, you can [Configure DoS Protection Against Flooding of New Sessions](#).
- On a hardware model 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 `layer7processing` value indicates completed for sessions offloaded or enabled for sessions not offloaded.

```
admin@PA-5060> show session id 68088184

Session      68088184

c2s flow:
  source:      1.1.42.15 [trust]
  dst:         1.2.27.99
  proto:       6
  sport:       55993      dport:    6881
  state:       ACTIVE     type:     FLOW
  src user:    unknown
  dst user:    unknown
  offload:     Yes

s2c flow:
  source:      1.2.27.99 [untrust]
  dst:         1.1.42.15
  proto:       6
  sport:       6881      dport:    55993
  state:       ACTIVE     type:     FLOW
  src user:    unknown
  dst user:    unknown
  offload:     Yes

DP
index(local):      : 2
start time         : 979320
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.

STEP 1 | In the CLI, execute the following operational command on any hardware model:

```
admin@PA-7050> request session-discard [timeout <seconds>] [reason <reason-
string>] id <session-id>
```

The default timeout is 3,600 seconds.

STEP 2 | Verify that sessions have been discarded.

```
admin@PA-7050> show session all filter state discard
```


Certifications

The following topics describe how to configure Palo Alto Networks® firewalls and appliances 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.

For details about product certifications and third-party validation, refer to the [Certifications page](#).

- > [Enable FIPS and Common Criteria Support](#)
- > [FIPS-CC Security Functions](#)
- > [Scrub Swap Memory on a Firewall or Appliances in FIPS-CC Mode](#)

Enable FIPS and Common Criteria Support

Use the following procedures 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.

FIPS-CC mode is supported on all Palo Alto Networks next-generation firewalls and appliances—including VM-Series firewalls. To enable FIPS-CC mode, first boot the firewall into the Maintenance Recovery Tool (MRT) and then change the operational mode from normal mode to FIPS-CC mode. The procedure to change the operational mode is the same for all firewalls and appliances but the procedure to access the MRT varies.



When you enable FIPS-CC mode, the firewall will reset to the factory default settings; all configuration will be removed.

- [Access the Maintenance Recovery Tool \(MRT\)](#)
- [Change the Operational Mode to FIPS-CC Mode](#)

Access the Maintenance Recovery Tool (MRT)

The Maintenance Recovery Tool (MRT) enables you to perform several tasks on Palo Alto Networks firewalls and appliances. For example, you can revert the firewall or appliance to factory default settings, revert PAN-OS or a content update to a previous version, run diagnostics on the file system, gather system information, and extract logs. Additionally, you can use the MRT to [Change the Operational Mode to FIPS-CC Mode](#) or from FIPS-CC mode to normal mode.

The following procedures describe how to access the Maintenance Recovery Tool (MRT) on various Palo Alto Networks products.

- Access the MRT on hardware firewalls and appliances (such as PA-220 firewalls, PA-7000 Series firewalls, or M-Series appliances).
 1. Establish a serial console session to the firewall or appliance.
 1. Connect a serial cable from the serial port on your computer to the console port on the firewall or appliance.



If your computer does not have a 9-pin serial port but does have a USB port, use a serial-to-USB converter to establish the connection. If the firewall has a [micro USB console port](#), connect to the port using a standard Type-A USB to micro USB cable.

2. Open terminal emulation software on your computer and set to 9600-8-N-1 and then connect to the appropriate COM port.



On a Windows system, you can go to the Control Panel to view the COM port settings for Device and Printers to determine which COM port is assigned to the console.

3. Log in using an administrator account. (The default username/password is admin/admin.)
2. Enter the following CLI command and press **y** to confirm:

```
debug system maintenance-mode
```

3. After the firewall or appliance boots to the MRT welcome screen (in approximately 2 to 3 minutes), press Enter on **Continue** to access the MRT main menu.



You can also access the MRT by rebooting the firewall or appliance and entering **maint** at the maintenance mode prompt. A direct serial console connection is required.

After the firewall or appliance boots into the MRT, you can access the MRT remotely by establishing an SSH connection to the management (MGT) interface IP address. At the login prompt, enter **maint** as the username and the firewall or appliance serial number as the password.

- Access the MRT on VM-Series firewalls deployed in a private cloud (such as on a VMware ESXi or KVM hypervisor).
 1. Establish an SSH session to the management IP address of the firewall and log in using an administrator account.
 2. Enter the following CLI command and press **y** to confirm:

```
debug system maintenance-mode
```



It will take approximately 2 to 3 minutes for the firewall to boot to the MRT. During this time, your SSH session will disconnect.

3. After the firewall boots to the MRT welcome screen, log in based on the operational mode:
 - **Normal mode**—Establish an SSH session to the management IP address of the firewall and log in using **maint** as the username and the firewall or appliance serial number as the password.
 - **FIPS-CC mode**—Access the virtual machine management utility (such as the vSphere client) and connect to the virtual machine console.
 4. From the MRT welcome screen, press Enter on **Continue** to access the MRT main menu.
- Access the MRT on VM-Series firewalls deployed in the public cloud (such as AWS or Azure).
 1. Establish an SSH session to the management IP address of the firewall and log in using an administrator account.
 2. Enter the following CLI command and press **y** to confirm:

```
debug system maintenance-mode
```



It will take approximately 2 to 3 minutes for the firewall to boot to the MRT. During this time, your SSH session will disconnect.

3. After the firewall boots to the MRT welcome screen, log in based on the virtual machine type:
 - **AWS**—Log in as **ec2-user** and select the SSH public key associated with the virtual machine when you deployed it.
 - **Azure**—Enter the credentials you created when you deployed the VM-Series firewall.
 - **GCP**—Log in as **gcp-user** and select the SSH public key associated with the virtual machine when you deployed it.
4. From the MRT welcome screen, press Enter on **Continue** to access the MRT main menu.

Change the Operational Mode to FIPS-CC Mode

The following procedure describes how to change the operational mode of a Palo Alto Networks product from normal mode to FIPS-CC mode.

STEP 1 | Connect to the firewall or appliance and [Access the Maintenance Recovery Tool \(MRT\)](#).

STEP 2 | Select **Set FIPS-CC Mode** from the menu.

STEP 3 | Select **Enable FIPS-CC Mode**. The mode change operation starts and a status indicator shows progress. After the mode change is complete, the status shows **Success**.

STEP 4 | When prompted, select **Reboot**.



*If you change the operational mode on a VM-Series firewall deployed in the public cloud (AWS or Azure) and you lose your SSH connection to the MRT before you are able to **Reboot**, you must wait 10-15 minutes for the mode change to complete, log back into the MRT, and then reboot the firewall to complete the operation.*

After you switch to FIPS-CC mode, you see the following status: FIPS-CC mode enabled successfully.

In addition, the following changes are in effect:

- FIPS-CC displays at all times in the status bar at the bottom of the web interface.
- The default administrator login credentials change to admin/paloalto.

See [FIPS-CC Security Functions](#) for details on the security functions that are enforced in FIPS-CC mode.

FIPS-CC Security Functions

When FIPS-CC mode is enabled, the following security functions are enforced on all firewalls and appliances:

- ❑ To log in, the browser must be TLS 1.1 (or later) compatible; on a WF-500 appliance, you manage the appliance only through the CLI and you must connect using an SSHv2-compatible client application.
- ❑ All passwords must be at least six characters.
- ❑ You must ensure that **Failed Attempts** and **Lockout Time (min)** are 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 ensure that the **Idle Timeout** is greater than 0 in authentication settings. If a login session is idle for more than the specified time, the administrator is automatically logged out.
- ❑ The firewall or appliance 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—they are 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 2,048 bits (or more) or ECDSA 256 bits (or more); you must also use a digest of SHA256 or greater.



You cannot use a [hardware security module \(HSM\)](#) to store the private ECDSA keys used for [SSL Forward Proxy](#) or [SSL Inbound Inspection](#).

- ❑ Telnet, TFTP, and HTTP management connections are not available.
- ❑ You must enable encryption for the [HA1 control link](#). You must set automatic rekeying parameters; you must set the data parameter to a value no greater than 1000 MB (you cannot let it default) and you must set a time interval (you cannot leave it disabled).
- ❑ The serial console port in FIPS-CC mode functions as a limited status output port only; CLI access is not available.
- ❑ The serial console port on hardware and private-cloud VM-Series firewalls booted into the MRT provides interactive access to the MRT.
- ❑ Interactive console access is not supported in the hypervisor environment private-cloud VM-Series firewalls booted into the MRT; you can access the MRT only using SSH.

Scrub the Swap Memory on Firewalls or Appliances Running in FIPS-CC Mode

You should ensure that sensitive information is removed from the swap memory before you decommission a firewall or appliance (in FIPS-CC mode) or before you send it in for repair. Use this procedure to remove all cryptographic security parameter (CSP) information from swap partitions.



If you send a firewall that is managed by Panorama in for repair, see [Before Starting RMA Firewall Replacement](#).

STEP 1 | Open an SSH management session to the firewall or appliance.

STEP 2 | Run the following operational command:

```
request [restart | shutdown] system with-swap-scrub [dod | nnsa]
```

For example, to shut down the firewall or appliance and perform a Department of Defense (DoD) scrub, run the following command:

```
request shutdown system with-swap-scrub dod
```

STEP 3 | Press **y** at the warning prompt to start the scrub.

STEP 4 | Verify that the scrub completed successfully. View the **System** log and filter on the word **swap**. The **System** log indicates the scrub status for each swap partition (either one or two partitions depending on the model) and also displays a log entry that indicates the overall status of the scrub. If the scrub completed successfully on all swap partitions, the **System** log shows **Swap space scrub was successful**.

If the scrub failed on one or more swap partitions, the **System** log shows **Swap space scrub was unsuccessful**. The following screen capture shows the log results for a firewall that has two partitions.

06/08 10:24:02	general	medium	general	Swap space scrub was successful
06/08 10:24:02	general	medium	general	Scrub performed on swap space /opt/panlogs
06/08 10:24:02	general	medium	general	/secondary_swapfile
06/08 10:24:02	general	medium	general	Scrub performed on swap space /dev/sda7



To view the scrub logs using the CLI, run the **show log system | match swap** command.



If you initiate the scrub using the shutdown command, the firewall or appliance will power off after the scrub completes. Before you can power on the firewall or appliance, you must first disconnect and reconnect the power source.

2.1	Requisitos gerais.		
	Especificações Básicas:		
2.2.1.	Possível aquisição de solução de proteção de rede características de Próxima Geração (NGFW) para segurança de informação perimetral que inclui filtro de pacote, controle de aplicação, administração de largura de banda (QoS), VPN IPSec e SSL, IPS, prevenção contra ameaças de vírus, spywares e malwares "Zero Day", Filtro de URL, bem como controle de transmissão de dados e acesso a internet compondo uma plataforma de segurança integrada e robusta;	Conforme proposta.	
2.1.2.	Ser compatível com a ferramenta de gerenciamento centralizada de segurança de rede panorama.	Produto do mesmo fabricante conforme proposta.	
2.1.3.	Por plataforma de segurança entende-se hardware e software integrados do tipo appliance.	Conforme proposta.	
2.2.	Características gerais:		
2.2.1.	A plataforma de segurança deve possuir a capacidade e as características abaixo, por equipamento:		
2.2.1.1.	Throughput de 16 Gbps com a funcionalidade de controle de aplicação habilitada para todas as assinaturas que o fabricante possuir;	palo-alto-networks-product-summary-specsheet	1
2.2.1.2.	Throughput de 8 Gbps com as seguintes funcionalidade habilitadas simultaneamente para todas as assinaturas que a plataforma de segurança possuir devidamente ativadas e atuantes: controle de aplicação IPS, Antivírus e Antispyware. Caso o fabricante divulgue múltiplos números de desempenho para qualquer uma destas funcionalidades, somente o de menor valor será aceito;	palo-alto-networks-product-summary-specsheet	1
2.2.1.3.	Os throughputs devem ser comprovados por documento de domínio público do fabricante. A ausência de tais documentos comprobatórios reservará ao órgão o direito de aferir a performance dos equipamentos em bancada, assim como atendimento de todas as funcionalidades especificadas neste edital. Caso seja comprovado o não atendimento das especificações mínimas nos testes de bancada, serão considerados inabilitados e sujeitos as sanções previstas em lei;	palo-alto-networks-product-summary-specsheet	1
2.2.1.4.	Os documentos públicos devem comprovar os throughputs aferidos com tráfego HTTP ou blend de protocolos definidos pelo fabricante como tráfego real (real-word traffic blend);	palo-alto-networks-product-summary-specsheet	2
2.2.1.5.	Não será aceito aceleração de pacotes na placa de rede limitando a análise somente até camada 4.	Não se aplica.	
2.2.1.6.	Suporte a, no mínimo, 3.800.000 de conexões simultâneas;	palo-alto-networks-product-summary-specsheet	1
2.2.1.7.	Suporte a, no mínimo, 110.000 novas conexões HTTP por segundo;	palo-alto-networks-product-summary-specsheet	1
2.2.1.8.	Fonte 120/240 AC ou DC, redundante e hot-swappable;	palo-alto-networks-product-summary-specsheet	2
2.2.1.9.	Cooler hot-swappable;	palo-alto-networks-product-summary-specsheet	2
2.2.1.10.	Disco Solid State Drive (SSD) redundante de, no mínimo, 240 GB.	pa-5200-series	2

2.2.1.11.	Discos de, no mínimo, 2 TB em RAID 1 para armazenamento de logs interno ou externo a solução de firewall;	palo-alto-networks-product-summary-specsheet	2
2.2.1.12.	No mínimo, 04 (quatro) interfaces de rede 1 Gbps em portas cobre;	palo-alto-networks-product-summary-specsheet	1
2.2.1.13.	No mínimo, 08 (oito) interfaces de rede 1 Gbps SFP;	palo-alto-networks-product-summary-specsheet	1
2.2.1.14.	No mínimo, 08 (oito) interfaces de rede 10 Gbps SFP+;	palo-alto-networks-product-summary-specsheet	1
2.2.1.15.	No mínimo, 04 (quatro) interfaces de rede 40 Gbps QSFP+;	palo-alto-networks-product-summary-specsheet	1
2.2.1.16.	2 (duas) Gbps interfaces dedicadas para alta disponibilidade sendo pelo menos uma do tipo 40 Gbps QSFP+;	palo-alto-networks-product-summary-specsheet	1
2.2.1.17.	1 (uma) interface de rede 1 Gbps dedicada para gerenciamento;	pa-5000-series-ds	1
2.2.1.18.	1 (uma) interface do tipo console ou similar;	palo-alto-networks-product-summary-specsheet	1
2.2.1.19.	Suporte a, no mínimo, 60 (sessenta) zonas de segurança;	pan-os-new-features.pdf	141
2.2.1.20.	Estar licenciada para ou suportar sem o uso de licença, 10.000 (dez mil) clientes de VPN SSL simultâneos;	performance	5
2.2.1.21.	Estar licenciada para ou suportar sem o uso de licença, 3.000 (três mil) túneis de VPN IPSEC simultâneos;	palo-alto-networks-product-summary-specsheet	1
2.2.2.	Deve suportar, no mínimo, 10 sistemas virtuais lógicos (Contextos) no firewall Físico;	palo-alto-networks-product-summary-specsheet	1
2.2.3.	Os contextos virtuais devem suportar as funcionalidades nativas do gateway de proteção incluindo: Firewall, IPS, Antivírus, Anti-Spyware, Filtro de URL, Filtro de Dados VPN, Controle de Aplicações, QOS, NAT e Identificação de usuários;	pan-os	104 5, 106 2
2.2.4.	Por cada equipamento que compõe a plataforma de segurança, entende-se o hardware e as licenças de softwares necessárias para o seu funcionamento;	Conforme proposta.	
2.2.5.	Por console de gerência e monitoração, entende-se as licenças de software necessárias para as duas funcionalidades, bem como hardware dedicado para o funcionamento das mesmas;	Conforme proposta.	
2.2.6.	A console de gerência e monitoração podem residir no mesmo appliance de proteção de rede, desde que possuam recurso de CPU, interface de rede e sistema operacional dedicados para esta função;	Conforme proposta.	
2.2.7.	Nenhum dos modelos ofertados poderão estar listados no site do fabricante em listas de end-of-life e end-of-sale;	Conforme proposta.	
2.2.8.	A solução deve consistir de appliance de proteção de rede com funcionalidades de Next Generation Firewall (NGFW) e console de gerência e monitoração;	Conforme proposta.	
2.2.9.	Por funcionalidades de NGFW entende-se: reconhecimento de aplicações, prevenção de ameaças, identificação de usuários e controle granular de permissões;	palo-alto-networks-product-summary-specsheet	3
2.2.10.	As funcionalidades de proteção de rede que compõe a plataforma de segurança, podem funcionar em múltiplos appliances desde que obedeçam a todos os requisitos desta especificação;	nao se aplica	

2.2.11.	A plataforma deve ser otimizada para análise de conteúdo de aplicações em camada 7;	pan-os	263
2.2.12.	O hardware e software que executem as funcionalidades de proteção de rede, bem como a console de gerência e monitoração, devem ser do tipo appliance. Não serão aceitos equipamentos servidores e sistema operacional de uso genérico;	Conforme proposta.	
2.2.13.	Todos os equipamentos fornecidos com trilho para montagem em rack 19 incluindo o kit de adaptação se necessário e cabos de alimentação;	palo-alto-networks-product-summary-specsheet	1
2.2.14.	O software deverá ser fornecido em sua versão mais atualizada;	sera fornecida a ultima versão	
2.2.15.	Os dispositivos de proteção de rede devem possuir pelo menos as seguintes funcionalidades:		
2.2.15.1.	Suporte a 4094 VLAN Tags 802.1q;	pa-5200-series	2
2.2.15.2.	Agregação de links 802.3ad e LACP;	pa-5200-series	2
2.2.15.3.	Policy based routing ou policy based forwarding;	PAN-OS-9-Admin-Guide.pdf	122 1
2.2.15.4.	Roteamento multicast (PIM-SM);	pa-5200-series	2
2.2.15.5.	DHCP Relay;	pa-5200-series	810
2.2.15.6.	DHCP Server;	pa-5200-series	810
2.2.15.7.	Jumbo Frames;	pan-os	300
2.2.15.8.	Suporte a criação de objetos de rede que possam ser utilizados como endereço IP de interfaces L3;	pan-os	994
2.2.16.	Suportar sub-interfaces ethernet logicas.	pan-os	301
2.2.16.1.	Suporte a, no mínimo, 15 (quinze) roteadores virtuais na mesma instância de firewall;	performance	3
2.2.17.	O firewall deve ter a capacidade de testar o funcionamento de rotas estáticas e rota default com a definição de um endereço IP de destino que deve estar comunicável através de uma rota. Caso haja falha na comunicação o firewall deve ter a capacidade de usar rota alternativa para estabelecer a comunicação;	pan-os	829
2.2.17.1. 1.	Deve suportar os seguintes tipos de NAT:	pan-os	
2.2.17.1. 2.	Nat dinâmico (Many-to-1);	pan-os	900
2.2.17.1. 3.	Nat dinâmico (Many-to-Many);	pan-os	900
2.2.17.1. 4.	Nat estático (1-to-1);	pan-os	900
2.2.17.1. 5.	NAT estático (Many-to-Many);	pan-os	900
2.2.17.1. 6.	Nat estático bidirecional 1-to-1;	pan-os	900
2.2.17.1. 7.	Tradução de porta (PAT);	pan-os	900
2.2.17.1. 8.	NAT de Origem;	pan-os	900
2.2.17.1. 9.	NAT de Destino;	pan-os	900
2.2.17.1. 10.	Suportar NAT de Origem e NAT de Destino simultaneamente;	pan-os	900
2.2.17.1. 11.	Deve implementar Network Prefix Translation (NPTv6), prevenindo problemas de roteamento assimétrico;	pan-os	900
2.2.18.	Deve implementar o protocolo ECMP;	pan-os	937
2.2.19.	Deve implementar balanceamento de link por hash do IP de origem;	pan-os	937- 938

2.2.20.	Deve implementar balanceamento de link por hash do IP de origem e destino;	pan-os	937-938
2.2.21.	Deve implementar balanceamento de link através do método round-robin;	pan-os	939
2.2.22.	Deve implementar balanceamento de link por peso. Nesta opção deve ser possível definir o percentual de tráfego que será escoado por cada um dos links. Deve suportar o balanceamento de, no mínimo, quatro links;	pan-os	937-938
2.2.23.	Deve implementar balanceamento de link através de políticas por usuário e grupos de usuários do LDAP/AD;	pan-os	937-938
2.2.24.	Deve implementar balanceamento de link através de políticas por aplicação e porta de destino;	pan-os	939
2.2.25.	Deve implementar o protocolo Link Layer Discovery (LLDP), permitindo que o appliance e outros ativos da rede se comuniquem para identificação da topologia da rede em que estão conectados e a função dos mesmos facilitando o processo de troubleshooting. As informações aprendidas e armazenadas pelo appliance devem ser acessíveis via SNMP;	pan-os	942
2.2.26.	Enviar log para sistemas de monitoração externos, simultaneamente;	pan-os	382
2.2.27.	Deve haver a opção de enviar logs para os sistemas de monitoração externos via protocolo TCP e SSL;	pan-os	382
2.2.28.	Deve permitir configurar certificado caso necessário para autenticação no sistema de monitoração externo de logs;	pan-os	383
2.2.29.	Proteção contra anti-spoofing;	pan-os	802
2.2.30.	Deve permitir bloquear sessões TCP que usem variações do 3-way hand-shake, como 4 way e 5 way split hand-shake, prevenindo desta forma possíveis tráfegos maliciosos;	pan-os	973
2.2.31.	Deve permitir bloquear conexões que contenham dados no payload de pacotes TCP-SYN e SYN-ACK durante o three-way handshake;	pan-os	973
2.2.32.	Deve exibir nos logs de tráfego o motivo para o término da sessão no firewall, incluindo sessões finalizadas onde houver de-criptografia de SSL e SSH;	pan-os	351
2.2.33.	Para IPv4, deve suportar roteamento estático e dinâmico (RIPv2, BGP e OSPFv2);	pan-os	832, 843, 834
2.2.34.	Para IPv6, deve suportar roteamento estático e dinâmico (OSPFv3);	pan-os	838
2.2.35.	Suportar a OSPF graceful restart;	pan-os	840
2.2.36.	Deve suportar o protocolo MP-BGP (Multiprotocol BGP) permitindo que o firewall possa anunciar rotas multicast para IPv4 e rotas unicast para IPv6;	pan-os	843

2.2.37.	Suportar no mínimo as seguintes funcionalidades em IPv6: SLAAC (address auto configuration), NAT64, Identificação de usuários a partir do LDAP/AD, Captive Portal, IPv6 over IPv4 IPSec, Regras de proteção contra DoS (Denial of Service), De-criptografia SSL e SSH, PBF (Policy Based Forwarding), QoS, DHCPv6 Relay, IPSEC, VPN SSL, Ativo/Ativo, Ativo/Passivo, SNMP, NTP, SYSLOG, DNS, Neighbor Discovery (ND), Recursive DNS Server (RDNSS), DNS Search List (DNSSL) e controle de aplicação;	compatibility-matrix	105-108
2.2.38.	O dispositivos de proteção devem ter a capacidade de operar de forma simultânea em uma única instância de firewall, mediante o uso de suas interfaces físicas nos seguintes modos: Modo sniffer (monitoramento e análise do tráfego de rede), camada 2 (L2) e camada 3 (L3);	pan-os	799
2.2.39.	Modo Sniffer, para inspeção via porta espelhada do tráfego de dados da rede;	pan-os	799
2.2.40.	Modo Camada – 2 (L2), para inspeção de dados em linha e ter visibilidade e controle do tráfego em nível de aplicação;	pan-os	799
2.2.41.	Modo Camada – 3 (L3), para inspeção de dados em linha e ter visibilidade e controle do tráfego em nível de aplicação operando como default gateway das redes protegidas;	pan-os	809
2.2.42.	Modo misto de trabalho Sniffer, L2 e L3 em diferentes interfaces físicas;	pan-os	799
2.2.43.	Suporte a configuração de alta disponibilidade Ativo/Passivo e Ativo/Ativo:	pan-os	252
2.2.43.1.	Em modo transparente;	pan-os	252
2.2.43.2.	Em layer 3;	pan-os	252
2.2.44.	A configuração em alta disponibilidade deve sincronizar:		
2.2.44.1.	Sessões;	pan-os	302, 303
2.2.44.2.	Configurações, incluindo, mas não limitado a políticas de Firewall, NAT, QOS e objetos de rede;	pan-os	302, 303
2.2.44.3.	Certificados de-criptografados;	pan-os	302, 303
2.2.44.4.	Associações de Segurança das VPNs;	pan-os	302, 303
2.2.44.5.	Tabelas FIB;	pan-os	302, 303
2.2.44.6.	O HA (modo de Alta-Disponibilidade) deve possibilitar monitoração de falha de link.	pan-os	256
2.2.44.7.	As funcionalidades de controle de aplicações, VPN IPSec e SSL, QOS, SSL e SSH Decryption e protocolos de roteamento dinâmico devem operar em caráter permanente, podendo ser utilizadas por tempo indeterminado, mesmo que não subsista o direito de receber atualizações ou que não haja contrato de garantia de software com o fabricante.	Conforme licenciamento	
2.3.	Controle de Política por Firewall		
2.3.1.	Deverá suportar controles por zona de segurança.	pan-os	989
2.3.2.	Controles de políticas por porta e protocolo.	pan-os	990

2.3.3.	Controle de políticas por aplicações grupos estáticos de aplicações, grupos dinâmicos de aplicações (baseados em características e comportamento das aplicações) e categorias de aplicações.	pan-os	518
2.3.4.	Controle de políticas por usuários, grupos de usuários, IPs, redes e zonas de segurança.	pan-os	988-990
2.3.5.	Deve suportar a consulta a fontes externas de endereços IP, domínios e URLs podendo ser adicionados nas políticas de firewall para bloqueio ou permissão do tráfego.	pan-os	33
2.3.6.	Deve permitir autenticação segura através de certificado nas fontes externas de endereços IP, domínios e URLs;	pan-os	155
2.3.7.	Deve permitir consultar e criar exceção para objetos das listas externas a partir da interface de gerência do próprio firewall;	pan-os	161
2.3.8.	Controle de políticas por código de País (Por exemplo: BR, USA, UK, RUS).	pan-os	331
2.3.8.	Controle, inspeção e de-criptografia de SSL por política para tráfego de entrada (Inbound) e Saída (Outbound).	pan-os	589
2.3.9.	Deve suportar offload de certificado em inspeção de conexões SSL de entrada (Inbound);	pan-os	625
2.3.10.	Deve de-criptografar tráfego Inbound e Outbound em conexões negociadas com TLS 1.2;	pan-os	591
2.3.11.	Deve de-criptografar sites e aplicações que utilizam certificados ECC, incluindo Elliptical Curve Digital Signature Algorithm (ECDSA);	pan-os	604
2.3.12.	Controle de inspeção e de-criptografia de SSH por política;	pan-os	615
2.3.13.	A de-criptografia de SSH deve possibilitar a identificação e bloqueio de tráfego caso o protocolo esteja sendo usado para tunelar aplicações como técnica evasiva para burlar os controles de segurança;	pan-os	615
2.3.14.	A plataforma de segurança deve implementar espelhamento de tráfego de-criptografado (SSL e TLS) para soluções externas de análise (Forense de rede, DLP, Análise de Ameaças, entre outras);	pan-os	604
2.3.15.	É permitido uso de appliance externo, específico para a de-criptografia de (SSL e TLS), com espelhamento de cópia do tráfego de-criptografado tanto para o firewall, quanto para as soluções de análise.	nao se aplica	
2.3.16.	Bloqueios dos seguintes tipos de arquivos: bat, cab, dll, exe, pif, e reg	pan-os	998
2.3.17.	Traffic shaping QoS baseado em Políticas (Prioridade, Garantia e Máximo)	pan-os	705
2.3.18.	QoS baseado em políticas para marcação de pacotes (diffserv marking), inclusive por aplicações.	pan-os	705
2.3.19.	Suporte a objetos e regras IPV6.	pa-5200-series	2
2.3.20.	Suporte a objetos e regras multicast.	pa-5200-series	2
2.3.21.	Deve suportar no mínimo três tipos de negação de tráfego nas políticas de firewall: Drop sem notificação do bloqueio ao usuário, Drop com opção de envio de ICMP Unreachable para máquina de origem do tráfego, TCP-Reset para o client, TCP-Reset para o server ou para os dois lados da conexão;	pan-os	991

2.3.22.	Suportar a atribuição de agendamento as políticas com o objetivo de habilitar e desabilitar políticas em horários pré-definidos automaticamente.	pan-os	990
2.4.	Controle de Aplicações		
2.4.1.	Os dispositivos de proteção de rede deverão possuir a capacidade de reconhecer aplicações, independente de porta e protocolo, com as seguintes funcionalidades:	pan-os	511
2.4.1.1.	Deve ser possível a liberação e bloqueio somente de aplicações sem a necessidade de liberação de portas e protocolos.	techbrief-app-id	2
2.4.1.2.	Reconhecer pelo menos 1700 aplicações diferentes, incluindo, mas não limitado: a tráfego relacionado a peer-to-peer, redes sociais, acesso remoto, update de software, protocolos de rede, voip, áudio, vídeo, proxy, mensageiros instantâneos, compartilhamento de arquivos, e-mail;	site-1	1
2.4.1.3.	Reconhecer pelo menos as seguintes aplicações: bittorrent, gnutella, skype, facebook, linked-in, twitter, citrix, logmein, teamviewer, ms-rdp, vnc, gmail, youtube, http-proxy, http-tunnel, facebook chat, gmail chat, whatsapp, 4shared, dropbox, google drive, skydrive, db2, mysql, oracle, active directory, kerberos, ldap, radius, itunes, dhcp, ftp, dns, wins, msrpc, ntp, snmp, rpc over http, gotomeeting, webex, evernote, google-docs, etc;	https://aplipedia.paloaltonetworks.com	
2.4.1.4.	Deve inspecionar o payload de pacote de dados com o objetivo de detectar através de expressões regulares assinaturas de aplicações conhecidas pelo fabricante independente de porta e protocolo. A checagem de assinaturas também deve determinar se uma aplicação está utilizando a porta default ou não, incluindo, mas não limitado a RDP na porta 80 ao invés de 389;	techbrief-app-id	2
2.4.1.5.	Deve aplicar heurística a fim de detectar aplicações através de análise comportamental do tráfego observado, incluindo, mas não limitado a Encrypted Bittorrent e aplicações VOIP que utilizam criptografia proprietária;	techbrief-app-id	2
2.4.1.6.	Identificar o uso de táticas evasivas, ou seja, deve ter a capacidade de visualizar e controlar as aplicações e os ataques que utilizam táticas evasivas via comunicações criptografadas, tais como Skype e ataques mediante a porta 443.	techbrief-app-id	1
2.4.1.7.	Para tráfego criptografado SSL, deve de-criptografar pacotes a fim de possibilitar a leitura de payload para checagem de assinaturas de aplicações conhecidas pelo fabricante;	techbrief-app-id	2
2.4.1.8.	Deve realizar decodificação de protocolos com o objetivo de detectar aplicações encapsuladas dentro do protocolo e validar se o tráfego corresponde com a especificação do protocolo, incluindo, mas não limitado a Yahoo Instant Messenger usando HTTP. A decodificação de protocolo também deve identificar funcionalidades específicas dentro de uma aplicação, incluindo, mas não limitado a compartilhamento de arquivo dentro do Webex. Além de detectar arquivos e outros conteúdos que devem ser inspecionados de acordo as regras de segurança implementadas;	techbrief-app-id	3

2.4.1.9.	Deve permitir a utilização de aplicativos para um determinado grupo de usuário e bloquear para o restante, incluindo, mas não limitado a Skype. Deve permitir também a criação de políticas de exceção concedendo o acesso a aplicativos como Skype apenas para alguns usuários;	pan-os	988-990
2.4.1.10.	Deve permitir habilitar aplicações SAAS apenas no modo corporativo e bloquea-las para uso pessoal, tais como: office 365, Skype, aplicativos do Google e etc.	pan-os	988-990
2.4.1.11.	Identificar o uso de táticas evasivas via comunicações criptografadas;	techbrief-app-id	1
2.4.1.12.	Atualizar a base de assinaturas de aplicações automaticamente;	techbrief-app-id	4
2.4.1.13.	Reconhecer aplicações em IPv6;	compatibility-matrix	105
2.4.1.14.	Limitar a banda (download/upload) usada por aplicações (traffic shaping), baseado no IP de origem, usuários e grupos do LDAP/AD;	pan-os	705
2.4.1.15.	Os dispositivos de proteção de rede devem possuir a capacidade de identificar o usuário de rede com integração ao Microsoft Active Directory, sem a necessidade de instalação de agente no Domain Controller, nem nas estações dos usuários;	pan-os	447
2.4.1.16.	Deve ser possível adicionar controle de aplicações em todas as regras de segurança do dispositivo, ou seja, não se limitando somente a possibilidade de habilitar controle de aplicações em algumas regras;	pan-os	988-990
2.4.1.17.	Deve suportar múltiplos métodos de identificação e classificação das aplicações, por pelo menos checagem de assinaturas, decodificação de protocolos e análise heurística;	pan-os	447, 989
2.4.1.18.	Para manter a segurança da rede eficiente, deve suportar o controle sobre aplicações desconhecidas e não somente sobre aplicações conhecidas;	techbrief-app-id	2
2.4.1.19.	Permitir nativamente a criação de assinaturas personalizadas para reconhecimento de aplicações proprietárias na própria interface gráfica da solução, sem a necessidade de ação do fabricante, mantendo a confidencialidade das aplicações do órgão;	pan-os	519
2.4.1.20.	A criação de assinaturas personalizadas deve permitir o uso de expressões regulares, contexto (sessões ou transações), usando posição no payload dos pacotes TCP e UDP e usando decoders de pelo menos os seguintes protocolos:	pan-os	519
2.4.1.21.	HTTP, FTP, SMB, SMTP, Telnet, SSH, MS-SQL, IMAP, IMAP, MS-RPC, RTSP e File body.	pan-os	523-525
2.4.1.22.	O fabricante deve permitir a solicitação de inclusão de aplicações na base de assinaturas de aplicações;	https://aplipedia.paloaltonetworks.com	
2.4.1.23.	Deve alertar o usuário quando uma aplicação for bloqueada;	pan-os	381
2.4.1.24.	Deve possibilitar que o controle de portas seja aplicado para todas as aplicações;	pan-os	988-990
2.4.1.25.	Deve permitir criar filtro na tabela de regras de segurança para exibir somente:	techbrief-app-id	4

2.4.1.25.1.	Regras que permitem passagem de tráfego baseado na porta e não por aplicação, exibindo quais aplicações estão trafegando nas mesmas, o volume em bytes trafegado por cada aplicação por, pelo menos, os últimos 30 dias e o primeiro e último registro de log de cada aplicação trafegada por esta de- terminada regra;	techbrief-app-id	4
2.4.1.25.2.	Aplicações permitidas em regras de forma desnecessária, pois não há tráfego da mesma na determinada regra;	techbrief-app-id	4
2.4.1.25.3.	Regras de segurança onde não houve passagem de tráfego nos últimos 90 dias.	techbrief-app-id	4
2.4.1.26.	Deve possibilitar a diferenciação de tráfegos Peer2Peer (Bittorrent, emule, ne- onet, etc.) possuindo granularidade de controle/políticas para os mesmos;	techbrief-app-id	4
2.4.1.27.	Deve possibilitar a diferenciação de tráfegos de Instant Messaging (AIM, Gtalk, Facebook Chat, etc.) possuindo granularidade de controle/políticas para os mesmos;	techbrief-app-id	4
2.4.1.28.	Deve possibilitar a diferenciação e controle de partes das aplicações como por exemplo permitir o Gtalk chat e bloquear a transferência de arquivos;	techbrief-app-id	4
2.4.1.29.	Deve possibilitar a diferenciação de aplicações Proxies (ghostsurf, freegate, etc.) possuindo granularidade de controle/políticas para os mesmos;	techbrief-app-id	4
2.4.1.30.	Deve ser possível a criação de grupos estáticos de aplicações e grupos dinâmicos de aplicações baseados em características das aplicações como:	techbrief-app-id	4
2.4.1.30.1.	Tecnologia utilizada na aplicações (Client-Server, Browse Based, Network Protocol, etc).	pan-os	518-522
2.4.1.30.2.	Nível de risco da aplicação.	pan-os	518-522
2.4.1.30.3.	Categoria e sub-categoria de aplicações.	pan-os	518-522
2.4.1.30.4.	Aplicações que usem técnicas evasivas, utilizadas por malwares, como transferência de arquivos e/ou uso excessivo de banda, etc.	pan-os	518-522
2.5	Prevenção de ameaças		
2.5.1.	Para proteção do ambiente contra ataques, os dispositivos de proteção devem possuir módulo de IPS, Antivírus e Anti-Spyware integrados no próprio appliance de Firewall ou entregue através de composição com outro equipamento ou fabricante.	PAN-OS-6.0-Admin-Guide.pdf	pag. 291
2.5.2.	Deve incluir assinaturas de prevenção de intrusão (IPS) e bloqueio de arquivos maliciosos (Antivírus e Anti-Spyware);	PAN-OS-6.0-Admin-Guide.pdf	pag. 291
2.5.3.	As funcionalidades de IPS, Antivírus e Anti-Spyware devem operar em caráter permanente, podendo ser utilizadas por tempo indeterminado, mesmo que não subsista o direito de receber atualizações ou que não haja contrato de garantia de software com o fabricante.	De acordo com as condições do edital	
2.5.4.	Deve sincronizar as assinaturas de IPS, Antivírus, Anti-Spyware quando implementado em alta disponibilidade ativo/ativo e ativo/passivo	PAN-OS-6.0-Admin-Guide.pdf	pag. 53

2.5.5.	Deve implementar os seguintes tipos de ações para ameaças detectadas pelo IPS e Antispyware: permitir, permitir e gerar log, bloquear, bloquear IP do atacante por um intervalo de tempo e enviar tcp-reset;	Threat-Prevention-Deployment-TechNote-Version-3.0-RevA.pdf	pag. 53
2.5.6.	Deve possuir a capacidade de detectar e prevenir contra ameaças em tráfegos HTTP/2;	Threat-Prevention-Deployment-TechNote-Version-3.0-RevA.pdf	pag. 53 e 54
2.5.7.	As assinaturas devem poder ser ativadas ou desativadas, ou ainda habilitadas apenas em modo de monitoração;	Threat-Prevention-Deployment-TechNote-Version-3.0-RevA.pdf	pag. 53 e 54
2.5.8.	Exceções por IP de origem ou de destino devem ser possíveis nas regras, de forma geral e assinatura a assinatura;	PAN-OS-6.0-Admin-Guide.pdf	pag. 299
2.5.9.	Deve suportar granularidade nas políticas de IPS Antivírus e Anti-Spyware, possibilitando a criação de diferentes políticas por zona de segurança, endereço de origem, endereço de destino, serviço e a combinação de todos esses itens.	PAN-OS-6.0-Admin-Guide.pdf	pag. 27
2.5.10.	Deve permitir o bloqueio de vulnerabilidades.	PAN-OS-6.0-Admin-Guide.pdf	pag. 291
2.5.11.	Deve permitir o bloqueio de exploits conhecidos.	PAN-OS-6.0-Admin-Guide.pdf	pag. 291
2.5.12.	Deve incluir proteção contra ataques de negação de serviços.	PAN-OS-6.0-Admin-Guide.pdf	pag. 297
2.5.13.	Deve suportar a inspeção e criação de regras de proteção de DOS e QOS para o conteúdo de tráfego tunelado pelo protocolo GRE;	panorama-ds.pdf	650
2.5.14.	Deverá possuir os seguintes mecanismos de inspeção de IPS:		
2.5.14.1.	Análise de padrões de estado de conexões;	Content_ID_tech.pdf	pag. 3
2.5.14.2.	Análise de decodificação de protocolo;	Content_ID_tech.pdf	pag. 3
2.5.14.3.	Análise para detecção de anomalias de protocolo;	Content_ID_tech.pdf	pag. 3
2.5.14.4.	Análise heurística;	Content_ID_tech.pdf	pag. 3
2.5.14.5.	IP Defragmentation;	Content_ID_tech.pdf	pag. 3
2.5.14.6.	Remontagem de pacotes de TCP;	Content_ID_tech.pdf	pag. 3
2.5.14.7.	Bloqueio de pacotes malformados.	Content_ID_tech.pdf	pag. 3
2.5.15.	Ser imune e capaz de impedir ataques básicos como: Synflood, ICMPflood, UDPflood, etc;	PAN-OS-6.0-Admin-Guide.pdf	pag. 297
2.5.16.	Detectar e bloquear a origem de portscans com possibilidade de criar exceções para endereços IPs de ferramentas de monitoramento da organização;	panorama-ds.pdf	107 4
2.5.17.	Bloquear ataques efetuados por worms conhecidos, permitindo ao administrador acrescentar novos padrões;	PAN-OS-6.0-Admin-Guide.pdf	pag. 294
2.5.18.	Suportar os seguintes mecanismos de inspeção contra ameaças de rede: análise de padrões de estado de conexões, análise de decodificação de protocolo, análise para detecção de anomalias de protocolo, análise heurística, IP Defragmentation, remontagem de pacotes de TCP e bloqueio de pacotes malformados;	Content_ID_tech.pdf	pag. 3
2.5.19.	Possuir assinaturas específicas para a mitigação de ataques DoS e DDoS;	Application DDoS Mitigation.pdf / https://threatvault.paloaltonetworks.com/	pag. 1 a 27 / pag. 1

2.5.20.	Possuir assinaturas para bloqueio de ataques de buffer overflow;	Content_ID_tech.pdf	pag. 3
2.5.21.	Deverá possibilitar a criação de assinaturas customizadas pela interface gráfica do produto;	PAN-OS-5.0-Admin-Guide.pdf	pag. 243
2.5.22.	Deve permitir usar operadores de negação na criação de assinaturas customizadas de IPS e anti-spyware, permitindo a criação de exceções com granularidade nas configurações;	PAN-OS-9-Admin-Guide.pdf	Pag 619
2.5.23.	Permitir o bloqueio de vírus e spywares em, pelo menos, os seguintes protocolos: HTTP, FTP, SMB, SMTP e POP3;	PAN-OS-6.0-Admin-Guide.pdf	pag. 294
2.5.24.	É permitido uso de appliance externo (antivírus de rede), para o bloqueio de vírus e spywares em protocolo SMB de forma a conter malwares se espalhando horizontalmente pela rede;	De acordo com as condições do edital.	
2.5.25.	Suportar bloqueio de arquivos por tipo;	PAN-OS-6.0-Admin-Guide.pdf	pag. 297
2.5.26.	Identificar e bloquear comunicação com botnets;	ThreatPrevetion.pdf	pag. 3 e 4
2.5.27.	Deve suportar varias técnicas de prevenção, incluindo Drop e tcp-rst (Cliente, Servidor e ambos);	panorama-ds.pdf	995, 996
2.5.28.	Deve suportar referencia cruzada com CVE;	PAN-OS-6.0-Admin-Guide.pdf	pag. 309
2.5.29.	Registrar na console de monitoração as seguintes informações sobre ameaças identificadas:		
2.5.29.1.	O nome da assinatura ou do ataque, aplicação, usuário, origem e o destino da comunicação, além da ação tomada pelo dispositivo;	PAN-OS-6.0-Admin-Guide.pdf	pag 157 a 159
2.5.29.2.	Deve suportar a captura de pacotes (PCAP), por assinatura de IPS e Antispyware;	PAN-OS-5.0-Admin-Guide.pdf	pag. 278
2.5.29.3.	Deve permitir que na captura de pacotes por assinaturas de IPS e Antispyware seja definido o número de pacotes a serem capturados. Esta captura deve permitir selecionar, no mínimo, 50 pacotes;	PAN-OS-6.0.13-RN.pdf	pag. 6
2.5.29.4.	Deve possuir a função resolução de endereços via DNS, para que conexões com destino a domínios maliciosos sejam resolvidas pelo Firewall com endereços (IPv4 e IPv6), previamente definidos;	PAN-OS-6.0-Admin-Guide.pdf	pag. 171
2.5.29.5.	Permitir o bloqueio de vírus, pelo menos, nos seguintes protocolos: HTTP, FTP, SMB, SMTP e POP3;	PAN-OS-6.0-Admin-Guide.pdf	pag. 294
2.5.29.6.	Os eventos devem identificar o país de onde partiu a ameaça;	panorama-ds.pdf	367
2.5.29.7.	Deve incluir proteção contra vírus em conteúdo HTML e javascript, software espião (spyware) e worms.	panorama-ds.pdf	995
2.5.29.8.	Proteção contra downloads involuntários usando HTTP de arquivos executáveis. maliciosos.	ThreatPrevetion.pdf	pag. 3
2.5.29.9.	Rastreamento de vírus em pdf.	Specs_summary_sept2015.pdf	pag. 2
2.5.29.10	Deve permitir a inspeção em arquivos comprimidos que utilizam o algoritmo deflate (zip, gzip, etc.)	PAN-OS-5.0-Admin-Guide.pdf	pag. 422
2.5.29.11	Deve ser possível a configuração de diferentes políticas de controle de ameaças e ataques baseado em políticas do firewall considerando Usuários, Grupos de usuários, origem, destino, zonas de segurança, etc, ou seja, cada política de firewall poderá ter uma configuração diferentes de IPS, sendo essas politicas por Usuários, Grupos de usuário, origem, destino, zonas de segurança.	PAN-OS-6.0-Admin-Guide.pdf	pag. 27 e 274

2.6.	Análise de malwares modernos		
2.6.1.	Devido aos Malwares hoje em dia serem muito dinâmicos e um antivírus comum reativo não ser capaz de detectar os mesmos com a mesma velocidade que suas variações são criadas, a solução ofertada deve possuir funcionalidades para análise de Malwares não conhecidos incluídas na própria ferramenta ou entregue com composição com outro fabricante;	wildfire.pdf	pag. 1
2.6.2.	O dispositivo de proteção deve ser capaz de enviar arquivos trafegados de forma automática para análise "In Cloud" ou local, onde o arquivo será executado e simulado em ambiente controlado;	wildfire.pdf	pag. 1
2.6.3.	Selecionar através de políticas granulares quais tipos de arquivos sofrerão esta análise incluindo, mas não limitado a: endereço IP de origem/destino, usuário/grupo do AD/LDAP, aplicação, porta, URL/categoria de URL de destino, tipo de arquivo e todas estas opções simultaneamente;	PAN-OS-6.0-Admin-Guide.pdf	pag. 50
2.6.4.	Deve possuir a capacidade de diferenciar arquivos analisados em pelo menos três categorias: malicioso, não malicioso e arquivos não maliciosos, mas com características indesejáveis como softwares que deixa o sistema operacional lento, que alteram parâmetros do sistema, etc.;	wildfire-Admin-Guide.pdf	pag. 13
2.6.5.	Suportar a análise com pelo menos 100 (cem) tipos de comportamentos maliciosos para a análise da ameaça não conhecida;	wildfire.pdf	pag. 2
2.6.6.	Suportar a análise de arquivos maliciosos em ambiente controlado com, no mínimo, sistema operacional Windows XP e Windows 7;	wildfire.pdf	pag. 1
2.6.7.	Deve suportar a monitoração de arquivos trafegados na internet (HTTPs, FTP, HTTP, SMTP) como também arquivos trafegados internamente entre servidores de arquivos usando SMB em todos os modos de implementação: sniffer, transparente e L3;	wildfire.pdf	pag. 1 e 2
2.6.8.	A solução deve possuir a capacidade de analisar em sand-box links (http e https) presentes no corpo de e-mails trafegados em SMTP e POP3. Deve ser gerado um relatório caso a abertura do link pela sand-box o identifique como site hospedeiro de exploits;	wildfire-Admin-Guide.pdf	pag. 14
2.6.9.	A análise de links em sand-box deve ser capaz de classificar sites falsos na categoria de phishing e atualizar a base de filtro de URL da solução;	wildfire-Admin-Guide.pdf	13
2.6.10.	Para ameaças trafegadas em protocolo SMTP e POP3, a solução deve ter a capacidade de mostrar nos relatórios o remetente, destinatário e assunto dos e-mails permitindo identificação ágil do usuário vítima do ataque;	wildfire-Admin-Guide.pdf	pag. 14
2.6.11.	O sistema de análise "In Cloud" ou local deve prover informações sobre as ações do Malware na máquina infectada, informações sobre quais aplicações são utilizadas para causar/propagar a infecção, detectar aplicações não confiáveis utilizadas pelo Malware, gerar assinaturas de Antivírus e Anti-spyware automaticamente, definir URLs não confiáveis utilizadas pelo novo Malware e prover informações sobre o usuário infectado (seu endereço ip e seu login de rede);	wildfire.pdf	pag. 3

2.6.12.	O sistema automático de análise "In Cloud" ou local deve emitir relatório com identificação de quais soluções de antivírus existentes no mercado possuem assinaturas para bloquear o malware;	wildfire.pdf	pag. 3
2.6.13.	Deve permitir exportar o resultado das análises de malwares de dia Zero em PDF e CSV a partir da própria interface de gerência;	whats_new_PAN-OS-6.0.pdf	pag. 1
2.6.14.	Deve permitir o download dos malwares identificados a partir da própria interface de gerência;	wildfire.pdf	pag. 3
2.6.15.	Deve permitir visualizar o resultados das análises de malwares de dia zero nos diferentes sistemas operacionais suportados;	wildfire.pdf	pag. 3
2.6.16.	Deve permitir informar ao fabricante quanto a suspeita de ocorrências de falso-positivo e falso-negativo na análise de malwares de dia Zero a partir da própria interface de gerência.	whats_new_PAN-OS-6.0.pdf	pag. 1
2.6.17.	Caso a solução seja fornecida em appliance local, deve possuir, no mínimo, 28 ambientes controlados (sand-box) independentes para execução simultânea de arquivos suspeitos;	De acordo com o Edital.	
2.6.18.	Caso seja necessário licenças de sistemas operacional e softwares para execução de arquivos no ambiente controlado (sand-box), as mesmas devem ser fornecidas em sua totalidade, sem custos adicionais para a contratante;	De acordo com o Edital.	
2.6.19.	Suportar a análise de arquivos executáveis, DLLs, ZIP e criptografados em SSL no ambiente controlado;	wildfire.pdf	1 e 2
2.6.20.	Suportar a análise de arquivos do pacote office (.doc, .docx, xls, xlsx, ppt, pptx), arquivos java (.jar e class), Android APKs, MacOS (mach-0, DMG e PKG), Linux (ELF), RAR e 7zip no ambiente sandbox;	wildfire.pdf	pag. 3
2.6.21.	Deve atualizar a base com assinaturas para bloqueio dos malwares identificados em sand-box com frequencia de pelo menos 5 minutos;	wildfire.pdf	pag. 3
2.6.22.	Permitir o envio de arquivos e links para analise no mambiente controlado de forma automatica via API	wildfire.pdf	pag. 3
2.6.23.	Deve permitir o envio para analise em sand-box de malwares bloqueados pelo antivirus da solução.	wildfire.pdf	pag. 3
2.7.	Filtro de url		
2.7.1.	A plataforma de segurança deve possuir as seguintes funcionalidades de filtro de URL:		
2.7.1.1.	Permite especificar política por tempo, ou seja, a definição de regras para um determinado horário ou período (dia, mês, ano, dia da semana e hora);	PAN-OS-6.0-Admin-Guide.pdf	pág. 32 e 33
2.7.1.2.	Deve ser possível a criação de políticas por Usuários, Grupos de Usuários, Ips, Redes e Zonas de segurança.	PAN-OS-6.0-Admin-Guide.pdf	pág. 26, 27, 31 a 33 e 274
2.7.1.3.	Deverá incluir a capacidade de criação de políticas baseadas na visibilidade e controle de quem está utilizando quais URLs através da integração com serviços de diretório, autenticação via ldap, Active Directory, E-directory e base de dados local.	PAN-OS-6.0-Admin-Guide.pdf	pag. 230
2.7.1.4.	Permite popular todos os logs de URL com as informações dos usuários conforme descrito na integração com serviços de diretório;	PAN-OS-6.0-Admin-Guide.pdf	pag. 360
2.7.1.5.	Suporta a capacidade de criação de políticas baseadas no controle por URL e Categoria de URL;	PAN-OS-6.0-Admin-Guide.pdf	pag. 342

2.7.1.6.	Deve bloquear o acesso a sites de busca (Google, Bing e Yahoo), caso a opção Safe Search esteja desabilitada. Deve ainda exibir página de bloqueio fornecendo instruções ao usuário de como habilitar a função;	PAN-OS-6.0-Admin-Guide.pdf	pag. 366
2.7.1.7.	Suporta base ou cache de URLs local no appliance, evitando delay de comunicação/validação das URLs;	PAN-OS-6.0-Admin-Guide.pdf	pag. 340
2.7.1.8.	Possui pelo menos 60 categorias de URLs;	PAN-OS-6.0-Admin-Guide.pdf	pag. 340
2.7.1.9.	Deve classificar o nível de risco de URLs em, pelo menos, três níveis: baixo, médio e alto;	PAN-OS-6.0-Admin-Guide.pdf	Pa g 344
2.7.1.10.	Deve possuir categoria específica para classificar domínios recém registrados (com menos de 32 dias);	PAN-OS-6.0-Admin-Guide.pdf	Pa g 344
2.7.1.11.	A solução deve ter a capacidade de classificar sites em mais de uma categoria, de acordo com a necessidade;	PAN-OS-6.0-Admin-Guide.pdf	Pa g 345
2.7.1.12.	A categorização de URL deve analisar toda a URL e não somente até o nível de diretório;	PAN-OS-6.0-Admin-Guide.pdf	Pag 340, 351
2.7.1.13.	Suporta a criação categorias de URLs customizadas;	PAN-OS-6.0-Admin-Guide.pdf	Pa g 344
2.7.1.14.	Suporta a exclusão de URLs do bloqueio, por categoria;	PAN-OS-6.0-Admin-Guide.pdf	pag. 340 e 341
2.7.1.15.	Permite a customização de página de bloqueio;	PAN-OS-6.0-Admin-Guide.pdf	pag. 343
2.7.1.16.	Deve proteger contra o roubo de credenciais, usuários e senhas identificadas através da integração com Active Directory submetidos em sites não corporativos. Deve ainda permitir criação de regra onde usuários do Active Directory só possam enviar informações de login para sites autorizados na solução;	panorama-ds.pdf	557
2.7.1.17.	Deve permitir bloquear o acesso do usuário caso o mesmo tente fazer o envio de suas credenciais em sites classificados como phishing pelo filtro de URL da solução;	panorama-ds.pdf	557
2.7.1.18.	Permite o bloqueio e continuação (possibilitando que o usuário acesse um site potencialmente bloqueado informando o mesmo na tela de bloqueio e possibilitando a utilização de um botão "Continuar" para permitir o usuário continuar acessando o site);	panorama-ds.pdf	562
2.7.1.19.	Suporta a inclusão nos logs do produto de informações das atividades dos usuários;	panorama-ds.pdf	336
2.7.1.20.	Deve salvar nos logs as informações dos seguintes campos do cabeçalho HTTP nos acessos a URLs: UserAgent, Referer, e X-Forwarded For;	panorama-ds.pdf	652
2.8.	Identificação de Usuários		
2.8.1.	Deve incluir a capacidade de criação de políticas baseadas na visibilidade e controle de quem está utilizando quais aplicações através da integração com serviços de diretório, autenticação via ldap, Active Directory, E-directory e base de dados local;	pan-os	447, 449
2.8.2.	Deve possuir integração com Microsoft Active Directory para identificação de usuários e grupos permitindo granularidade de controle/políticas baseadas em usuários e grupos de usuários;	pan-os	447, 449

2.8.3.	Deve possuir integração com Radius para identificação de usuários e grupos permitindo granularidade de controle/políticas baseadas em usuários e grupos de usuários;	pan-os	158, 159
2.8.4.	Deve implementar a criação de políticas de segurança baseada em atributos específicos do Radius, incluindo mas não limitado a: baseado no sistema operacional do usuário remoto exigir autenticação padrão Windows e on-time password (OTP) para usuários Android;	pan-os	158, 159
2.8.5.	Deve possuir integração com Ldap para identificação de usuários e grupos permitindo granularidade de controle/políticas baseadas em Usuários e Grupos de usuários;	pan-os	160
2.8.6.	Deve suportar o recebimento eventos de autenticação de controladoras wireless, dispositivos 802.1x e soluções NAC via syslog, para a identificação de endereços IP e usuários;	pan-os	459
2.8.7.	4.1.1.122. Deve permitir o controle, sem instalação de cliente de software, em equipamentos que solicitem saída a internet para que antes de iniciar a navegação, expanda-se um portal de autenticação residente no firewall (Captive Portal);	pan-os	482
2.8.8.	Suporte a autenticação Kerberos;	pan-os	157
2.8.9.	Deve suportar autenticação via Kerberos para administradores da plataforma de segurança, captive Portal e usuário de VPN SSL;	pan-os	157
2.8.10.	Deve possuir suporte a identificação de múltiplos usuários conectados em um mesmo endereço IP em ambientes Citrix e Microsoft Terminal Server, permitindo visibilidade e controle granular por usuário sobre o uso das aplicações que estão nestes serviços;	pan-os	450
2.8.11.	Deve identificar usuários através de leitura do campo x-forwarded-for, populando nos logs do firewall o endereço IP, bem como o usuário de rede responsável pelo acesso;	pan-os	450
2.8.12.	Deve permitir a criação de políticas de segurança baseadas em usuários de rede com reconhecimento dos mesmos através de leitura do campo x-forwarded-for;	pan-os	450
2.8.13.	Deve suportar e operar com Security Assertion Markup Language (SAML) 2.0 com single sign-on e single logout para as funcionalidades de Captive Portal e VPN SSL (client to server) permitindo login único e interativo para determinados acessos internos e externos a organização.	pan-os	156
2.8.14.	Deve implementar a criação de grupos customizados de usuários no firewall, baseado em atributos do LDAP/AD;	pan-os	557
2.8.15.	Deve possuir suporte a identificação de múltiplos usuários conectados em um mesmo endereço IP em servidores acessados remotamente, mesmo que não sejam servidores Windows.	pan-os	449
2.9.	QOS		

2.9.1.	Com a finalidade de controlar aplicações e tráfego cujo consumo possa ser excessivo, (como youtube, ustream, etc) e ter um alto consumo de largura de banda, se requer que a solução, além de poder permitir ou negar esse tipo de aplicações, deve ter a capacidade de controlá-las por políticas de máximo de largura de banda quando forem solicitadas por diferentes usuários ou aplicações, tanto de áudio como de vídeo streaming.	pan-os	703
2.9.2.	Suportar a criação de políticas de QoS por:	pan-os	
2.9.2.1.	Endereço de origem	pan-os	705
2.9.2.2.	Endereço de destino	pan-os	705
2.9.2.3.	Por usuário e grupo do LDAP/AD.	pan-os	705
2.9.2.4.	Por aplicações, incluindo, mas não limitado a Skype, Bittorrent, YouTube e Azureus;	pan-os	705
2.9.2.5.	Por porta;	pan-os	705
2.9.3.	O QoS deve possibilitar a definição de classes por:	pan-os	
2.9.3.1.	Banda Garantida	pan-os	706
2.9.3.2.	Banda Máxima	pan-os	706
2.9.3.3.	Fila de Prioridade	pan-os	706
2.9.4.	Suportar priorização RealTime de protocolos de voz (VOIP) como H.323, SIP, SCCP, MGCP e aplicações como Skype.	pan-os	706
2.9.5.	Suportar marcação de pacotes Diffserv, inclusive por aplicação;	pan-os	709
2.9.6.	Deve implementar QOS (traffic-shapping), para pacotes marcados por outros ativos na rede (DSCP). A priorização e limitação do tráfego deve ser efetuada nos dois sentidos da conexão (inbound e outbound);	pan-os	719
2.9.7.	Disponibilizar estatísticas RealTime para classes de QoS.	pan-os	723
2.9.8.	Deve suportar QOS (traffic-shapping), em interface agregadas;	pan-os	706
2.9.9.	Deverá permitir o monitoramento do uso que as aplicações fazem por bytes, sessões e por usuário.	pan-os	712
2.10.	Filtro de dados		
2.10.1.	Permite a criação de filtros para arquivos e dados pré-definidos;	pan-os	538
2.10.2.	Os arquivos devem ser identificados por extensão e assinaturas;	pan-os	540
2.10.3.	Permite identificar e opcionalmente prevenir a transferência de vários tipos de arquivos (MS Office, PDF, etc) identificados sobre aplicações (P2P, InstantMessaging, SMB, etc);	pan-os	540
2.10.4.	Suportar identificação de arquivos compactados e a aplicação de políticas sobre o conteúdo desses tipos de arquivos;	pan-os	540
2.10.5.	Permitir identificar e opcionalmente prevenir a transferência de informações sensíveis, incluindo, mas não limitado a número de cartão de crédito, possibilitando a criação de novos tipos de dados via expressão regular;	pan-os	540
2.10.6.	Permitir listar o número de aplicações suportadas para controle de dados;	pan-os	540

2.10.7.	Permitir listar o número de tipos de arquivos suportados para controle de dados;	pan-os	540
2.11.	Geo-Localização		
2.11.1.	Suportar a criação de políticas por Geo Localização, permitindo o tráfego de determinado País/Países sejam bloqueados.	pan-os	93
2.11.2.	Deve possibilitar a visualização dos países de origem e destino nos logs dos acessos.	pan-os	324
2.11.3.	Deve permitir visualizar nos logs e criar políticas para liberar e bloquear tráfego de países por: tipo de arquivo, aplicação e categoria de URL;	pan-os	327
2.11.4.	Deve possibilitar a criação de regiões geográficas pela interface gráfica e criar políticas utilizando as mesmas.	pan-os	327
2.12.	VPN		
2.12.1.	Suportar VPN Site-to-Site e Cliente-To-Site;	pan-os	730
2.12.2.	Suportar IPSec VPN;	pan-os	745
2.12.3.	Suportar SSL VPN;	globalprotect-admin-guide	135
2.12.4.	A VPN IPSEc deve suportar:		
2.12.4.1.	DES e 3DES;	pan-os	733
2.12.4.2.	Autenticação MD5 e SHA-1;	pan-os	733
2.12.4.3.	4.1.1.154.3. Diffie-Hellman Group 1 , Group 2, Group 5 e Group 14;	pan-os	733
2.12.4.4.	Algoritmo Internet Key Exchange (IKEv1 e v2);	pan-os	733
2.12.4.5.	4.1.1.154.5. AES 128, 192 e 256 (Advanced Encryption Standard)	pan-os	733
2.12.4.6.	Autenticação via certificado IKE PKI.	pan-os	733
2.12.5.	Deve possuir interoperabilidade com os seguintes fabricantes:		
2.12.5.1.	Cisco;	pan-os	727
2.12.5.2.	Checkpoint;	pan-os	727
2.12.5.3.	Juniper;	pan-os	727
2.12.5.4.	Palo Alto Networks;	pan-os	727
2.12.5.5.	Fortinet;	pan-os	727
2.12.5.6.	Sonic Wall;	pan-os	727
2.12.6.	Deve permitir habilitar, desabilitar, reiniciar e atualizar IKE gateways e túneis de VPN IPSEc a partir da interface gráfica da solução, facilitando o processo de troubleshooting;	pan-os	749
2.12.7.	A VPN SSL deve suportar:		
2.12.7.1.	O usuário realizar a conexão por meio de cliente instalado no sistema operacional do equipamento ou por meio de interface WEB;	globalprotect-admin-guide	135
2.12.7.2.	A funcionalidades de VPN SSL devem ser atendidas com ou sem o uso de agente;	globalprotect-admin-guide	135
2.12.7.3.	Atribuição de endereço IP nos clientes remotos de VPN SSL;	globalprotect-admin-guide	203
2.12.7.4.	Deve permitir a atribuição de IPs fixos nos usuários remotos de VPN SSL;	globalprotect-admin-guide	75

2.12.7.5.	Deve permitir a criação de rotas de acesso e faixas de endereços IP atribuídas a clientes remotos de VPN de forma customizada por usuário AD/LDAP e grupo de usuário AD/LDAP;	globalprotect-admin-guide	31, 141
2.12.7.6.	Deve permitir que todo o tráfego dos usuários remotos de VPN seja escoado para dentro do túnel de VPN, impedindo comunicação direta com dispositivos locais como proxies;	globalprotect-admin-guide	730
2.12.7.7.	Atribuição de DNS nos clientes remotos de VPN;	globalprotect-admin-guide	77
2.12.7.8.	Deve permitir que seja definido métodos de autenticação distintos por sistema operacional do dispositivo remoto de VPN (Android, IOS, Mac, Windows e Chrome OS);	globalprotect-admin-guide	85
2.12.7.9.	A solução de VPN deve verificar se o client que está conectando é o mesmo para o qual o certificado foi emitido inicialmente. O acesso deve ser bloqueado caso o dispositivo não seja o correto;	globalprotect-admin-guide	143
2.12.7.10	Deve possuir lista de bloqueio para dispositivos que forem reportados com roubado ou perdido pelo usuário;	globalprotect-admin-guide	190
2.12.7.11	Deve haver a opção de ocultar o agente de VPN instalado no cliente remoto, tornando o mesmo invisível para o usuário;	globalprotect-admin-guide	94
2.12.7.12	Deve exibir mensagens de notificação customizada toda vez que um usuário remoto se conectar a VPN. Deve permitir que o usuário desabilite a exibição da mensagem nas conexões seguintes;	globalprotect-admin-guide	77
2.12.7.13	Deve avisar ao usuário remoto de VPN quanto a proximidade da expiração de senha LDAP. Deve permitir também a customização da mensagem com informações relevantes para o usuário;	globalprotect-admin-guide	77
2.12.7.14	Dever permitir criar políticas de controle de aplicações, IPS, Antivírus, Antispyware e filtro de URL para tráfego dos clientes remotos conectados na VPN SSL;	globalprotect-admin-guide	16
2.12.7.15	A VPN SSL deve suportar proxy arp e uso de interfaces PPPOE;	https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIUuCAK	1
2.12.7.16	Suportar autenticação via AD/LDAP, OTP (One Time Password), certificado e base de usuários local;	globalprotect-admin-guide	27
2.12.7.17	Deve permitir a distribuição de certificado para o usuário de remoto através do portal de VPN de forma automatizada;	globalprotect-admin-guide	45
2.12.7.18	Deve possuir lista de bloqueio para dispositivos em casos quando, por exemplo, o usuário reportar que o dispositivo foi perdido ou roubado;	globalprotect-admin-guide	190
2.12.7.19	Permite estabelecer um túnel VPN client-to-site do cliente a plataforma de segurança, fornecendo uma solução de single-sign-on aos usuários, integrando-se com as ferramentas de Windows-logon;	globalprotect-admin-guide	93
2.12.7.20	Suporta leitura e verificação de CRL (certificate revocation list);	globalprotect-admin-guide	143
2.12.7.21	Permite a aplicação de políticas de segurança e visibilidade para as aplicações que circulam dentro dos túneis SSL;	globalprotect-admin-guide	143

2.12.7.22	O agente de VPN a ser instalado nos equipamentos desktop e laptops, deve ser capaz de ser distribuído de maneira automática via Microsoft SMS, Active Directory e ser descarregado diretamente desde o seu próprio portal, o qual residirá no centralizador de VPN;	globalprotect-admin-guide	103
2.12.7.23	O agente deverá comunicar-se com o portal para determinar as políticas de segurança do usuário,	globalprotect-admin-guide	86
2.12.8.	Deve permitir que a conexão com a VPN SSL seja estabelecida das seguintes formas:		
2.12.8.1.	Antes do usuário autenticar na estação;	globalprotect-admin-guide	22
2.12.8.2.	Após autenticação do usuário na estação;	globalprotect-admin-guide	22
2.12.8.3.	Sob demanda do usuário;	globalprotect-admin-guide	115
2.12.9.	Deve manter uma conexão segura com o portal durante a sessão.	globalprotect-admin-guide	14
2.12.10.	O agente de VPN SSL client-to-site deve ser compatível com pelo menos: Windows XP, Vista Windows 7, Windows 8, Mac OSx e Chrome OS;	globalprotect-admin-guide	11
2.12.11.	O portal de VPN deve enviar ao cliente remoto, a lista de gateways de VPN ativos para estabelecimento da conexão, os quais devem poder ser administrados centralmente;	globalprotect-admin-guide	225
2.12.12.	Deve haver a opção do cliente remoto escolher manualmente o gateway de VPN e de forma automática através da melhor rota entre os gateways disponíveis com base no tempo de resposta mais rápido;	globalprotect-admin-guide	86
2.12.13.	Deve possuir a capacidade de identificar se a origem da conexão de VPN é externa ou interna;	globalprotect-admin-guide	68
2.13.	Console de gerência e monitoração		
2.13.1.	Centralizar a administração de regras e políticas do cluster, usando uma única interface de gerenciamento;	Panorama.pdf	pag. 1
2.13.2.	O gerenciamento da solução deve suportar acesso via SSH, cliente ou WEB (HTTPS) e API aberta;	Panorama_AdminGuide.pdf	pag. 73
2.13.3.	Deve permitir substituir o certificado de fábrica no acesso HTTPS a gerência do firewall como possibilidade de uso de certificado criado localmente na própria solução ou importado de fonte externa;	Panorama.pdf	238
2.13.4.	Caso haja a necessidade de instalação de cliente para administração da solução o mesmo deve ser compatível com sistemas operacionais Windows e Linux;	PAN-OS-6.0-Admin-Guide.pdf	pag. 62
2.13.5.	O gerenciamento deve permitir/possuir:		
2.13.5.1.	Criação e administração de políticas de firewall e controle de aplicação;	Panorama.pdf	pag. 1 e 2
2.13.5.2.	Criação e administração de políticas de IPS, Antivírus e Anti-Spyware;	Panorama.pdf	pag. 1 e 2
2.13.5.3.	Criação e administração de políticas de Filtro de URL;	Panorama.pdf	pag. 1 e 2
2.13.5.4.	Monitoração de logs;	Panorama.pdf	pag. 1
2.13.5.5.	Ferramentas de investigação de logs;	Panorama.pdf	pag. 2
2.13.5.6.	Debugging;	PAN-OS-5.0-Admin-Guide.pdf,	pag. 322, 323,

			354 e 370
2.13.5.7.	Captura de pacotes.	PAN-OS-6.0-Admin-Guide.pdf	pag. 174
2.13.6.	Acesso concorrente de administradores;	PAN-OS-7.0-Admin-Guide.pdf	pag. 69
2.13.7.	Deve permitir que administradores concorrentes façam modificações, valide configurações e reverta configurações do firewall simultaneamente e que cada administrador consiga aplicar apenas as suas alterações de forma independente das realizadas por outro administrador	Panorama.pdf	64
2.13.8.	Deve mostrar ao administrador do firewall a hora e data do último login e tentativas de login com falha para acessos a partir da interface gráfica e CLI.	Panorama.pdf	62
2.13.9.	Deve possuir mecanismo busca global na solução onde possa se consultar por uma string tais como: nome de objetos, ID ou nome de ameaças, nome de aplicações, nome de políticas, endereços IPs, permitindo a localização e uso dos mesmo na configuração do dispositivo;	Panorama_AdminGuide.pdf	186 a 189
2.13.10.	Deve possuir um mecanismo de busca por comandos no gerenciamento via SSH, facilitando a localização de comandos;	whats_new_PAN-OS-6.0.pdf	pag. 6
2.13.11.	Deve permitir usar palavras chaves e cores para facilitar identificação de regras;	whats_new_PAN-OS-6.0.pdf	pag. 6
2.13.12.	Deve permitir monitorar via SNMP falhas de hardware, inserção ou remoção de fontes, discos e coolers, uso de recursos por número elevado de sessões, número de túneis estabelecidos na VPN cliente-to-site, porcentagem de utilização em referência ao número total suportado/licenciado e número de sessões estabelecidas, estatísticas/taxa de logs, uso de disco, período de retenção dos logs e status do envio de logs para soluções externas;	whats_new_PAN-OS-6.0.pdf	pag. 6
2.13.13.	Deve suportar também o monitoramento dos seguintes recursos via SNMP: IP fragmentation, TCP state e dropped packets;	Panorama_AdminGuide.pdf	19, 182, 183 e 189
2.13.14.	Bloqueio de alterações, no caso acesso simultâneo de dois ou mais administradores;	PAN-OS-7.0-Admin-Guide.pdf	pag. 69
2.13.15.	Definição de perfis de acesso à console com permissões granulares como: acesso de escrita, acesso de leitura, criação de usuários, alteração de configurações;	PAN-OS-7.0-Admin-Guide.pdf	pag. 75
2.13.16.	Autenticação integrada ao Microsoft Active Directory e servidor Radius;	PAN-OS-7.0-Admin-Guide.pdf	pag. 77
2.13.17.	Localização de em quais regras um endereço IP, IP Range, subnet ou objetos estão sendo utilizados;	PAN-OS-5.0-Admin-Guide.pdf,	pag. 185
2.13.18.	Deve atribuir sequencialmente um número a cada regra de firewall, NAT, QOS e regras de DOS;	whats_new_PAN-OS-6.0.pdf	pag. 6
2.13.19.	Criação de regras que fiquem ativas em horário definido;	PAN-OS-5.0-Admin-Guide.pdf,	pag. 249
2.13.20.	Criação de regras com data de expiração;	PAN-OS-5.0-Admin-Guide.pdf,	pag. 249
2.13.21.	Backup das configurações e rollback de configuração para a última configuração salva;	Panorama_AdminGuide.pdf	pag. 61
2.13.22.	Suportar Rollback de Sistema Operacional para a última versão local;	PAN-OS-7.0-Admin-Guide.pdf	pag. 752

2.13.23.	Habilidade de upgrade via SCP, TFTP e interface de gerenciamento;	TFTP Upgrade Instructions - Live Community.pdf / SCP Upgrade Instructions - Live Community.pdf / How to Upgrade PAN-OS on a Palo Alto Networks Devi...pdf	pag. 1
2.13.24.	Deve possuir mecanismo de análise de impacto na política de segurança antes de atualizar a base com novas aplicações disponibilizadas pelo fabricante;	Panorama_AdminGuide.pdf	210
2.13.25.	Validação de regras antes da aplicação	Panorama_AdminGuide.pdf	pag. 199
2.13.26.	Deve implementar mecanismo de validação de configurações antes da aplicação das mesmas permitindo identificar erros, tais como: rota de destino inválida, regras em shadowing etc.	Panorama_AdminGuide.pdf	211
2.13.27.	É permitido o uso de appliance externo para permitir a validação de regras antes da aplicação.	De acordo com as condições do edital	
2.13.28.	Validação da políticas, avisando quando houver regras que, ofusquem ou conflitem com outras (shadowing);	Panorama_AdminGuide.pdf	pag. 199
2.13.29.	É permitido o uso de appliance externo para permitir a validação de políticas, avisando quando houver regras que, ofusquem ou conflitem com outras (shadowing);	Panorama_AdminGuide.pdf	pag. 199
2.13.30.	Deve possibilitar a visualização e comparação de configurações Atuais, configuração anterior e configurações antigas.	Panorama_AdminGuide.pdf	pag. 198
2.13.31.	Deve permitir auditar regras de segurança exibindo quadro comparativo das alterações de uma regra em relação a versão anterior;	PAN-OS-9-Admin-Guide.pdf	Pa g 417
2.13.32.	Deve possibilitar a integração com outras soluções de SIEM de mercado (third-party SIEM vendors)	PAN-OS-6.0-Admin-Guide.pdf	pag. 190
2.13.33.	Geração de logs de auditoria detalhados, informando a configuração realizada, o administrador que a realizou e o horário da alteração;	PAN-OS-6.0-Admin-Guide.pdf	pag. 178
2.13.34.	Deverá ter a capacidade de gerar um relatório gráfico que permita visualizar as mudanças na utilização de aplicações na rede no que se refere a um período de tempo anterior, para permitir comparar os diferentes consumos realizados pelas aplicações no tempo presente com relação ao passado;	PAN-OS-6.0-Admin-Guide.pdf	pag. 172
2.13.35.	Geração de relatórios com mapas geográficos gerados em tempo real para a visualização de origens e destinos do tráfego gerado na instituição;	PAN-OS-6.0-Admin-Guide.pdf	pag. 171
2.13.36.	Deve prover relatórios com visão correlacionada de aplicações e filtro de arquivos, para melhor diagnóstico e resposta a incidentes;	PAN-OS-6.0-Admin-Guide.pdf	pag. 209
2.13.37.	Deve permitir a criação de Dash-Boards customizados para visibilidades do tráfego de aplicativos, usuários e tráfego bloqueado;	Panorama_AdminGuide.pdf	80 e 180
2.13.38.	O gerenciamento da solução deve possibilitar a coleta de estatísticas de todo o tráfego que passar pelos dispositivos de segurança;	PAN-OS-6.0-Admin-Guide.pdf	pag. 47
2.13.39.	Dever permitir a visualização dos logs de tráfego e malwares modernos (IP de origem, destino, usuário e porta), aplicação e filtro de arquivos em uma única tela;	panorama-ds.pdf	355
2.13.40.	Deve possuir relatórios de utilização dos recursos por aplicações, URL, ameaças que passaram pela solução;	PAN-OS-6.0-Admin-Guide.pdf	pag. 158

2.13.41.	Prover uma visualização sumarizada de todas as aplicações que passaram pela solução;	PAN-OS-6.0-Admin-Guide.pdf	pag. 167
2.13.42.	Deve possuir mecanismo "Drill-Down" para navegação nos relatórios em RealTime;	PAN-OS-5.0-Admin-Guide.pdf,	pag. 255
2.13.43.	Nas opções de "Drill-Down", ser possível identificar o usuário que fez determinado acesso;	PAN-OS-6.0-Admin-Guide.pdf	pag. 280
2.13.44.	Deve possuir relatório de visibilidade e uso sobre aplicativos (SaaS). O relatório também deve mostrar os riscos para a segurança do ambiente, tais como a entrega de malwares através de aplicativos SaaS com a informação do usuário responsável pelo acesso;	panorama-ds.pdf	368
2.13.45.	Os relatórios de visibilidade e uso sobre aplicativos (SaaS) devem poder ser extraídos por grupo de usuários apresentando o uso e consumo de aplicações por grupo de usuário;	panorama-ds.pdf	370
2.13.46.	Deve ser possível exportar os logs em CSV;	PAN-OS-6.0-Admin-Guide.pdf	pag. 181
2.13.47.	Deverá ser possível acessar o equipamento a aplicar configurações durante momentos onde o tráfego é muito alto e a CPU e memória do equipamento estiver totalmente utilizada.	PAN-OS-6.0-Admin-Guide.pdf	pag. 12
2.13.48.	Rotação do log;	PAN-OS-6.0-Admin-Guide.pdf	pag. 12
2.13.49.	Deve permitir que os logs e relatórios sejam rotacionados automaticamente baseado no tempo em que estão armazenados na solução, assim como no espaço em disco usado;	Panorama_AdminGuide.pdf	pag. 217
2.13.50.	Deve permitir fazer o envio de logs para soluções externas de forma granular podendo selecionar quais campos dos logs serão enviados incluindo, mas não limitado a: tipo de ameaça, usuário, aplicação, etc;	pan-os	358
2.13.51.	Exibição das seguintes informações, de forma histórica e em tempo real (atualizado de forma automática e contínua a cada 1 minuto):	PAN-OS-6.0-Admin-Guide.pdf	pag. 158
2.13.51.1	Situação do dispositivo e do cluster	PAN-OS-6.0-Admin-Guide.pdf	pag. 158
2.13.51.2	Principais aplicações;	PAN-OS-6.0-Admin-Guide.pdf	pag. 158
2.13.51.3	Principais aplicações por risco;	PAN-OS-6.0-Admin-Guide.pdf	pag. 158
2.13.51.4	Administradores autenticados na gerência da plataforma de segurança;	PAN-OS-6.0-Admin-Guide.pdf	pag. 158
2.13.51.5	Número de sessões simultâneas	PAN-OS-6.0-Admin-Guide.pdf	pag. 158
2.13.51.6	Status das interfaces;	PAN-OS-6.0-Admin-Guide.pdf	pag. 158
2.13.51.7	Uso de CPU;	PAN-OS-6.0-Admin-Guide.pdf	pag. 158
2.13.52.	Geração de relatórios. No mínimo os seguintes relatórios devem ser gerados:	PAN-OS-6.0-Admin-Guide.pdf	pag. 200
2.13.53.	Resumo gráfico de aplicações utilizadas;	PAN-OS-6.0-Admin-Guide.pdf	pag. 200
2.13.54.	Principais aplicações por utilização de largura de banda de entrada e saída;	PAN-OS-6.0-Admin-Guide.pdf	pag. 200
2.13.55.	Principais aplicações por taxa de transferência de bytes;	PAN-OS-6.0-Admin-Guide.pdf	pag. 200
2.13.56.	Principais hosts por número de ameaças identificadas;	PAN-OS-6.0-Admin-Guide.pdf	pag. 200
2.13.57.	Atividades de um usuário específico e grupo de usuários do AD/LDAP, incluindo aplicações acessadas de rede vinculadas a este tráfego;	PAN-OS-6.0-Admin-Guide.pdf	pag. 203 a 206

2.13.58.	Deve permitir a criação de relatórios personalizados;	PAN-OS-6.0-Admin-Guide.pdf	pag. 203
2.14.	Em cada critério de pesquisa do log deve ser possível incluir múltiplas entradas (ex. 10 redes e IP's distintos; serviços HTTP, HTTPS e SMTP), exceto no campo horário, onde deve ser possível definir um faixa de tempo como critério de pesquisa;	PAN-OS-6.0-Admin-Guide.pdf	pag. 203
2.15.	Gerar alertas automáticos via:		
2.15.1.	Email;	PAN-OS-6.0-Admin-Guide.pdf	pag. 185
2.15.2.	SNMP;	PAN-OS-6.0-Admin-Guide.pdf	pag. 185
2.15.3.	Syslog;	PAN-OS-6.0-Admin-Guide.pdf	pag. 185
2.16.	A plataforma de segurança deve permitir através de API-XML (Application Program Interface) a integração com sistemas existentes no ambiente da contratante de forma a possibilitar que aplicações desenvolvidas na contratante possam interagir em RealTime com a solução possibilitando assim que regras e políticas de segurança de possam ser modificadas por estas aplicações com a utilização de scripts em linguagens de programação como Perl ou PHP.	PAN-OS-6.0-Admin-Guide.pdf	pag. 70