

PAN-OS[®] Administrator's Guide

Version 6.1

Contact Information

Corporate Headquarters:

Palo Alto Networks

4401 Great America Parkway

Santa Clara, CA 95054

www.paloaltonetworks.com/company/contact-support

About this Guide

This guide takes you through the configuration and maintenance of your Palo Alto Networks next-generation firewall. For additional information, refer to the following resources:

- For information on how to configure other components in the Palo Alto Networks Next-Generation Security Platform, go to the Technical Documentation portal: <https://www.paloaltonetworks.com/documentation> or [search](#) the documentation.
- For access to the knowledge base and community forums, refer to <https://live.paloaltonetworks.com>.
- For contacting support, for information on support programs, to manage your account or devices, or to open a support case, refer to <https://www.paloaltonetworks.com/support/tabs/overview.html>.
- For the most current PAN-OS and Panorama 6.1 release notes, go to <https://www.paloaltonetworks.com/documentation/61/pan-os/pan-os-release-notes.html>.

To provide feedback on the documentation, please write to us at: documentation@paloaltonetworks.com.

Palo Alto Networks, Inc.

www.paloaltonetworks.com

© 2014–2017 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at <https://www.paloaltonetworks.com/company/trademarks.html>. All other marks mentioned herein may be trademarks of their respective companies.

Revision Date: May 9, 2017



Table of Contents

Getting Started	13
Integrate the Firewall into Your Management Network	14
Determine Your Management Strategy	14
Perform Initial Configuration	15
Set Up Network Access for External Services	17
Register the Firewall	21
Activate Licenses and Subscriptions	21
Manage Content Updates	23
Install Software Updates	26
Create the Security Perimeter	28
Basic Interface Deployments	28
About Security Policy	30
Plan the Deployment	31
Configure Interfaces and Zones	32
Set Up Basic Security Policies	34
Enable Basic Threat Prevention Features	41
Enable WildFire	41
Scan Traffic for Threats	43
Control Access to Web Content	47
Best Practices for Completing the Firewall Deployment	51
Device Management	53
Management Interfaces	54
Use the Web Interface	55
Use the Command Line Interface (CLI)	60
Use the XML API	62
Manage Firewall Administrators	64
Administrative Roles	64
Administrative Authentication	65
Create an Administrative Account	65
Reference: Web Interface Administrator Access	69
Web Interface Access Privileges	69
Panorama Web Interface Access	105
Reference: Port Numbers Used by Palo Alto Networks Devices	108
Ports Used for Management Functions	108
Ports Used for HA	109
Ports Used for Panorama	110
Ports Used for User-ID	111
Reset the Firewall to Factory Default Settings	112

Certificate Management	113
Keys and Certificates	114
Certificate Revocation	116
Certificate Revocation List (CRL)	116
Online Certificate Status Protocol (OCSP)	117
Certificate Deployment	118
Set Up Verification for Certificate Revocation Status	119
Configure an OCSP Responder	119
Configure Revocation Status Verification of Certificates Used for User/Device Authentication	120
Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption	121
Configure the Master Key	123
Obtain Certificates	124
Create a Self-Signed Root CA Certificate	125
Generate a Certificate on the Firewall	126
Import a Certificate and Private Key	127
Obtain a Certificate from an External CA	128
Configure a Certificate Profile	130
Configure the Key Size for SSL Forward Proxy Server Certificates	133
Revoke and Renew Certificates	134
Revoke a Certificate	134
Renew a Certificate	134
Secure Keys with a Hardware Security Module	135
Set up Connectivity with an HSM	136
Encrypt a Master Key Using an HSM	142
Store Private Keys on an HSM	144
Manage the HSM Deployment	145
High Availability	147
HA Overview	148
HA Concepts	149
HA Modes	149
HA Links and Backup Links	149
Device Priority and Preemption	151
Failover Triggers	152
HA Timers	153
Set Up Active/Passive HA	155
Prerequisites for Active/Passive HA	156
Configuration Guidelines for Active/Passive HA	157
Configure Active/Passive HA	159
Define HA Failover Conditions	165
Verify Failover	166
HA Resources	167

Reports and Logging	169
Use the Dashboard	170
Use the Application Command Center	172
ACC Risk Level	173
ACC Charts	174
ACC Detail Pages	176
Use the ACC	177
Use App-Scope	178
Summary Report	179
Change Monitor Report	180
Threat Monitor Report	182
Threat Map Report	183
Network Monitor Report	184
Traffic Map Report	185
Take Packet Captures	186
Monitor the Firewall	188
Monitor Applications and Threats	189
Monitor Log Data	190
Monitor the Dashboard	194
View Reports	195
Forward Logs to External Services	196
Define Remote Logging Destinations	197
Enable Log Forwarding	204
Log Forwarding Profiles	205
Monitor the Firewall Using SNMP	208
Monitor the Firewall Using NetFlow	210
NetFlow Templates	211
Identify Firewall Interfaces in External Monitoring Systems	214
Manage Reporting	216
About Reports	217
View Reports	218
Disable Predefined Reports	219
Generate Custom Reports	220
Generate Botnet Reports	226
Manage PDF Summary Reports	228
Generate User/Group Activity Reports	230
Manage Report Groups	231
Schedule Reports for Email Delivery	232
Syslog Field Descriptions	233
User-ID	245
User-ID Overview	246
User-ID Concepts	248
Group Mapping	248
User Mapping	248
Enable User-ID	253

Map Users to Groups	254
Map IP Addresses to Users	256
Configure User Mapping Using the Windows User-ID Agent	257
Configure User Mapping Using the PAN-OS Integrated User-ID Agent	263
Configure User-ID to Receive User Mappings from a Syslog Sender	266
Map IP Addresses to User Names Using Captive Portal	277
Configure User Mapping for Terminal Server Users	283
Send User Mappings to User-ID Using the XML API	291
Configure a Firewall to Share User Mapping Data with Other Firewalls	292
Enable User- and Group-Based Policy	294
Verify the User-ID Configuration	297
App-ID	301
App-ID Overview	302
Manage Custom or Unknown Applications	303
Best Practices for Using App-ID in Policy	305
Applications with Implicit Support	306
Application Level Gateways	309
Disable the SIP Application-level Gateway (ALG)	310
Threat Prevention	311
Set Up Security Profiles and Policies	312
Set Up Antivirus, Anti-spyware, and Vulnerability Protection	313
Set Up Data Filtering	315
Set Up File Blocking	319
Prevent Brute Force Attacks	321
Customize the Action and Trigger Conditions for a Brute Force Signature	322
Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions	325
Enable Passive DNS Collection for Improved Threat Intelligence	328
Use DNS Queries to Identify Infected Hosts on the Network	329
DNS Sinkholing	330
Configure DNS Sinkholing	331
Identify Infected Hosts	335
Content Delivery Network Infrastructure for Dynamic Updates	337
Threat Prevention Resources	338
Decryption	339
Decryption Overview	340
Decryption Concepts	341
Keys and Certificates for Decryption Policies	342
SSL Forward Proxy	344
SSL Inbound Inspection	346
SSH Proxy	347
Decryption Exceptions	348
Decryption Port Mirroring	349

Configure SSL Forward Proxy	350
Configure SSL Inbound Inspection	355
Configure SSH Proxy	357
Configure Decryption Exceptions	358
Exclude Traffic From Decryption	358
Exclude a Server From Decryption	359
Enable Users to Opt Out of SSL Decryption	360
Configure Decryption Port Mirroring	362
Temporarily Disable SSL Decryption	365
URL Filtering	367
URL Filtering Overview	368
URL Filtering Vendors	368
Interaction Between App-ID and URL Categories	369
URL Filtering Concepts	370
URL Categories	371
URL Filtering Profile	372
URL Filtering Profile Actions	373
Block and Allow Lists	374
Safe Search Enforcement	375
Container Pages	377
HTTP Header Logging	378
URL Filtering Response Pages	379
URL Category as Policy Match Criteria	381
PAN-DB Categorization Workflow	382
PAN-DB URL Categorization Components	382
PAN-DB URL Categorization Workflow	384
Enable a URL Filtering Vendor	385
Enable PAN-DB URL Filtering	386
Enable BrightCloud URL Filtering	387
Determine URL Filtering Policy Requirements	389
Monitor Web Activity	391
Interpret the URL Filtering Logs	392
Use the ACC to Monitor Web Activity	394
View URL Filtering Reports	395
View the User Activity Report	396
Configure Custom URL Filtering Reports	398
Configure URL Filtering	399
Log Client IP Addresses for Proxied HTTP/HTTPS Requests	401
Use XFF Values for Policies and Logging Source Users	401
Add XFF Values to URL Filtering Logs	402
Customize the URL Filtering Response Pages	403
Configure URL Admin Override	404

Enable Safe Search Enforcement	406
Block Search Results that are not Using Strict Safe Search Settings	406
Enable Transparent Safe Search Enforcement	410
URL Filtering Use Case Examples	413
Use Case: Control Web Access	414
Use Case: Use URL Categories in Policy	418
Troubleshoot URL Filtering	420
Problems Activating PAN-DB	420
PAN-DB Cloud Connectivity Issues	421
URLs Classified as Not-Resolved	422
Incorrect Categorization	423
URL Database Out of Date	424
Quality of Service	425
QoS Overview	426
QoS Concepts	428
QoS for Applications and Users	428
QoS Profile	429
QoS Classes	430
QoS Policy	431
QoS Egress Interface	432
QoS Cleartext and Tunneled Traffic	433
Configure QoS	434
Configure QoS for a Virtual System	439
QoS Use Case Examples	446
Use Case: QoS for a Single User	447
QoS for Voice and Video Applications	450
VPNs	453
VPN Deployments	454
Site-to-Site VPN Overview	455
Site-to-Site VPN Concepts	456
IKE Gateway	456
Tunnel Interface	456
Tunnel Monitoring	457
Internet Key Exchange (IKE) for VPN	457
Set Up Site-to-Site VPN	461
Set up an IKE Gateway	461
Define Cryptographic Profiles	463
Set up an IPSec Tunnel	465
Set up Tunnel Monitoring	468
Test VPN Connectivity	470
Interpret VPN Error Messages	470
Site-to-Site VPN Quick Configs	472
Site-to-Site VPN with Static Routing	472
Site-to-Site VPN with OSPF	476
Site-to-Site VPN with Static and Dynamic Routing	481

Large Scale VPN (LSVPN)	489
LSVPN Overview	490
Create Interfaces and Zones for the LSVPN	491
Enable SSL Between GlobalProtect LSVPN Components	493
About Certificate Deployment	493
Deploy Server Certificates to the GlobalProtect LSVPN Components	493
Configure the Portal to Authenticate Satellites	497
Configure GlobalProtect Gateways for LSVPN	499
Prerequisite Tasks	499
Configure the Gateway	499
Configure the GlobalProtect Portal for LSVPN	502
Prerequisite Tasks	502
Configure the Portal	503
Define the Satellite Configurations	504
Prepare the Satellite Device to Join the LSVPN	507
Verify the LSVPN Configuration	510
LSVPN Quick Configs	511
Basic LSVPN Configuration with Static Routing	512
Advanced LSVPN Configuration with Dynamic Routing	515
 Networking	 519
Interface Deployments	520
Virtual Wire Deployments	520
Layer 2 Deployments	523
Layer 3 Deployments	523
Tap Mode Deployments	524
Virtual Routers	525
Static Routes	527
RIP	529
OSPF	531
OSPF Concepts	531
Configure OSPF	534
Configure OSPFv3	539
Configure OSPF Graceful Restart	542
Confirm OSPF Operation	543
BGP	546
Session Settings and Timeouts	551
Transport Layer Sessions	552
TCP	552
UDP	554
ICMP	555
Configure Session Timeouts	557
Configure Session Settings	559

DHCP	562
DHCP Overview	562
DHCP Messages	563
DHCP Addressing	563
DHCP Options	565
Restrictions to Palo Alto Networks Implementation of DHCP	566
Configure an Interface as a DHCP Server	566
Configure an Interface as a DHCP Client	570
Configure an Interface as a DHCP Relay Agent	571
Monitor and Troubleshoot DHCP	571
NAT	573
Purpose of NAT	573
NAT Rules and Security Policies	573
Source NAT and Destination NAT	574
NAT Rule Capacities	575
Dynamic IP and Port NAT Oversubscription	576
Dataplane NAT Memory Statistics	577
Configure NAT	579
NAT64	587
NAT64 Overview	588
IPv4-Embedded IPv6 Address	589
DNS64 Server	590
Path MTU Discovery	591
Configure NAT64	592
LACP	601
LACP Settings	601
Configure LACP	604
Policy	607
Policy Types	608
Security Policy	609
Components of a Security Policy Rule	610
Security Policy Best Practices	613
Policy Objects	614
Security Profiles	615
Antivirus Profiles	616
Anti-spyware Profiles	617
Vulnerability Protection Profiles	618
URL Filtering Profiles	619
Data Filtering Profiles	620
File Blocking Profiles	622
DoS Protection Profile	623
Zone Protection Profiles	624
Security Profile Group	625
Enumeration of Rules Within a Rulebase	629
Use Tags to Visually Distinguish Objects	631

Table of Contents

Tips and Tricks for Searching for Objects in Security Policy.....	633
Search Security Policy Rules for Security Profiles.....	633
Search for Disabled Rules in the Security Policy.....	637
Search for Log Forwarding Details in Security Policy.....	637
Search for Log Details in Security Policy.....	638
Search for Scheduled Rules in Security Policy.....	639
Use a Dynamic Block List in Policy.....	640
View the IP Address Limit For Your Firewall Model.....	640
Formatting Guidelines for Dynamic Block Lists.....	641
Enforce Policy with a Dynamic Block List.....	641
View the List of IP addresses in the Dynamic Block List.....	642
Retrieve a Dynamic Block List from Web Server.....	643
Register IP Addresses and Tags Dynamically.....	644
Monitor Changes in the Virtual Environment.....	645
Enable VM Monitoring to Track Changes on the Virtual Network.....	646
Attributes Monitored in the AWS and VMware Environments.....	649
Use Dynamic Address Groups in Policy.....	650
CLI Commands for Dynamic IP Addresses and Tags.....	653
Policy-Based Forwarding.....	655
PBF.....	656
Create a Policy-Based Forwarding Rule.....	658
Use Case: PBF for Outbound Access with Dual ISPs.....	660
DoS Protection Against Flooding of New Sessions.....	668
DoS Protection Against Flooding of New Sessions.....	669
Configure DoS Protection Against Flooding of New Sessions.....	673
Use the CLI to End a Single Attacking Session.....	676
Identify Sessions That Use an Excessive Percentage of the Packet Buffer.....	677
Discard a Session Without a Commit.....	679
Virtual Systems.....	681
Virtual Systems Overview.....	682
Virtual System Components and Segmentation.....	682
Benefits of Virtual Systems.....	683
Use Cases for Virtual Systems.....	683
Platform Support and Licensing for Virtual Systems.....	684
Restrictions on Virtual Systems.....	684
Administrative Roles for Virtual Systems.....	684
Shared Objects for Virtual Systems.....	684
Communication Between Virtual Systems.....	685
Inter-VSYS Traffic That Must Leave the Firewall.....	686
Inter-VSYS Traffic That Remains Within the Firewall.....	687
Inter-VSYS Communication Uses Two Sessions.....	690
Shared Gateway.....	691
External Zones and Shared Gateway.....	691
Networking Considerations for a Shared Gateway.....	692
Configure Virtual Systems.....	693
Configure Inter-Virtual System Communication within the Firewall.....	696

Configure a Shared Gateway 697
Virtual System Functionality with Other Features 698



Getting Started

The following sections provide detailed steps to help you deploy a new Palo Alto Networks next-generation firewall. They provide details for integrating a new firewall into your network and configuring basic security policies and threat prevention features.

After you perform the basic configuration steps required to integrate the firewall into your network, you can use the rest of the topics in this guide to help you deploy the comprehensive enterprise security platform features as necessary to address your network security needs.

- ▲ [Integrate the Firewall into Your Management Network](#)
- ▲ [Create the Security Perimeter](#)
- ▲ [Enable Basic Threat Prevention Features](#)
- ▲ [Best Practices for Completing the Firewall Deployment](#)

Integrate the Firewall into Your Management Network

All Palo Alto Networks firewalls provide an out-of-band management port (MGT) that you can use to perform the firewall administration functions. By using the MGT port, you separate the management functions of the firewall from the data processing functions, safeguarding access to the firewall and enhancing performance. When using the web interface, you must perform all initial configuration tasks from the MGT port even if you plan to use an in-band port for managing your device going forward.

Some management tasks, such as retrieving licenses and updating the threat and application signatures on the firewall require access to the Internet. If you do not want to enable external access to your MGT port, you will need to either set up a data port to provide access to required external services or plan to manually upload updates regularly.

The following topics describe how to perform the initial configuration steps that are necessary to integrate a new firewall into the management network and deploy it in a basic security configuration.

- ▲ [Determine Your Management Strategy](#)
- ▲ [Perform Initial Configuration](#)
- ▲ [Set Up Network Access for External Services](#)
- ▲ [Register the Firewall](#)
- ▲ [Activate Licenses and Subscriptions](#)
- ▲ [Manage Content Updates](#)
- ▲ [Install Software Updates](#)



The following topics describe how to integrate a single Palo Alto Networks next-generation firewall into your network. However, for redundancy, consider deploying a pair of firewalls in a [High Availability](#) configuration.

Determine Your Management Strategy


The Palo Alto Networks firewall can be configured and managed locally or it can be managed centrally using [Panorama](#), the Palo Alto Networks centralized security management system. If you have six or more firewalls deployed in your network, use Panorama to achieve the following benefits:



- Reduce the complexity and administrative overhead in managing configuration, policies, software and dynamic content updates. Using device groups and templates on Panorama, you can effectively manage device specific configuration locally on a device and enforce shared policies across all devices or device groups.
- Aggregate data from all managed firewalls and gain visibility across all the traffic on your network. The Application Command Center (**ACC**) on Panorama provides a single glass pane for unified reporting across all the firewalls, allowing you to centrally analyze, investigate and report on network traffic, security incidents and administrative modifications.

The procedures in this document describe how to manage the firewall using the local web interface. If you want to use Panorama for centralized management, after you complete the instructions in the [Perform Initial Configuration](#) section of this guide and verify that the firewall can establish a connection to Panorama. From that point on you can use [Panorama](#) to configure your firewall centrally.

Perform Initial Configuration

By default, the firewall has an IP address of 192.168.1.1 and a username/password of admin/admin. For security reasons, you must change these settings before continuing with other firewall configuration tasks. You must perform these initial configuration tasks either from the MGT interface, even if you do not plan to use this interface for your firewall management, or using a direct serial connection to the console port on the device.

Set Up Network Access to the Firewall	
<p>Step 1 Gather the required information from your network administrator.</p>	<ul style="list-style-type: none"> <input type="checkbox"/> IP address for MGT port <input type="checkbox"/> Netmask <input type="checkbox"/> Default gateway <input type="checkbox"/> DNS server address
<p>Step 2 Connect your computer to the firewall.</p>	<p>You can connect to the firewall in one of the following ways:</p> <ul style="list-style-type: none"> • Connect a serial cable from your computer to the Console port and connect to the firewall using terminal emulation software (9600-8-N-1). Wait a few minutes for the boot-up sequence to complete; when the device is ready, the prompt changes to the name of the firewall, for example <code>PA-500 login</code>. • Connect an RJ-45 Ethernet cable from your computer to the MGT port on the firewall. From a browser, go to <code>https://192.168.1.1</code>. Note that you may need to change the IP address on your computer to an address in the 192.168.1.0 network, such as 192.168.1.2, in order to access this URL.
<p>Step 3 When prompted, log in to the firewall.</p>	<p>You must log in using the default username and password (admin/admin). The firewall will begin to initialize.</p>
<p>Step 4 Configure the MGT interface.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Management and then edit the Management Interface Settings. 2. Enter the IP Address, Netmask, and Default Gateway. 3. Set the Speed to auto-negotiate. 4. Select which management services to allow on the interface. <ul style="list-style-type: none">  Make sure Telnet and HTTP are not selected because these services use plaintext and are not as secure as the other services. 5. Click OK.

Set Up Network Access to the Firewall (Continued)	
<p>Step 5 (Optional) Configure general firewall settings.</p>	<ol style="list-style-type: none"> 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 the Latitude and Longitude to enable accurate placement of the firewall on the world map. 4. Click OK.
<p>Step 6 Configure DNS, time and date settings.</p> <p> 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.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Services and edit the Services. 2. On the Services tab, enter the IP address of the Primary DNS Server and optionally a Secondary DNS Server. 3. To use the virtual cluster of time servers on the Internet, enter the hostname pool.ntp.org as the Primary NTP Server or add the IP address of your Primary NTP Server and optionally a Secondary NTP Server. 4. To authenticate time updates from an NTP server, select the NTP tab, enter the NTP Server Address, and select the Authentication Type for the firewall to use. 5. Click OK to save your settings.
<p>Step 7 Set a secure password for the admin account.</p>	<ol style="list-style-type: none"> 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.
<p>Step 8 Commit your changes.</p> <p> When the configuration changes are saved, you will lose connectivity to the web interface because the IP address will have changed.</p>	<p>Click Commit. The device may take up to 90 seconds to save your changes.</p> 
<p>Step 9 Connect the firewall to your network.</p>	<ol style="list-style-type: none"> 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.
<p>Step 10 Open an SSH management session to the firewall.</p>	<p>Using a terminal emulation software, such as PuTTY, launch an SSH session to the firewall using the new IP address you assigned to it.</p>

Set Up Network Access to the Firewall (Continued)

Step 11 Verify network access to external services required for firewall management, such as the Palo Alto Networks Update Server.

You can do this in one of the following ways:

- If you do not want to allow external network access to the MGT interface, you will need to set up a data port to retrieve required service updates. Continue to [Set Up Network Access for External Services](#).
- If you do plan to allow external network access to the MGT interface, verify that you have connectivity and then proceed to [Register the Firewall](#) and [Activate Licenses and Subscriptions](#).

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

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



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

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

```
request support check
```

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

```
Contact Us
https://www.paloaltonetworks.com/company/contact-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 a data port to provide access to these required external services.




This task requires familiarity with firewall interfaces, zones, and policies. For more information on these topics, see [Create the Security Perimeter](#).

Set Up a Data Port for Access to External Services	
Step 1 Decide which port you want to use for access to external services and connect it to your switch or router port.	The interface you use will need to have a static IP address.
Step 2 Log in to the web interface.	Using a secure connection (https) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>). You will see a certificate warning; that is okay. Continue to the web page.
Step 3 (Optional) The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and zones). If you do not plan to use this virtual wire configuration, you must manually delete the configuration to prevent it from interfering with other interface settings you define.	<p>You must delete the configuration in the following order:</p> <ol style="list-style-type: none"> 1. To delete the default security policy, select Policies > Security, select the rule, and click Delete. 2. Next, delete the default virtual wire by selecting Network > Virtual Wires, selecting the virtual wire and clicking Delete. 3. To delete the default trust and untrust zones, select Network > Zones, select each zone and click Delete. 4. Finally, delete the interface configurations by selecting Network > Interfaces and then select each interface (ethernet1/1 and ethernet1/2) and click Delete. 5. Commit the changes.
Step 4 Configure the interface.	<ol style="list-style-type: none"> 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, define a Name for new zone, for example L3-trust, 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. 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. These services provide management access to the device, so only select the services that correspond to the management activities you want to allow on this interface. For example, if you plan to use the MGT interface for device configuration tasks through the web interface or CLI, you would not want to enable HTTP, HTTPS, SSH, or Telnet so that you could prevent unauthorized access through this interface. For the purposes of allowing access to the external services you probably only need to enable Ping and then click OK. 8. To save the interface configuration, click OK.

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

Step 5 Because the firewall uses the MGT interface by default to access the external services it requires, you must change the interface the firewall uses to send these requests by editing the service routes.

1. Select **Device > Setup > Services > Service Route Configuration**.



For the purposes of activating your licenses and getting the most recent content and software updates, you will want to change the service route for **DNS, Palo Alto Updates, URL Updates, and WildFire**.

2. Click the **Customize** radio button, and select one of the following:
 - For a predefined service, select **IPv4** or **IPv6** and click the link for the service for which you want to modify the **Source Interface** and select the interface you just configured.

If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you select an IP address.
 - To create a service route for a custom destination, select **Destination**, and click **Add**. Enter a **Destination** name and select a **Source Interface**. If more than one IP address is configured for the selected interface, the **Source Address** drop-down allows you select an IP address.
3. Click **OK** to save the settings.
4. Repeat steps 2-3 above for each service route you want to modify.
5. **Commit** your changes.

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

Step 6 Configure an external-facing interface and an associated zone and then create security and NAT policy rules to allow the firewall to send service requests from the internal zone to the external zone.

1. Select **Network > Interfaces** and then select your external-facing interface. Select **Layer3** as the **Interface Type**, **Add** the **IP** address (on the **IPv4** or **IPv6** tab), and create the associated **Security Zone** (on the **Config** tab), such as l3-untrust. You do not need to set up management services on this interface.
2. To set up a security rule that allows traffic from your internal network to the Palo Alto Networks update server and external DNS servers, select **Policies > Security** and click **Add**. For the purposes of initial configuration, you can create a simple rule that allows all traffic from l3-trust to l3-untrust as follows:

Name	Tag	Source				Destination				Action
		Zone	Address	User	HIP Profile	Zone	Address	Application	Service	
rule1	none	l3-trust	any	any	any	l3-untrust	any	any	any	

3. If you are using a private IP address on the internal-facing interface, you will need to create a **Source NAT** rule to translate the address to a publicly routable address. Select **Policies > NAT** and then click **Add**. At a minimum you must define a name for the rule (**General** tab), specify a source and destination zone, l3-trust to l3-untrust in this case (**Original Packet** tab), and define the source address translation settings (**Translated Packet** tab) and then click **OK**.

Name	Tag	Original Packet					
		Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service
source NAT	none	l3-trust	l3-untrust	any	any	any	any

4. **Commit** your changes.

Step 7 Verify that you have connectivity from the data port to the external services, including the default gateway, DNS server, and the Palo Alto Networks Update Server.

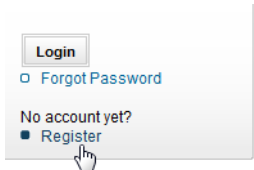
After you verify you have the required network connectivity, continue to [Register the Firewall and Activate Licenses and Subscriptions](#).

Launch the CLI and use the ping utility to verify that you have connectivity. Keep in mind that by default pings are sent from the MGT interface, so in this case you must specify the source interface for the ping requests as follows:

```
admin@PA-200> ping source 192.168.1.254 host
updates.paloaltonetworks.com
PING updates.paloaltonetworks.com (67.192.236.252) from
192.168.1.254 : 56(84) bytes of data.
64 bytes from 67.192.236.252: icmp_seq=1 ttl=242 time=56.7 ms
64 bytes from 67.192.236.252: icmp_seq=2 ttl=242 time=47.7 ms
64 bytes from 67.192.236.252: icmp_seq=3 ttl=242 time=47.6 ms
^C
```

After you have verified connectivity, press Ctrl+C to stop the pings.

Register the Firewall

Register the Firewall	
Step 1 Log in to the web interface.	Using a secure connection (https) from your web browser, log in using the new IP address and password you assigned during initial configuration (https://<IP address>). You will see a certificate warning; that is okay. Continue to the web page.
Step 2 Locate your serial number and copy it to the clipboard.	On the Dashboard , locate your Serial Number in the General Information section of the screen.
Step 3 Go to the Palo Alto Networks Support site.	In a new browser tab or window, go to https://support.paloaltonetworks.com .
Step 4 Register the device. The way you register depends on whether you already have a login to the support site.	<ul style="list-style-type: none"> If this is the first Palo Alto Networks device you are registering and you do not yet have a login, click Register on the right side of the page. To register, you must provide your sales order number or customer ID, and the serial number of your firewall (which you can paste from your clipboard) or the authorization code you received with your order. You will also be prompted to set up a username and password for access to the Palo Alto Networks support community.  <ul style="list-style-type: none"> If you already have a support account, log in and then click My Devices. Scroll down to Register Device section at the bottom of the screen and enter the serial number of your firewall (which you can paste from your clipboard), your city and postal code and then click Register Device.

Activate Licenses and Subscriptions


Before you can start using your firewall to secure the traffic on your network, you must activate the licenses for each of the services you purchased. Available licenses and subscriptions include the following:

- **Threat Prevention**—Provides antivirus, anti-spyware, and vulnerability protection.
- **Decryption Port Mirroring**—Provides the ability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures—such as NetWitness or Solera—for archiving and analysis.
- **URL Filtering**—In order to create policy rules based on dynamic URL categories, you must purchase and install a subscription for one of the supported URL filtering databases: PAN-DB or BrightCloud. For more information about URL filtering, see [Control Access to Web Content](#).

- **Virtual Systems**—This license is required to enable support for multiple virtual systems on PA-2000 and PA-3000 Series firewalls. In addition, you must purchase a Virtual Systems license if you want to increase the number of virtual systems beyond the base number provided by default on PA-4000 Series, PA-5000 Series, and PA-7050 firewalls (the base number varies by platform). The PA-500, PA-200, and VM-Series firewalls do not support virtual systems.
- **WildFire**—Although basic WildFire support is included as part of the Threat Prevention license, the WildFire subscription service provides enhanced services for organizations that require immediate coverage for threats, enabling sub-hourly WildFire signature updates, advanced file type forwarding (APK, PDF, Microsoft Office, and Java Applet), as well as the ability to upload files using the WildFire API. A WildFire subscription is also required if your firewalls will be forwarding files to a private WF-500 WildFire appliance.
- **GlobalProtect**—Provides mobility solutions and/or large-scale VPN capabilities. By default, you can deploy a single GlobalProtect portal and gateway (without HIP checks) without a license. However, if you want to deploy multiple gateways, you must purchase a portal license (one-time, permanent license). If you want to use host checks you will also need gateway licenses (subscription) for each gateway.

Activate Licenses	
<p>Step 1 Locate the activation codes for the licenses you purchased.</p>	<p>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.</p>
<p>Step 2 Launch the web interface and go to the license page.</p>	<p>Select Device > Licenses.</p>
<p>Step 3 Activate each license you purchased.</p>	<p>After purchasing your licenses/subscriptions activate them in one of the following ways:</p> <ul style="list-style-type: none"> • Retrieve license keys from license server—Use this option if you activated your license on the support portal. • Activate feature using authorization code—Use this option to enable purchased subscriptions using an authorization code for licenses that have not been previously activated on the support portal. When prompted, enter the Authorization Code and then click OK. • Manually upload license key—Use this option if your device does not have connectivity to the Palo Alto Networks support site. In this case, you must download a license key file from the support site on an Internet connected computer and then upload to the device.
<p>Step 4 Verify that the license was successfully activated</p>	<p>On the Device > Licenses page, verify that the license was successfully activated. For example, after activating the WildFire license, you should see that the license is valid:</p> <div data-bbox="678 1675 1198 1801" style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>WildFire License</p> <p>Date Issued September 11, 2012</p> <p>Date Expires September 11, 2015</p> <p>Description WildFire License</p> </div>

Activate Licenses (Continued)

<p>Step 5 (WildFire subscriptions only) Perform a commit to complete WildFire subscription activation.</p>	<p>After activating a WildFire subscription, a commit is required for the firewall to begin forwarding advanced file types:</p> <ul style="list-style-type: none"> • Commit any pending changes. • Make a minor change and perform a commit. For example, update a rule description and commit the change. <p> Check the file blocking profile rule used for WildFire forwarding includes the advanced file types that are now supported with the WildFire subscription. If the rule is already set to forward Any file type, or no change to the rule is required, make a minor edit to the rule description and perform a commit.</p>
---	---

Manage Content Updates

In order to stay ahead of the changing threat and application landscape, Palo Alto Networks maintains a Content Delivery Network (CDN) infrastructure for delivering content updates to the Palo Alto Networks devices. The devices access the web resources in the CDN to perform various App-ID and Content-ID functions. By default, the devices use the management port to access the CDN infrastructure for application updates, threat and antivirus signature updates, BrightCloud and PAN-DB database updates and lookups, and access to the Palo Alto Networks WildFire Cloud. To ensure that you are always protected from the latest threats (including those that have not yet been discovered), you must ensure that you keep your devices up-to-date with the latest updates published by Palo Alto Networks.

The following content updates are available, depending on which subscriptions you have:



Although you can manually download and install content updates at any time, as a best practice you should [Schedule each update](#). Scheduled updates occur automatically.

- **Antivirus**—Includes new and updated antivirus signatures, including signatures discovered by the WildFire cloud service. You must have a Threat Prevention subscription to get these updates. New antivirus signatures are published daily.
- **Applications**—Includes new and updated application signatures. This update does not require any additional subscriptions, but it does require a valid maintenance/support contract. New application updates are published weekly.
- **Applications and Threats**—Includes new and updated application and threat signatures. This update is available if you have a Threat Prevention subscription (and you get it instead of the Applications update). New Applications and Threats updates are published weekly.
- **GlobalProtect Data File**—Contains the vendor-specific information for defining and evaluating host information profile (HIP) data returned by GlobalProtect agents. You must have a GlobalProtect portal and GlobalProtect gateway license in order to receive these updates. In addition, you must create a schedule for these updates before GlobalProtect will function.

- **BrightCloud URL Filtering**—Provides updates to the BrightCloud URL Filtering database only. You must have a BrightCloud subscription to get these updates. New BrightCloud URL database updates are published daily. If you have a PAN-DB license, scheduled updates are not required as devices remain in-sync with the servers automatically.
- **WildFire**—Provides near real-time malware and antivirus signatures created as a result of the analysis done by the WildFire cloud service. Without the subscription, you must wait 24 to 48 hours for the signatures to roll into the Applications and Threat update.



If your firewall does not have Internet access from the management port, you can download content updates from the [Palo Alto Networks Support](#) portal and then Upload them to your firewall.

If your firewall is deployed behind existing firewalls or proxy servers, access to these external resources might be restricted using access control lists that allow the firewall to only access a hostname or an IP address. In such cases, to allow access to the CDN, set the update server address to use the hostname `staticupdates.paloaltonetworks.com` or the IP address `199.167.52.15`.

Download the Latest Databases	
<p>Step 1 Verify that the firewall points to the CDN infrastructure.</p>	<p>Select Device > Setup > Services.</p> <ul style="list-style-type: none"> • As a best practice, set the Update Server to access <code>updates.paloaltonetworks.com</code>. This allows the firewall to receive content updates from the server to which it is closest in the CDN infrastructure. • (Optional) If the firewall has restricted access to the Internet, set the update server address to use the hostname <code>staticupdates.paloaltonetworks.com</code> or the IP address <code>199.167.52.15</code>. • For additional security, select Verify Update Server Identity. The firewall verifies that the server from which the software or content package is download has an SSL certificate signed by a trusted authority.
<p>Step 2 Launch the web interface and go to the Dynamic Updates page.</p>	<p>Select Device > Dynamic Updates.</p>

Download the Latest Databases (Continued)

Step 3 Check for the latest updates.

Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates. The link in the **Action** column indicates whether an update is available:

- **Download**—Indicates that a new update file is available. Click the link to begin downloading the file directly to the firewall. After successful download, the link in the **Action** column changes from **Download** to **Install**.

WildFire		Last checked: 2012/09/11 17:09:37		Schedule: None	
4451-7738	panup-all-wildfire-4451-7738	Full	3 MB	2012/09/11 16:33:15	Download Release Notes



You cannot download the antivirus database until you have installed the Application and Threats database.

- **Upgrade**—Indicates that there is a new version of the BrightCloud database available. Click the link to begin the download and installation of the database. The database upgrade begins in the background; when completed a check mark displays in the **Currently Installed** column. Note that if you are using PAN-DB as your URL filtering database you will not see an upgrade link because the PAN-DB database automatically stays in sync with the server.

URL Filtering		Schedule: Every day at 01:02 (download-and-install)	
3940			Upgrade
3932		✓	



To check the status of an action, click **Tasks** (on the lower right-hand corner of the window).

- **Revert**—Indicates that the corresponding software version has been downloaded previously. You can choose to revert to the previously installed version of the update.

Step 4 Install the updates.



Installation can take up to 20 minutes on a PA-200, PA-500, or PA-2000 device and up to two minutes on a PA-3000 Series, PA-4000 Series, PA-5000 Series, PA-7050, or VM-Series firewall.

Click the **Install** link in the **Action** column. When the installation completes, a check mark displays in the **Currently Installed** column.


WildFire		Last checked: 2012/09/11 17:17:26		Schedule: None	
4451-7738	panup-all-wildfire-4451-7738	Full	3 MB	2012/09/11 16:33:15	Install Release Notes <input checked="" type="checkbox"/>

Download the Latest Databases (Continued)**Step 5** Schedule each update.

Repeat this step for each update you want to schedule.



Stagger the update schedules because the firewall can only download one update at a time. If you schedule the updates to download during the same time interval, only the first download will succeed.

1. Set the schedule of each update type by clicking the **None** link.2. Specify how often you want the updates to occur by selecting a value from the **Recurrence** drop-down. The available values vary by content type (WildFire updates are available **Every 15 minutes**, **Every 30 minutes** or **Every Hour** whereas all other content types can be scheduled for **Daily** or **Weekly** update).3. Specify the **Time** and (or, minutes past the hour in the case of WildFire), if applicable depending on the **Recurrence** value you selected, **Day** of the week that you want the updates to occur.4. Specify whether you want the system to **Download Only** or, as a best practice, **Download And Install** the update.5. In rare instances, errors in content updates may be found. For this reason, you may want to delay installing new updates until they have been released for a certain number of hours. You can specify how long after a release to wait before performing a content update by entering the number of hours to wait in the **Threshold (Hours)** field.6. Click **OK** to save the schedule settings.7. Click **Commit** to save the settings to the running configuration.

Install Software Updates

When installing a new firewall, it is a good idea to [upgrade](#) to the latest software update (or to the update version recommended by your reseller or Palo Alto Networks Systems Engineer) to take advantage of the latest fixes and security enhancements. Note that before updating the software, you should first make sure you have the latest content updates as detailed in the previous section (the [Release Notes](#) for a software update specify the minimum Content Release version supported in the release).

Update PAN-OS

Step 1 Launch the web interface and go to the Software page.

Select **Device > Software**.

Step 2 Check for software updates.

Click **Check Now** to check for the latest updates. If the value in the **Action** column is **Download** it indicates that an update is available.

Update PAN-OS (Continued)**Step 3** Download the update.

If your firewall does not have Internet access from the management port, you can download the software update from the [Palo Alto Networks Support](#) portal. You can then manually **Upload** them to your firewall.


Locate the version you want and then click **Download**. When the download completes, the value in the Action column changes to **Install**.

Version	Size	Release Date	Downloaded	Currently Installed	Action
5.0.0	259 MB	2012/11/01 19:58:24	✓		Install
4.1.9	169 MB	2012/11/05 23:40:31			Download
4.1.8	168 MB	2012/09/22 21:01:08	✓	✓	Download
4.1.8-h3	168 MB	2012/10/18 23:49:21			Download
4.1.7	152 MB	2012/07/29 09:30:58			Download

Step 4 Install the update.

1. Click **Install**.
2. Reboot the firewall:

- If you are prompted to reboot, click **Yes**.

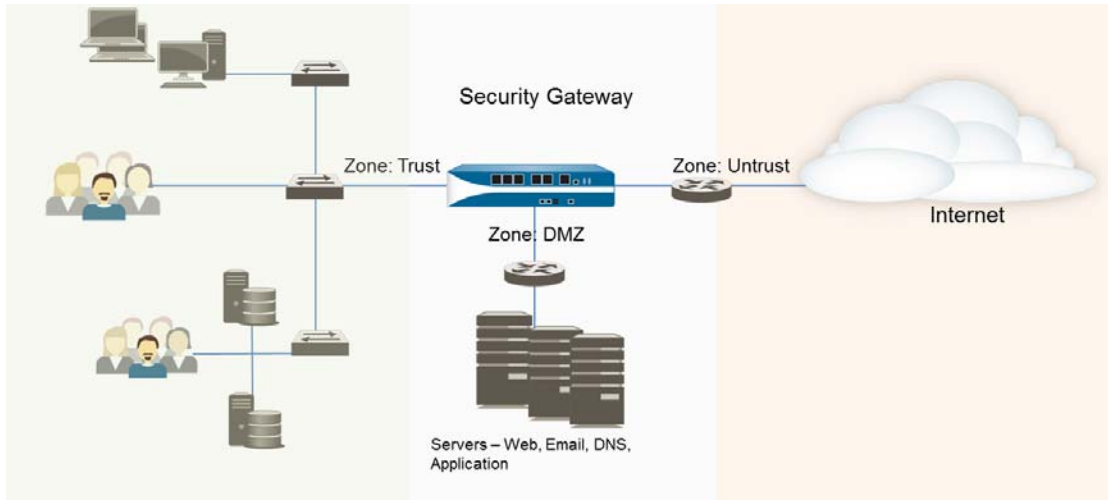


- If you are not prompted to reboot, select **Device > Setup > Operations** and click **Reboot Device** in the Device Operations section of the screen.



Create the Security Perimeter

Traffic must pass through the firewall in order for the firewall to manage and control it. Physically, traffic enters and exits the firewall through *interfaces*. The firewall determines how to act on a packet based on whether the packet matches a *security policy rule*. At the most basic level, each security policy rule must identify where the traffic came from and where it is going. On a Palo Alto Networks next-generation firewall, security policy rules are applied between zones. A *zone* is a grouping of interfaces (physical or virtual) that provides an abstraction for an area of trust for simplified policy enforcement. For example, in the following topology diagram, there are three zones: Trust, Untrust, and DMZ. Traffic can flow freely within a zone, but traffic will not be able to flow between zones until you define a security policy rule that allows it.



The following topics describe the components of the security perimeter and provide steps for configuring the firewall interfaces, defining zones, and setting up a basic security policy that allows traffic from your internal zone to the Internet and to the DMZ. By initially creating a basic security policy rulebase like this, you will be able to analyze the traffic running through your network and use this information to define more granular policies for safely enabling applications while preventing threats.

- ▲ [Basic Interface Deployments](#)
- ▲ [About Security Policy](#)
- ▲ [Plan the Deployment](#)
- ▲ [Configure Interfaces and Zones](#)
- ▲ [Set Up Basic Security Policies](#)

If you use private IP addresses in your internal networks, you will also need to configure network address translation (NAT); see [Networking](#) for NAT concepts and configuration tasks.

Basic Interface Deployments

All Palo Alto Networks next-generation firewalls provide a flexible networking architecture that includes support for dynamic routing, switching, and VPN connectivity, enabling you to deploy the firewall into nearly any networking environment. When configuring the Ethernet ports on your firewall, you can choose from

virtual wire, Layer 2, or Layer 3 interface deployments. In addition, to allow you to integrate into a variety of network segments, you can configure different types of interfaces on different ports. The following sections provide basic information on each type of deployment.

- ▲ [Virtual Wire Deployments](#)
- ▲ [Layer 2 Deployments](#)
- ▲ [Layer 3 Deployments](#)

For more detailed deployment information, refer to [Designing Networks with Palo Alto Networks Firewalls](#).

Virtual Wire Deployments

In a virtual wire deployment, the firewall is installed transparently on a network segment by binding two ports together. By using a virtual wire, you can install the firewall in any network environment without reconfiguring adjacent devices. If necessary, a virtual wire can block or allow traffic based on the virtual LAN (VLAN) tag values. You can also create multiple subinterfaces and classify traffic according to an IP Address (address, range, or subnet), VLAN, or a combination of the two.

By default, the virtual wire (named **default-vwire**) binds Ethernet ports 1 and 2 and allows all untagged traffic. Choose this deployment to simplify installation and configuration and/or avoid configuration changes to surrounding network devices.

A virtual wire is the default configuration, and should be used only when no switching or routing is needed. If you do not plan to use the default virtual wire, you should manually delete the configuration before proceeding with interface configuration to prevent it from interfering with other interface settings you define. For instructions on how to delete the default virtual wire and its associated security policy and zones, see [Step 3 in Set Up a Data Port for Access to External Services](#).

Layer 2 Deployments

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

For more information on Layer 2 deployments, refer to the [Layer 2 Networking Tech Note](#) and/or the [Securing Inter VLAN Traffic Tech Note](#).

Layer 3 Deployments

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

You must assign an IP address to each physical Layer 3 interface you configure. You can also create logical subinterfaces for each physical Layer 3 interface that allows you to segregate the traffic on the interface based on VLAN tag (when VLAN trunking is in use) or by IP address, for example for multi-tenancy.

In addition, because the firewall must route traffic in a Layer 3 deployment, you must configure a virtual router. You can configure the virtual router to participate with dynamic routing protocols (BGP, OSPF, or RIP) as well as adding static routes. You can also create multiple virtual routers, each maintaining a separate set of routes that are not shared between virtual routers, enabling you to configure different routing behaviors for different interfaces.

The configuration example in this chapter illustrates how to integrate the firewall into your Layer 3 network using static routes. For information on other types of routing integrations, refer to the following documents:

- [How to Configure OSPF Tech Note](#)
- [How to Configure BGP Tech Note](#)

About Security Policy

Security Policy protects network assets from threats and disruptions and aids in optimally allocating network resources for enhancing productivity and efficiency in business processes. On the Palo Alto Networks firewall, security policy rules determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service.

For traffic that doesn't match any defined rules, the default rules apply. The default rules—displayed at the bottom of the security rulebase—are predefined to allow all intrazone (within the zone) traffic and deny all interzone (between zones) traffic. Although these rules are part of the pre-defined configuration and are read-only by default, you can override them and change a limited number of settings, including the tags, action (allow or deny), log settings, and security profiles.

Security policies rules are evaluated left to right and from top to bottom. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated. Therefore, the more specific rules must precede more generic ones in order to enforce the best match criteria. Traffic that matches a rule generates a log entry at the end of the session in the traffic log, if logging is enabled for that rule. The logging options are configurable for each rule, and can for example be configured to log at the start of a session instead of, or in addition to, logging at the end of a session.

- ▲ [About Policy Objects](#)
- ▲ [About Security Profiles](#)

About Policy Objects

A *policy object* is a single object or a collective unit that groups discrete identities such as IP addresses, URLs, applications, or users. With **Policy Objects** that are a collective unit, you can reference the object in security policy instead of manually selecting multiple objects one at a time. Typically, when creating a policy object, you group objects that require similar permissions in policy. For example, if your organization uses a set of server IP addresses for authenticating users, you can group the set of server IP addresses as an *address group* policy object and reference the address group in the security policy. By grouping objects, you can significantly reduce the administrative overhead in creating policies.

Some examples of address and application policy objects are shown in the security policies that are included in [Create Security Rules](#). For information on the other policy objects, see [Enable Basic Threat Prevention Features](#).

About Security Profiles

While security policies enable you to allow or deny traffic on your network, security profiles help you define an *allow but scan* rule, which scan allowed applications for threats. When traffic matches the allow rule defined in the security policy, the [Security Profiles](#) that are attached to the rule are applied for further content inspection rules such as antivirus checks and data filtering.



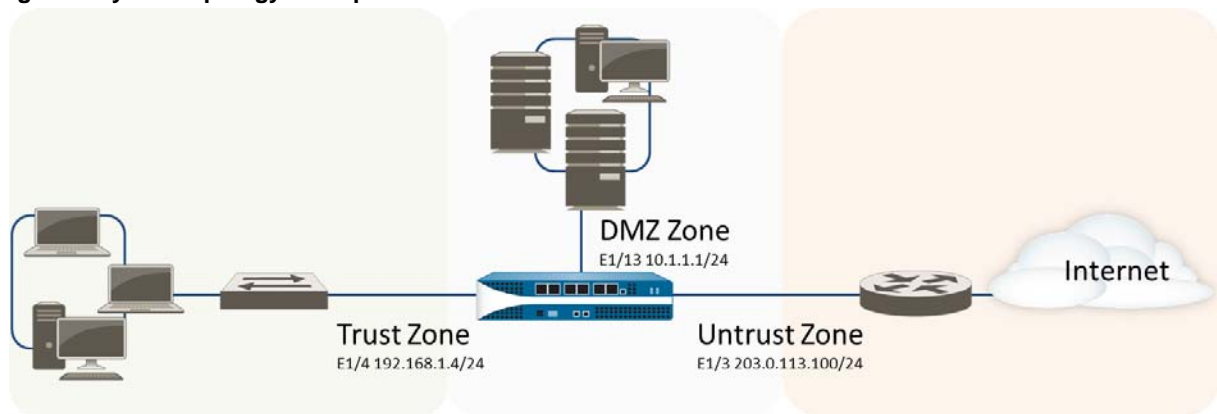
Security profiles are not used in the match criteria of a traffic flow. The security profile is applied to scan traffic after the application or category is allowed by the security policy.

The different types of security profiles that can be attached to security policies are: Antivirus, Anti-spyware, Vulnerability Protection, URL Filtering, File Blocking, and Data Filtering. The firewall provides default security profiles that you can use out of the box to begin protecting your network from threats. See [Create Security Rules](#) for information on using the default profiles in your security policy. As you get a better understanding about the security needs on your network, you can create custom profiles. See [Scan Traffic for Threats](#) for more information.

Plan the Deployment

Before you begin configuring interfaces and zones, take some time to plan the zones you will need based on the different usage requirements within your organization. In addition, you should gather all of the configuration information you will need ahead of time. At a basic level, you should plan which interfaces will belong to which zones. For Layer 3 deployments you'll also need to obtain the required IP addresses and network configuration information from your network administrator, including information on how to configure the routing protocol or static routes required for the virtual router configuration. The example in this chapter will be based on the following topology:

Figure: Layer 3 Topology Example



The following table shows the information we will use to configure the Layer 3 interfaces and their corresponding zones as shown in the sample topology.

Zone	Deployment Type	Interface(s)	Configuration Settings
Untrust	L3	Ethernet1/3	IP address: 203.0.113.100/24 Virtual router: default Default route: 0.0.0.0/0 Next hop: 203.0.113.1
Trust	L3	Ethernet1/4	IP address: 192.168.1.4/24 Virtual router: default
DMZ	L3	Ethernet1/13	IP address: 10.1.1.1/24 Virtual router: default

Configure Interfaces and Zones

After you plan your zones and the corresponding interfaces, you can configure them on the device. The way you configure each interface depends on your network topology.


The following procedure shows how to configure a Layer 3 deployment as depicted in [Figure: Layer 3 Topology Example](#).

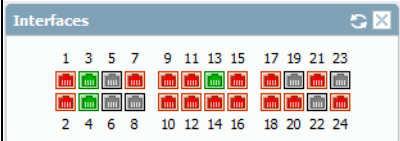


The firewall comes preconfigured with a default virtual wire interface between ports Ethernet 1/1 and Ethernet 1/2 (and a corresponding default security policy and virtual router). If you do not plan to use the default virtual wire, you must manually delete the configuration and commit the change before proceeding to prevent it from interfering with other settings you define. For instructions on how to delete the default virtual wire and its associated security policy and zones, see Step 3 in [Set Up a Data Port for Access to External Services](#).

Set Up Interfaces and Zones	
<p>Step 1 Configure a default route to your Internet router.</p>	<ol style="list-style-type: none"> 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, 203.00.113.1). 4. Click OK twice to save the virtual router configuration.

Set Up Interfaces and Zones (Continued)

<p>Step 2 Configure the external interface (the interface that connects to the Internet).</p>	<ol style="list-style-type: none"> 1. Select Network > Interfaces and then select the interface you want to configure. In this example, we are configuring Ethernet1/3 as the external interface. 2. Select the Interface Type. Although your choice here depends on your network topology, this example shows the steps for Layer3. 3. On the Config tab, select New Zone from the Security Zone drop-down. In the Zone dialog, define a Name for new zone, for example Untrust, and then click OK. 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 208.80.56.100/24. 6. To enable you to ping the interface, select Advanced > Other Info, expand the Management Profile drop-down, and select New Management Profile. Enter a Name for the profile, select Ping and then click OK. 7. To save the interface configuration, click OK.
<p>Step 3 Configure the interface that connects to your internal network.</p> <p> In this example, the interface connects to a network segment that uses private IP addresses. Because private IP addresses cannot be routed externally, you will have to configure NAT.</p>	<ol style="list-style-type: none"> 1. Select Network > Interfaces and select the interface you want to configure. In this example, we are configuring Ethernet1/4 as the internal interface. 2. Select Layer3 from the Interface Type drop-down. 3. On the Config tab, expand the Security Zone drop-down and select New Zone. In the Zone dialog, define a Name for new zone, for example Trust, and then click OK. 4. Select the same Virtual Router you used in Step 2, 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 created in Step 2-6. 7. To save the interface configuration, click OK.

Set Up Interfaces and Zones (Continued)	
<p>Step 4 Configure the interface that connects to the DMZ.</p>	<ol style="list-style-type: none"> 1. Select the interface you want to configure. 2. Select Layer3 from the Interface Type drop-down. In this example, we are configuring Ethernet1/13 as the DMZ interface. 3. On the Config tab, expand the Security Zone drop-down and select New Zone. In the Zone dialog, define a Name for new zone, for example DMZ, and then click OK. 4. Select the Virtual Router you used in Step 2, default in this example. 5. To assign an IP address to the interface, select the IPv4 tab, click Add in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.1.1.1/24. 6. To enable you to ping the interface, select the management profile that you created in Step 2-6. 7. To save the interface configuration, click OK.
<p>Step 5 Save the interface configuration.</p>	<p>Click Commit.</p>
<p>Step 6 Cable the firewall.</p>	<p>Attach straight through cables from the interfaces you configured to the corresponding switch or router on each network segment.</p>
<p>Step 7 Verify that the interfaces are active.</p>	<p>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.</p>  <p>The screenshot shows a window titled 'Interfaces' with a grid of 24 interface status icons arranged in two rows of 12. The top row is labeled with interface numbers 1, 3, 5, 7, 9, 11, 13, 15, 17, 19, 21, 23. The bottom row is labeled with interface numbers 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24. Each icon consists of a small red square with a white symbol, and a green square with a white symbol, representing the link state of each interface.</p>

Set Up Basic Security Policies

Policies allow you to enforce rules and take action. The different types of policy rules that you can create on the firewall are: Security, NAT, Quality of Service (QoS), Policy Based Forwarding (PBF), Decryption, Application Override, Captive Portal, Denial of Service, and Zone protection policies. All these different policies work together to allow, deny, prioritize, forward, encrypt, decrypt, make exceptions, authenticate access, and reset connections as needed to help secure your network. This section covers basic security policies and the default security profiles:

- ▲ [Create Security Rules](#)
- ▲ [Test Your Security Policies](#)
- ▲ [Monitor the Traffic on Your Network](#)

Create Security Rules

Security policies reference security zones and enable you to allow, restrict, and track traffic on your network. Because each zone implies a level of trust, the implicit rule for passing traffic between two different zones is deny, and the traffic within a zone is permitted. To allow traffic between two different zones, you must create a security rule that allows traffic to flow between them.

While setting up the basic framework for securing the enterprise perimeter, it's good idea to start with a simple security policy that allows traffic between the different zones without being too restrictive. As illustrated in the following section, our objective is to minimize the likelihood of breaking applications that users on the network need access to, while providing visibility into the applications and the potential threats for your network.



When defining policies make sure that you do not create a policy that denies all traffic from *any* source zone to *any* destination zone as this will break intra-zone traffic that is implicitly allowed. By default, intra-zone traffic is permitted because the source and destination zones are the same and therefore share the same level of trust.

Define Basic Security Rules

Step 1 Permit Internet access for all users on the enterprise network.

Zone: Trust to Untrust




By default, the firewall includes a security rule named *rule1* that allows all traffic from Trust zone to Untrust zone. You can either delete the rule or modify the rule to reflect your zone-naming convention.

To safely enable applications that are required for day-to-day business operations we will create a simple rule that allows access to the Internet. To provide basic threat protection, we will attach the default security profiles available on the firewall.

1. Select **Policies > Security** and click **Add**.
2. Give the rule a descriptive name in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination** tab, Set the **Destination Zone** to Untrust.



To scan policy rules and visually identify the zones on each rule, create a tag with the same name as the zone. For example, to color code the Trust zone as green, select **Objects > Tags**, click **Add** and **Name** the tag Trust, and select the **Color** green.



5. In the **Service/ URL Category** tab, select service-http and service-https.
6. In the **Actions** tab, complete these tasks:
 - a. Set the **Action Setting** to **Allow**.
 - b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection and URL filtering, under **Profile Setting**.
7. Verify that logging is enabled at the end of a session under **Options**. Only traffic that matches a security rule will be logged.



Step 2 Permit users on the internal network to access the servers in the DMZ.

Zone: Trust to DMZ



If using IP addresses for configuring access to the servers in the DMZ, always refer to the original IP addresses in the packet (i.e. the pre-NAT addresses), and the post-NAT zone.

1. Click **Add** in the **Policies > Security** section.
2. Give the rule a descriptive name in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination** tab, set the **Destination Zone** to DMZ.
5. In the **Service/ URL Category** tab, make sure the **Service** is set to **application-default**.
6. In the **Actions** tab, set the **Action Setting** to Allow.
7. Leave all the other options at the default values.

Name	Zone	User	Zone	Address	Applic...	Service	Action	Profile	Options
Internal network to DMZ s...	Trust	any	DMZ	any	any	application-d...	Allow	none	

Define Basic Security Rules (Continued)

Step 3 Restrict access from the Internet to the servers on the DMZ to specific server IP addresses only.

For example, you might only allow users to access the webmail servers from outside.

Zone: Untrust to DMZ

To restrict inbound access to the DMZ from the Internet, configure a rule that allows access only to specific servers IP addresses and on the default ports that the applications use.

1. Click **Add** to add a new rule, and give it a descriptive name.
2. In the **Source** tab, set the **Source Zone** to Untrust.
3. In the **Destination** tab, set the **Destination Zone** to DMZ.
4. Set the **Destination Address** to the **Public web server** address object you created earlier. The public web server address object references the public IP address—208.80.56.11/24—of the web server that is accessible on the DMZ.
5. Select the webmail application in the **Application** tab.



The **Service** is set to **application-default** by default.


6. Set the **Action Setting** to **Allow**.

Name	Source		Destination		Application	Service	Action	Profile
	Zone	Address	Zone	Address				
Internet to DMZ	untrust	any	DMZ	Public web server	outlook-web	applicatio...	✓	none

Step 4 Allow access from the DMZ to your internal network (Trust zone). To minimize risk, you will allow traffic only between specific servers and destination addresses. For example, if you have an application server on the DMZ that needs to communicate with a specific database server in your Trust zone, create a rule to allow traffic between a specific source to a specific destination.

Zone: DMZ to Trust

1. Click **Add** to add a new rule, and give it a descriptive name.
2. Set the **Source Zone** to DMZ.
3. Set the **Destination Zone** to Trust.
4. Create a an address object that specifies the server(s) on your Trust zone that can be accessed from the DMZ.



5. In the **Destination** tab on the Security Policy rule, set the **Destination Address** to the Address object you created above.
6. In the **Actions** tab, complete these tasks:
 - a. Set the **Action Setting** to **Allow**.
 - b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection, under **Profile Setting**.
 - c. In the Other Settings section, select the option to **Disable Server Response Inspection**. This setting disables the antivirus and anti-spyware scanning on the server-side responses, and thus reduces the load on the firewall.


Name	Source		Destination		Application	Service	Action	Profile	Options
	Zone	Address	Zone	Address					
Internal network to D...	DMZ	any	untrust	Database acce...	any	any	✓	Antivirus, Anti-spyware, Vulnerability Protection	Server Response Inspection (disabled)


Define Basic Security Rules (Continued)

Step 5 Enable the servers on the DMZ to obtain updates and hot fixes from the Internet. Say, for example, you would like to allow the Microsoft Update service.

Zone: DMZ to Untrust

1. Add a new rule and give it a descriptive label.
2. Set the **Source Zone** to DMZ.
3. Set the **Destination Zone** to Untrust.
4. Create an application group to specify the applications that you would like to allow. In this example, we allow Microsoft updates (ms-updates) and dns.



 The **Service** is set to **application-default** by default. This allows the firewall to permit the applications only when they use the standard ports associated with these applications.

5. Set the **Action Setting** to **Allow**.
6. Attach the default profiles for antivirus, anti-spyware, and vulnerability protection, under **Profiles**.

Name	Zone	User	Zone	Address	Application	Service	Action	Profile	Options
Updates- DMZ to Internet	DMZ	any	Untrust	any	dns ms-update	application-d...	Allow	Antivirus, Anti-spyware, Vulnerability Protection	

Step 6 Save your policies to the running configuration on the device.

Click **Commit**.

Test Your Security Policies


To verify that you have set up your basic policies effectively, test whether your security policies are being evaluated and determine which security rule applies to a traffic flow.

Verify Policy Match Against a Flow	
<p>To verify the policy rule that matches a flow, use the following CLI command:</p> <pre>test security-policy-match source <IP_address> destination <IP_address> destination port <port_number> protocol <protocol_number></pre> <p>The output displays the best rule that matches the source and destination IP address specified in the CLI command.</p>	<p>For example, to verify the policy rule that will be applied for a server on the DMZ with the IP address 208.90.56.11 when it accesses the Microsoft update server, you will try the following command:</p> <pre>test security-policy-match source 208.80.56.11 destination 176.9.45.70 destination-port 80 protocol 6</pre> <pre>"Updates-DMZ to Internet" { from dmz; source any; source-region any; to untrust; destination any; destination-region any; user any; category any; application/service[dns/tcp/any/53 dns/udp/any/53 dns/udp/any/5353 ms-update/tcp/any/80 ms-update/tcp/any/443]; action allow; terminal yes;</pre>

Monitor the Traffic on Your Network

Now that you have a basic security policy in place, you can review the statistics and data in the Application Command Center (ACC), traffic logs, and the threat logs to observe trends on your network, to identify where you need to create more granular policies.

Unlike traditional firewalls that use port or protocol to identify applications, the Palo Alto Networks firewalls use the application signature (the [App-ID](#) technology) to monitor applications. The application signature is based on unique application properties and related transaction characteristics in combination with the port or protocol. Therefore, even when the traffic uses the right port/protocol, the firewall can deny access to content because the application signature is not a match. This feature allows you to safely enable applications by allowing parts of the application while blocking or controlling functions within the same application. For example, if you allow the application *web-browsing* a user will be able to access content on the Internet. Then, if a user goes to Facebook and then goes on to play Scrabble on Facebook, the firewall will identify the application shifts and recognize Facebook as an *application* and Scrabble as a *Facebook-app*. Therefore, if you create a specific rule that blocks Facebook applications, the user will be denied access to Scrabble while still being able to access Facebook.

Monitor Network Traffic	
<ul style="list-style-type: none"> • Use the Application Command Center. 	<p>In the ACC, review the most used applications and the high-risk applications on your network. The ACC graphically summarizes the log information to highlight the applications traversing the network, who is using them (with User-ID enabled), and the potential security impact of the content to help you identify what is happening on the network in real time. You can then use this information to create appropriate security policies that block unwanted applications, while allowing and enabling applications in a secure manner.</p>
<ul style="list-style-type: none"> • Determine what updates/modifications are required for your network security rules and implement the changes. 	<p>For example:</p> <ul style="list-style-type: none"> • Evaluate whether to allow content based on schedule, users, or groups • Allow or control certain applications or functions within an application • Decrypt and inspect content • Allow but scan for threats and exploits <p>For information on refining your security policies and for attaching custom security profiles, see Enable Basic Threat Prevention Features.</p>
<ul style="list-style-type: none"> • View the Log Files. 	<p>Specifically, view the traffic and threat logs (Monitor > Logs).</p> <p> Traffic logs are dependent on how your security policies are defined and setup to log traffic. The ACC tab, however, records applications and statistics regardless of policy configuration; it shows all traffic that is allowed on your network, therefore it includes the inter zone traffic that is allowed by policy and the same zone traffic that is allowed implicitly</p>
<ul style="list-style-type: none"> • Interpret the URL Filtering Logs 	<p>Review the URL filtering logs to scan through alerts, denied categories/URL. URL logs are generated when a traffic matches a security rule that has a URL filtering profile attached with an action of alert, continue, override or block.</p>

Enable Basic Threat Prevention Features

The Palo Alto Networks next-generation firewall has unique threat prevention capabilities that allow it to protect your network from attack despite evasive, tunneled, or circumvention techniques. The threat prevention features on the firewall include the WildFire service, the Security Profiles that support Antivirus, Anti-spyware, Vulnerability Protection, URL Filtering, File Blocking and Data Filtering capabilities and the Denial of Service (DoS) and Zone protection functionality.



Before you can apply threat prevention features, you must first configure zones—to identify one or more source or destination interfaces—and security policies. To configure interfaces, zones, and the policies that are needed to apply threat prevention features, see [Configure Interfaces and Zones](#) and [Set Up Basic Security Policies](#).

To begin protecting your network from threats start here:

- ▲ [Enable WildFire](#)
- ▲ [Scan Traffic for Threats](#)
- ▲ [Control Access to Web Content](#)

Enable WildFire




The [WildFire](#) service is included as part of the base product. The WildFire service enables the firewall to forward attachments to a sandbox environment where applications are run to detect any malicious activity. As new malware is detected by the WildFire system, malware signatures are automatically generated and are made available within 24-48 hours in the antivirus daily downloads. Your threat prevention subscription entitles you for antivirus signature updates that include signatures discovered by WildFire.

Consider purchasing the WildFire subscription service for these additional benefits:

- Sub-hourly (as often as every 15 minutes) WildFire signature updates
- Advanced file type forwarding (APK, Flash, PDF, Microsoft Office, and Java Applet)
- Ability to upload files using the WildFire API
- Ability to forward files to a private WF-500 WildFire appliance

While the ability to configure a file blocking profile to forward Portable Executable (PE) files to the WildFire cloud for analysis is free, in order to forward files to a private WildFire appliance, a WildFire subscription is required.

Enable WildFire	
<p>Step 1 Confirm that your device is registered and that you have a valid support account as well as any subscriptions you require.</p>	<ol style="list-style-type: none"> 1. Go to the Palo Alto Networks Support Site, log in, and select My Devices. 2. Verify that the firewall is listed. If it is not listed, see Register the Firewall. 3. (Optional) Activate Licenses and Subscriptions.

Enable WildFire (Continued)	
<p>Step 2 Set the WildFire forwarding options.</p> <p> If you do not have a WildFire subscription you can only forward executables.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > WildFire and edit the General Settings. 2. (Optional) Specify the WildFire Server to which to forward files. By default, the firewall will forward files to the public WildFire cloud hosted in the United States. To forward files to a different WildFire cloud, enter a new value as follows: <ul style="list-style-type: none"> • To forward to a private WildFire cloud, enter the IP address or FQDN of your WF-500 WildFire appliance. • To forward files to the public WildFire cloud running in Japan, enter <code>wildfire.paloaltonetworks.jp</code>. 3. Review the maximum file size that the firewall can forward for a specific type of file. <p> It is a WildFire best practice to set the File Size limit for PEs to the maximum file size supported: 10 MB.</p> 4. Click OK to save your changes.
<p>Step 3 Enable the firewall to forward decrypted SSL traffic for WildFire analysis.</p> <p> This is a WildFire best practice.</p>	<p>On a firewall with no virtual systems configured:</p> <ol style="list-style-type: none"> 1. Select Device > Setup > Content-ID. 2. Edit the Content-ID settings and Allow Forwarding of Decrypted Content. 3. Click OK to save the changes. <p>On a firewall with multiple virtual systems configured:</p> <p>Select Device > Virtual Systems, select the virtual system you want to modify, and Allow Forwarding of Decrypted Content.</p>
<p>Step 4 Set up a file blocking profile to forward files to WildFire.</p>	<ol style="list-style-type: none"> 1. Select Objects > Security Profiles > File Blocking and click Add. 2. Enter a Name and optionally a Description for the profile. 3. Click Add to create a forwarding rule and enter a name. 4. In the Action column, select forward. 5. Leave the other fields set to any to forward any supported file type from any application. 6. Click OK to save the profile.
<p>Step 5 Attach the file blocking profile to the security policies that allow access to the Internet.</p>	<ol style="list-style-type: none"> 1. Select Policies > Security and either select an existing policy or create a new policy as described in Create Security Rules. 2. Click the Actions tab within the security policy. 3. In the Profile Settings section, click the drop-down and select the file blocking profile you created for WildFire forwarding. (If you don't see a drop-down for selecting a profile, select Profiles from the Profile Type drop-down.
<p>Step 6 Save the configuration.</p>	<p>Click Commit.</p>

Enable WildFire (Continued)	
<p>Step 7 Verify that the firewall is forwarding files to WildFire.</p>	<ol style="list-style-type: none"> 1. Select Monitor > Logs > Data Filtering. 2. Check the Action column for the following actions: <ul style="list-style-type: none"> • Forward—Indicates that the file was successfully forwarded by the file blocking profile attached to the security policy. • Wildfire-upload-success—Indicates that the file was sent to WildFire. This means the file is not signed by a trusted file signer and it has not been previously analyzed by WildFire. • Wildfire-upload-skip—Indicates that the file was identified as eligible to be sent to WildFire by a file blocking profile/security policy, but did not need to be analyzed by WildFire because it has already been analyzed previously. In this case, the action will display as forward in the Data Filtering log because it was a valid forward action, but it was not sent to WildFire and analyzed because the file has already been sent to the WildFire cloud from another session, possibly from another firewall. 3. View the WildFire logs by selecting Monitor > Logs > WildFire Submissions. If new WildFire logs appear, the firewall is successfully forwarding files to WildFire and WildFire is returning file analysis reports.

Scan Traffic for Threats

[Security Profiles](#) provide threat protection in security policies. For example, you can apply an antivirus profile to a security policy and all traffic that matches the security policy will be scanned for viruses.

The following sections provide steps for setting up a basic threat prevention configuration:


- ▲ [Set Up Antivirus, Anti-Spyware, and Vulnerability Protection](#)
- ▲ [Set Up File Blocking](#)

Set Up Antivirus, Anti-Spyware, and Vulnerability Protection

Every Palo Alto Networks next-generation firewall comes with predefined Antivirus, Anti-Spyware, and Vulnerability Protection profiles that you can attach to security policies. There is one predefined Antivirus profile, **default**, which uses the default action for each protocol (block HTTP, FTP, and SMB traffic and alert on SMTP, IMAP, and POP3 traffic). There are two predefined Anti-spyware and Zone Protection profiles:

- **default**—Applies the default action to all client and server critical, high, and medium severity spyware/vulnerability protection events. It does not detect low and informational events.
- **strict**—Applies the block response to all client and server critical, high and medium severity spyware/vulnerability protection events and uses the default action for low and informational events.

To ensure that the traffic entering your network is free from threats, attach the predefined profiles to your basic web access policies. As you monitor the traffic on your network and expand your policy rulebase, you can then design more granular profiles to address your specific security needs.

Set up Antivirus/Anti-Spyware/Vulnerability Protection	
<p>Step 1 Verify that you have a Threat Prevention license.</p>	<ul style="list-style-type: none"> The Threat Prevention license bundles the Antivirus, Anti-Spyware, and the Vulnerability Protection features in one license. Select Device > Licenses to verify that the Threat Prevention license is installed and valid (check the expiration date).
<p>Step 2 Download the latest antivirus threat signatures.</p>	<ol style="list-style-type: none"> Select Device > Dynamic Updates and click Check Now at the bottom of the page to retrieve the latest signatures. In the Actions column, click Download to install the latest Antivirus, and Applications and Threats signatures.
<p>Step 3 Schedule signature updates.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>Perform a download-and-install on a daily basis for antivirus updates and weekly for applications and threats updates.</p> </div>	<ol style="list-style-type: none"> From Device > Dynamic Updates, click the text to the right of Schedule to automatically retrieve signature updates for Antivirus and Applications and Threats. Specify the frequency and timing for the updates and whether the update will be downloaded and installed or only downloaded. If you select Download Only, you would need to manually go in and click the Install link in the Action column to install the signature. When you click OK, the update is scheduled. No commit is required. (Optional) You can also enter the number of hours in the Threshold field to indicate the minimum age of a signature before a download will occur. For example, if you entered 10, the signature must be at least 10 hours old before it will be downloaded, regardless of the schedule. In an HA configuration, you can also click the Sync To Peer option to synchronize the content update with the HA peer after download/install. This will not push the schedule settings to the peer device, you need to configure the schedule on each device.

Recommendations for HA Configurations:

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

Set up Antivirus/Anti-Spyware/Vulnerability Protection (Continued)

Step 4 Attach the security profiles to a security policy.



Attach a clone of a predefined security profile to your basic security policies. That way, if you want to customize the profile you can do so without deleting the read-only predefined **strict** or **default** profile and attaching a customized profile.

1. Select **Policies > Security**, select the desired policy to modify it and then click the **Actions** tab.
2. In **Profile Settings**, click the drop-down next to each security profile you would like to enable. In this example we choose default for **Antivirus**, **Vulnerability Protection**, and **Anti-Spyware**.





If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.

Step 5 Save the configuration.

Click **Commit**.

Set Up File Blocking

File Blocking Profiles allow you to identify specific file types that you want to want to block or monitor. The following workflow shows how to set up a basic file blocking profile that prevents users from downloading executable files from the Internet.

Configure File Blocking											
<p>Step 1 Create the file blocking profile.</p>	<ol style="list-style-type: none"> 1. Select Objects > Security Profiles > File Blocking and click Add. 2. Enter a Name for the file blocking profile, for example <code>Block_EXE</code>. 3. Optionally enter a Description, such as <code>Block users from downloading exe files from websites</code>. 										
<p>Step 2 Configure the file blocking options.</p> <div data-bbox="245 1066 639 1188" style="border: 1px solid #ccc; padding: 5px; margin-top: 20px;">  The forward and continue-and-forward actions are for forwarding files to WildFire only. </div>	<ol style="list-style-type: none"> 1. Click Add to define the profile settings. 2. Enter a Name, such as <code>BlockEXE</code>. 3. Set the Applications to which to apply file blocking, or leave it set to any. 4. Set File Types to block. For example, to block download of executables, you would select exe. 5. Specify the Direction in which to block files: download, upload, or both. 6. Set the Action to one of the following: <ul style="list-style-type: none"> • continue—(web traffic only) Files matching the selected criteria will trigger a customizable response page that requires users to click Continue in order to proceed with the download/upload. You must enable response pages on the associated interfaces if you plan to use this option (Step 4). • block—Files matching the selected criteria will be blocked from download/upload. • alert—Files matching the selected criteria will be allowed, but will generate a log entry in the data filtering log. <div data-bbox="695 1339 1377 1528" style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>File Blocking Profile</p> <p>Name: <code>Block_EXE</code></p> <p>Description: <code>Block .exe files from being downloaded from websites.</code></p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th>Names</th> <th>Applications</th> <th>File Types</th> <th>Direction</th> <th>Action</th> </tr> </thead> <tbody> <tr> <td><input type="checkbox"/> BlockEXE</td> <td>web-browsing</td> <td>exe</td> <td>download</td> <td>continue</td> </tr> </tbody> </table> </div> 7. Click OK to save the profile. 	Names	Applications	File Types	Direction	Action	<input type="checkbox"/> BlockEXE	web-browsing	exe	download	continue
Names	Applications	File Types	Direction	Action							
<input type="checkbox"/> BlockEXE	web-browsing	exe	download	continue							
<p>Step 3 Attach the file blocking profile to the security policies that allow access to content.</p>	<ol style="list-style-type: none"> 1. Select Policies > Security and either select an existing policy or create a new policy as described in Create Security Rules. 2. Click the Actions tab within the security policy. 3. In the Profile Settings section, click the drop-down and select the file blocking profile you created. <div data-bbox="735 1787 1377 1843" style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  If you don't see drop-downs for selecting profiles, select Profiles from the Profile Type drop-down. </div> 										

Configure File Blocking (Continued)	
<p>Step 4 Enable response pages in the management profile for each interface on which you are attaching file blocking profile with a continue action.</p>	<ol style="list-style-type: none"> 1. Select Network > Network Profiles > Interface Mgmt and then select an interface profile to edit or click Add to create a new profile. 2. Select Response Pages, as well as any other management services required on the interface. 3. Click OK to save the interface management profile. 4. Select Network > Interfaces and select the interface to which to attach the profile. 5. On the Advanced > Other Info tab, select the interface management profile you just created. 6. Click OK to save the interface settings.
<p>Step 5 Test the file blocking configuration.</p>	<p>Access a client PC in the trust zone of the firewall and attempt to download an.exe file from a website in the untrust zone. Make sure the file is blocked as expected based on the action you defined in the file blocking profile:</p> <ul style="list-style-type: none"> • If you selected alert as the action, check the data filtering log to make sure you see a log entry for the request. • If you selected block as the action, the File Blocking Block Page response page should display. • If you selected the continue action, the File Blocking Continue Page response page should display. Click Continue to download the file. The following shows the default File Blocking Continue Page.

File Download Blocked

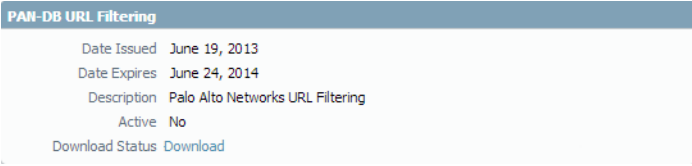
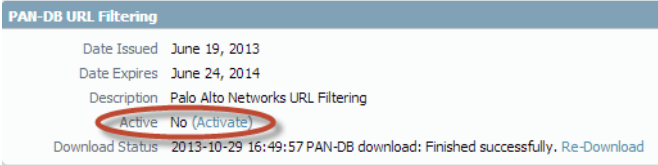

Access to the file you were trying to download has been blocked in accordance with company policy. Please contact your system administrator if you believe this is in error.

File name: Support_services_ds1.pdf

Please click to download/upload the file.

Control Access to Web Content

URL Filtering provides visibility and control over web traffic on your network. With URL filtering enabled, the firewall can categorize web traffic into one or more (from approximately 60) categories. You can then create policies that specify whether to allow, block, or log (alert) traffic based on the category to which it belongs. The following workflow shows how to enable PAN-DB for URL filtering, create security profiles, and attach them to security policies to enforce a basic URL filtering policy.

Configure URL Filtering	
<p>Step 1 Confirm license information for URL Filtering.</p>	<ol style="list-style-type: none"> Obtain and install a URL Filtering license. See Activate Licenses and Subscriptions for details. Select Device > Licenses and verify that the URL Filtering license is valid. 
<p>Step 2 Download the seed database and activate the license.</p>	<ol style="list-style-type: none"> To download the seed database, click Download next to Download Status in the PAN-DB URL Filtering section of the Licenses page. Choose a region (North America, Europe, APAC, Japan) and then click OK to start the download. After the download completes, click Activate. 
<p>Step 3 Create a URL filtering profile.</p> <p> Because the default URL filtering profile blocks risky and threat-prone content, clone this profile when creating a new profile in order to preserve the default settings.</p>	<ol style="list-style-type: none"> Select Objects > Security Profiles > URL Filtering. Select the default profile and then click Clone. The new profile will be named default-1. Select the new profile and rename it.

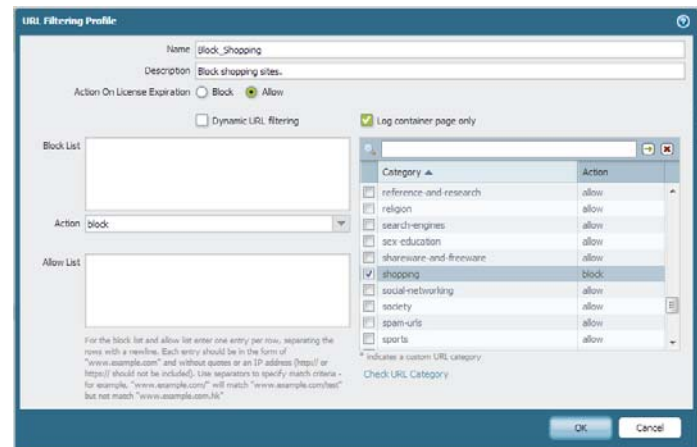
Configure URL Filtering (Continued)

Step 4 Define how to control access to web content.

If you are not sure what traffic you want to control, consider setting the categories (except for those blocked by default) to alert. You can then use the visibility tools on the firewall, such as the ACC and App Scope, to determine which web categories to restrict to specific groups or to block entirely. You can then go back and modify the profile to block and allow categories as desired.

You can also define specific sites to always allow or always block regardless of category and enable the safe search option to filter search results when defining the [URL Filtering](#) profile.

- For each category that you want visibility into or control over, select a value from the **Action** column as follows:
 - If you do not care about traffic to a particular category (that is you neither want to block it nor log it), select **allow**.
 - For visibility into traffic to sites in a category, select **alert**.
 - To present a response page to users attempting to access a particular category to alert them to the fact that the content they are accessing might not be work appropriate, select **continue**.
 - To prevent access to traffic that matches the associated policy, select **block** (this also generates a log entry).



- Click **OK** to save the URL filtering profile.

Step 5 Attach the URL filtering profile to a security policy.

- Select **Policies > Security**.
- Select the desired policy to modify it and then click the **Actions** tab.
- If this is the first time you are defining a security profile, select **Profiles** from the **Profile Type** drop-down.
- In the **Profile Settings** list, select the profile you just created from the **URL Filtering** drop-down. (If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.)
- Click **OK** to save the profile.
- Commit** the configuration.

Configure URL Filtering (Continued)	
<p>Step 6 Enable Response Pages in the management profile for each interface on which you are filtering web traffic.</p>	<ol style="list-style-type: none"> 1. Select Network > Network Profiles > Interface Mgmt and then select an interface profile to edit or click Add to create a new profile. 2. Select Response Pages, as well as any other management services required on the interface. 3. Click OK to save the interface management profile. 4. Select Network > Interfaces and select the interface to which to attach the profile. 5. On the Advanced > Other Info tab, select the interface management profile you just created. 6. Click OK to save the interface settings.
<p>Step 7 Save the configuration.</p>	<p>Click Commit.</p>
<p>Step 8 Test the URL filtering configuration.</p>	<p>Access a client PC in the trust zone of the firewall and attempt to access a site in a blocked category. Make sure URL filtering is applied based on the action you defined in the URL filtering profile:</p> <ul style="list-style-type: none"> • If you selected alert as the action, check the data filtering log to make sure you see a log entry for the request. • If you selected the continue action, the URL Filtering Continue and Override Page response page should display. Continue to the site. • If you selected block as the action, the URL Filtering and Category Match Block Page response page should display as follows: <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>Web Page Blocked</p> <p>Access to the web page you were trying to visit has been blocked in accordance with company policy. Please contact your system administrator if you believe this is an error.</p> <p>User: 192.168.2.10</p> <p>URL: amazon.com/</p> <p>Category: shopping</p> </div>

For More Information

For more detailed information on how to protect your enterprise from threats, see [Threat Prevention](#). For details on how to scan encrypted (SSH or SSL) traffic for threats, see [Decryption](#).

For information about the threats and applications that Palo Alto Networks products can identify, visit the following links:

- [Applipedia](#)—Provides details on the applications that Palo Alto Networks can identify.
- [Threat Vault](#)—Lists threats that Palo Alto Networks products can identify. You can search by Vulnerability, Spyware, or Virus. Click the Details icon next to the ID number for more information about a threat.

Best Practices for Completing the Firewall Deployment

Now that you have integrated the firewall into your network and enabled the basic security features, you can begin configuring more advanced features. Here are some things to consider next:

- ❑ Learn about the different [Management Interfaces](#) that are available to you and how to access and use them.
- ❑ Set up [High Availability](#)—High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration is synchronized to prevent a single point to failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up the firewalls in a two-device cluster provides redundancy and allows you to ensure business continuity.
- ❑ [Configure the Master Key](#)—Every Palo Alto Networks firewall has a default master key that encrypts private keys that are used to authenticate administrators when they access management interfaces on the firewall. As a best practice to safeguard the keys, configure the master key on each firewall to be unique. However, if you use Panorama, you must use the same master key on Panorama and all managed firewalls. Otherwise, Panorama cannot push configurations to the firewalls.
- ❑ [Manage Firewall Administrators](#)—Every Palo Alto Networks firewall and appliance is preconfigured with a default administrative account (admin) that provides full read-write access (also known as superuser access) to the device. As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This allows you to better protect the device from unauthorized configuration (or modification) and to enable logging of the actions of each individual administrator.
- ❑ Enable User Identification ([User-ID](#))—User-ID is a Palo Alto Networks next-generation firewall feature that allows you to create policies and perform reporting based on users and groups rather than individual IP addresses.
- ❑ Enable [Decryption](#)—Palo Alto Networks firewalls provide the capability to decrypt and inspect traffic for visibility, control, and granular security. Use decryption on a firewall to prevent malicious content from entering your network or sensitive content from leaving your network concealed as encrypted or tunneled traffic.
- ❑ [Enable Passive DNS Collection for Improved Threat Intelligence](#)—Enable this opt-in feature to enable the firewall to act as a passive DNS sensor and send select DNS information to Palo Alto Networks for analysis in order to improve threat intelligence and threat prevention capabilities.
- ❑ Follow the [Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions](#).



Device Management

Administrators can configure, manage, and monitor the Palo Alto Networks firewalls using the web interface, the CLI, and the API management interfaces. Role-based administrative access to the management interfaces can be customized in order to delegate specific tasks or permissions to certain administrators. See the following topics for information on device management options, including how to begin using the management interfaces and how to customize administrator roles:

- ▲ [Management Interfaces](#)
- ▲ [Manage Firewall Administrators](#)
- ▲ [Reference: Web Interface Administrator Access](#)
- ▲ [Reference: Port Numbers Used by Palo Alto Networks Devices](#)
- ▲ [Reset the Firewall to Factory Default Settings](#)

Management Interfaces

PAN-OS firewalls and Panorama provide three user interfaces: a web interface, a command line interface (CLI), and a XML-based management API. See the following topics for how to access and begin using each of the device management interfaces:

- [Use the Web Interface](#) to complete administrative tasks and generate reports from the web interface with relative ease. This graphical interface allows you to access the firewall using HTTPS and it is the best way to perform administrative tasks.
- [Use the Command Line Interface \(CLI\)](#) to type through the commands in rapid succession to complete a series of tasks. The CLI is a no-frills interface that supports two command modes and each mode has its own hierarchy of commands and statements. When you get familiar with the nesting structure and the syntax of the commands, the CLI allows quick response times and offers administrative efficiency.
- [Use the XML API](#) to streamline your operations and integrate with existing, internally developed applications and repositories. The XML API is provided as a web service that is implemented using HTTP/HTTPS requests and responses.

Use the Web Interface

The following topics describes how to begin using the firewall web interface. For detailed information about the tabs and fields that are available in the web interface, refer to the [Web Interface Reference Guide](#).


- ▲ [Launch the Web Interface](#)
- ▲ [Navigate the Web Interface](#)
- ▲ [Commit Changes](#)
- ▲ [Use Configuration Pages](#)
- ▲ [Required Fields](#)
- ▲ [Lock Transactions](#)

Launch the Web Interface

The following web browsers are supported for access to the web interface for PAN-OS firewalls and Panorama:

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

Launch an Internet browser and enter the firewall's IP address. Enter your user credentials. If logging in to the firewall for the first time, type the default **admin** into both the **Name** and **Password** fields.

To view information on how to use a specific page and an explanation of the fields and options on the page, click the **Help** icon  in the upper right area of the page to open the online help system. In addition to displaying context-sensitive help for a page, clicking the **Help** icon displays a help navigation pane with options to browse and search all help content.


Navigate the Web Interface

The following conventions apply when using the web interface.


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



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




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




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



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



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



- To modify an item, click its underlined link.

	Name	Location	Protocol
<input type="checkbox"/>	service-http	Predefined	TCP
<input type="checkbox"/>	service-https	Predefined	TCP

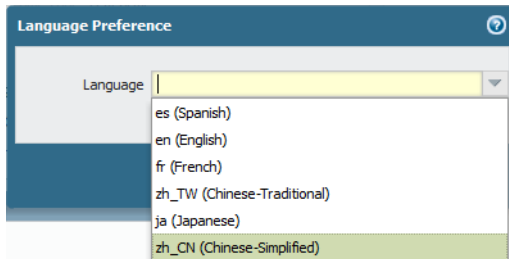
- To view the current list of tasks, click the **Tasks** icon in the lower right corner of the page. The Task Manager window opens to show the list of tasks, along with status, start times, associated messages, and actions. Use the **Show** drop-down to filter the list of tasks.

A screenshot of the Task Manager window titled 'Task Manager - All Tasks'. It contains a table with columns: Type, Status, Start Time, Messages, and Action.


Type	Status	Start Time	Messages	Action
Commit	Completed	09/09/11 07:57:23	<ul style="list-style-type: none"> In virtual-router vr1: address 10.40.1.1/24 on interface ethernet1/7 is duplicate with address 10.40.1.2/24 on interface ethernet1/8.(Module: routed) Commit failed 	
Commit	Completed	09/09/11 07:56:16	<ul style="list-style-type: none"> In virtual-router vr1: address 10.40.1.1/24 on interface ethernet1/7 is duplicate with address 10.40.1.2/24 on interface ethernet1/8.(Module: routed) Commit failed 	
Commit	Completed	09/09/11 07:52:26		
Commit	Completed	09/09/11 07:49:25	<ul style="list-style-type: none"> Interface ethernet1/7 has no ip pool(Module: dhcpd) Configuration committed successfully 	
Auto Commit	Completed	09/08/11 15:06:15	<ul style="list-style-type: none"> Configuration committed successfully Successfully committed last configuration L3 Service Configuration is changed. l3svc will be restarted. (Module: device) 	

At the bottom of the window, there is a 'Show' dropdown menu set to 'All Tasks' and a 'Close' button.

- The web interface language is controlled by the current language of the computer that is managing the device if a specific language preference has not been defined. For example, if the computer you use to manage the firewall has a locale of Spanish, when you log in to the firewall, the web interface will be in Spanish.
- To specify a language that will always be used for a given account regardless of the locale of the computer, click the **Language** icon in the lower right corner of the page and the Language Preference window opens. Click the drop-down to select the desired language and then click **OK** to save your change.



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



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

Commit Changes

Click **Commit** at the top of the web interface to open the Commit dialog box, which displays the following options. Certain options only display if you click **Advanced**.



If dependencies between the configuration changes you included and excluded cause a validation error, perform the commit with all the changes included. For example, if your changes introduce a new Log Forwarding profile (an object) that references a new Syslog server profile (a device setting), the commit must include both policy and object configurations and device and network configurations.

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

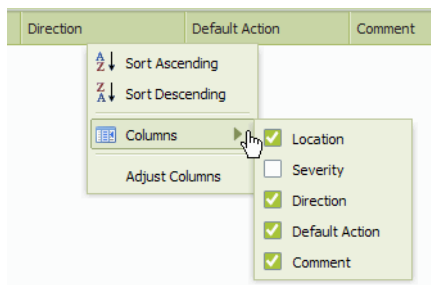
- **Preview Changes**—Click this button to bring up a two-pane window that shows proposed changes in the candidate configuration compared to the current running configuration. You can choose the number of lines of context to display, or show all lines. Changes are color coded based on items that you and other administrators added (green), modified (yellow), or deleted (red) since the last commit.



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


Use Configuration Pages

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



Required Fields

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




Lock Transactions

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



- **Config lock**—Blocks other administrators from making changes to the configuration. This type of lock can be set globally or for a virtual system. It can be removed only by the administrator who set it or by a superuser on the system.

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

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

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

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

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

You can arrange to automatically acquire a commit lock by selecting the **Automatically acquire commit lock** check box in the Management area of the **Device Setup** page.

Use the Command Line Interface (CLI)

The PAN-OS CLI allows you to access Firewall and Panorama devices, view status and configuration information, and modify configurations. Access to the PAN-OS CLI is provided through SSH, Telnet, or direct console access.

The following topics describe how to access and begin using the PAN-OS CLI:

- ▲ [Access the PAN-OS CLI](#)
- ▲ [Operational and Configuration Modes](#)

For more information on the CLI, refer to the *PAN-OS Command Line Interface Reference Guide*.

Access the PAN-OS CLI

You can connect to the CLI on the firewall either directly using a serial cable from your management computer to the Console port, or by opening a Secure Shell (SSH) connection to the firewall MGT port as follows:

Access the PAN-OS CLI

- Step 1** Launch a terminal emulator, such as PuTTY, and open a connection to the firewall using one of the following methods:
- **Console Connection**—Connect a serial cable from your computer to the Console port and connect use the terminal emulation software to open a Serial connection to the firewall. The firewall requires that you configure the terminal emulation software with the following settings:
 - Data rate: 9600
 - Data bits: 8
 - Parity: none
 - Stop bits: 1
 - Flow control: None
 - **SSH Connection**—Use the terminal emulation software to open an SSH connection (default [port 22](#)) to the IP address you assigned to the MGT port during [initial configuration](#).

Step 2 Enter the administrative username. The default is admin.

Step 3 Enter the administrative password. The default is admin.

Step 4 The PAN-OS CLI opens in Operational mode, and the CLI prompt is displayed:

Step 5 `username@hostname>`

Operational and Configuration Modes

When you log in, the PAN-OS CLI opens in Operational mode. You can move between Operational and Configuration modes at any time. Use Operational mode to view the state of the system, navigate the PAN-OS CLI, and enter configuration mode. Use Configuration mode to view and modify the configuration hierarchy.

- To enter Configuration mode from Operational mode, use the `configure` command:

```
username@hostname> configure
Entering configuration mode
[edit]
username@hostname#
```
- To leave Configuration mode and return to Operational mode, use the `quit` or `exit` command:

```
username@hostname# quit
Exiting configuration mode
username@hostname>
```
- To enter an Operational mode command while in Configuration mode, use the `run` command, for example:

```
username@hostname# run ping host 1.1.1.2
PING 1.1.1.1 (1.1.1.1) 56(84) bytes of data
...
username@hostname#
```
- To direct an Operational mode command to a particular VSYS, specify the target VSYS with the following command:

```
username@hostname# set system setting target-vsys <vsys_name>
```

Use the XML API

Palo Alto Networks XML API uses standard HTTP requests to send and receive data, allowing access to several types of data on the device so the data can be easily integrated with and used in other systems. Use the XML-based Management API to view a firewall or Panorama's configuration, extract report data in XML format, and execute operational commands. API calls can be made directly from command line utilities such as cURL or wget, or using any scripting or application framework that supports RESTful services. When using the API with command lines tools, both HTTP GET and POST methods are supported.

You must generate an API key in order to use the XML API. The API key authenticates the user to the firewall, application, or Panorama. After you have generated an API key, you can use the key to perform device configuration and operational tasks, retrieve reports and logs, and import and export files. See [Generate an API Key](#) for steps to generate an API key.

The following table shows the URL structure for API requests:

PAN-OS Version	XML API URL Structure
Prior to PAN-OS 4.1.0	http(s)://hostname/esp/restapi.esp?request-parameters-values
PAN-OS 4.1.0 and later	http(s)://hostname/api/?request-parameters-values

URL structure item definitions:

- **hostname**—Device's IP address or Domain name.
- **request-parameters-values**—A series of multiple 'parameter=value' pairs separated by the ampersand character (&). These values can either be keywords or data-values in standard or XML format (response data is always in XML format).

There are APIs for PAN-OS, User-ID, and WildFire products. For more information on how to use the API interface, refer to the [PAN-OS XML API Usage Guide](#). To access the online community for developing scripts, visit: <https://live.paloaltonetworks.com/community/devcenter>.

Generate an API Key

In order to use the API to manage a firewall or application, an API key is required to authenticate all API calls. Admin account credentials are used to generate API keys.



As a best practice, create a separate admin account for XML-based administration.

Generate an API key	
<p>Step 1 Create an administrator account.</p>	<ol style="list-style-type: none"> 1. In the web interface, on the Device > Administrators tab, click Add. 2. Enter a login Name for the admin. 3. Enter and confirm a Password for the admin. 4. Click OK and Commit.

Generate an API key (Continued)

<p>Step 2 Request an API key.</p>	<p>Replace the hostname, username and password parameters in the following URL with the appropriate values from your administrator account credentials:</p> <pre>http(s)://hostname/api/?type=keygen&user=username&password=password</pre> <p>The API key is displayed in an XML block. For example:</p> <pre><response status="success"> <result> <key>0RgWc420i0vDx2WRUIUM6A</key> </result> </response></pre>
<p>Step 3 (Optional) Revoke or change an API key.</p> <p>For PAN-OS 4.1.0 and later releases, generating an API key using the same administrator account credentials returns unique API keys every time, and all of the keys are valid.</p> <p>You can choose to revoke and then change an API key associated with an administrator account by changing the password associated with the administrator account. Any API keys that were generated using the previous credentials would no longer be valid.</p>	<ol style="list-style-type: none"> 1. On the Device > Administrators tab, open the administrator account associated with the API key. 2. Enter and confirm a new Password for the administrator account. 3. Click OK and Commit. <p>Any API keys associated with the admin account prior to the password change are revoked upon Commit.</p> <ol style="list-style-type: none"> 4. (Optional) Use the updated administrator account credentials to generate a new API key. See Step 2.

Example work flow using an API key:

Request an API key by entering the URL with the appropriate values in a web browser:

```
https://10.xx.10.50/esp/restapi.esp?type=keygen&user=admin&password=admin
```

Entering the URL displays an XML block that contains the API key:

```
<response status="success">
  <result>
    <key>0RgWc420i0vDx2WRUIUM6A</key>
  </result>
</response>
```

Continue to use the API key to create API requests. For example, to generate a report:

```
https://10.xx.10.50/esp/restapi.esp?type=report&reporttype=dynamic&reportname=top-app-summary&period=last-hour&topn=5&key=0RgWc420i0vDx2WRUIUM6A=
```

Manage Firewall Administrators

Every Palo Alto Networks firewall and appliance is preconfigured with a default administrative account (admin) that provides full read-write access (also known as superuser access) to the device.



As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of the firewall. This allows you to better protect the device from unauthorized configuration (or modification) and to enable logging of the actions of each individual administrator.

The following topics describe the various ways you can set up administrative accounts and provides procedures for setting up basic administrative access:

- ▲ [Administrative Roles](#)
- ▲ [Administrative Authentication](#)
- ▲ [Create an Administrative Account](#)

Administrative Roles

The way you configure administrator accounts depends on the security requirements within your organization, whether you have existing authentication services you want to integrate with, and how many different administrative roles you require. A *role* defines the type of access the associated administrator has to the system. There are two types of roles you can assign:

- **Dynamic Roles**—These are built-in roles that provide access to the firewall. When new features are added, the firewall automatically updates the definitions of dynamic roles; you never need to manually update them. The following table lists the access privileges associated with dynamic roles.

Dynamic Role	Privileges
Superuser	Full access to the firewall, including defining new administrator accounts and virtual systems. You must have superuser privileges to create an administrative user with superuser privileges.
Superuser (read-only)	Read-only access to the firewall.
Virtual system administrator	Full access to a selected virtual system (vsys) on the firewall.
Virtual system administrator (read-only)	Read-only access to a selected vsys on the firewall.
Device administrator	Full access to all firewall settings except for defining new accounts or virtual systems.
Device administrator (read-only)	Read-only access to all firewall settings except password profiles (no access) and administrator accounts (only the logged in account is visible).

- **Admin Role Profiles**—Custom roles you can configure for more granular access control over the functional areas of the web interface, CLI, and XML API. For example, you can create an Admin Role profile for your operations staff that provides access to the firewall and network configuration areas of the web interface and a separate profile for your security administrators that provides access to security policy definitions, logs, and reports. On a multi-vsyt firewall, you can select whether the role defines access for all

virtual systems or for a specific vsys. When new features are added to the product, you must update the roles with corresponding access privileges: the firewall does not automatically add new features to custom role definitions. For details on the privileges you can configure for custom administrator roles, see [Reference: Web Interface Administrator Access](#).

Administrative Authentication

There are four ways you can authenticate administrative users:

- **Local administrator account with local authentication**—Both the administrator account credentials and the authentication mechanisms are local to the firewall. You can further secure the local administrator account by creating a password profile that defines a validity period for passwords and by setting device-wide password complexity settings.
- **Local administrator account with SSL-based authentication**—With this option, you create the administrator accounts on the firewall, but authentication is based on SSH certificates (for CLI access) or client certificates/common access cards (for the web interface). Refer to the article [How to Configure Certificate-based Authentication for the WebUI](#) for details on how to configure this type of administrative access.
- **Local administrator account with external authentication**—The administrator accounts are managed on the local firewall, but the authentication functions are offloaded to an existing LDAP, Kerberos, or RADIUS service. To configure this type of account, you must first create an authentication profile that defines how to access the external authentication service and then create an account for each administrator that references the profile.
- **External administrator account and authentication**—Account administration and authentication are handled by an external RADIUS server. To use this option, you must define Vendor Specific Attributes (VSAs) on your RADIUS server that map to the admin role and, optionally, the virtual system objects you have defined on the Palo Alto Networks device. Refer to the [Radius Vendor Specific Attributes \(VSA\)](#) article for details on how to configure this type of administrative access.

Create an Administrative Account

Create administrative accounts to define access and administrative privileges for firewall administrators. Because it is common to delegate specific administrative tasks to specific administrators with varying roles, Palo Alto Networks recommends that you create admin role profiles that allow administrators access only to the areas of the management interface that are required to perform their jobs. You can assign the various roles you create to individual administrator accounts and specify access privileges to each management interface: the web interface, the Command Line Interface (CLI), and the XML-based Management API. By creating admin roles with very granular access privileges, you can ensure that sensitive company data is protected and end user privacy is ensured.

The following procedure describes how to create a local administrator account with local authentication, including how to set administrator access for each management interface.

Create a Local Administrator

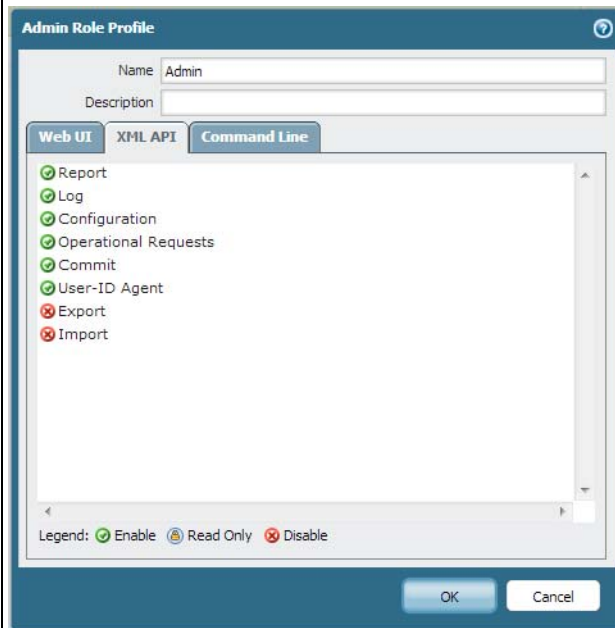
Step 1 Create the Admin Role Profiles that you plan to assign to your administrators (this does not apply if you plan to use Dynamic Roles). The Admin Role Profiles define what type of access to give to the different sections of the web interface, CLI, and XML API for each administrator that you assign a role to.

You can use this step to set particularly granular access privileges for web interface users. For details on what a specific option enables on the **Web UI** tab, see [Web Interface Access Privileges](#).

Complete the following steps for each role you want to create:

1. Select **Device > Admin Roles** and then click **Add**.
2. Enter a **Name** and optionally a **Description** for the role.
3. On the Web UI, Command Line and/or XML API tabs, specify the access to allow for each management interface:
 - On the **Web UI** and/or **XML API** tabs, set the access levels for each functional area of the interface by clicking the icon to toggle it to the desired setting: Enable, Read Only, or Disable.
 - On the **Command Line** tab, specify the type of access to allow to the CLI: **superreader**, **deviceadmin**, or **devicereader** (for Device roles); **vsysadmin** or **vsysreader** (for Virtual System roles); or **None** to disable CLI access entirely.
4. Click **OK** to save the profile.

For example, allow an admin full access to a device using the XML API, with the exception of importing or exporting files:



Create a Local Administrator (Continued)	
<p>Step 2 (Optional) Set requirements for local user-defined passwords.</p>	<ul style="list-style-type: none"> • Create Password Profiles—Define how often administrators must change their passwords. You can create multiple password profiles and apply them to administrator accounts as needed to enforce the desired security. To create a password profile, select Device > Password Profiles and then click Add. • Configure minimum password complexity settings—Define rules that govern password complexity, allowing you to force administrators to create passwords that are harder to guess, crack, or compromise. Unlike password profiles, which can be applied to individual accounts, these rules are device-wide and apply to all passwords. To configure the settings, select Device > Setup and then Edit in the Minimum Password Complexity section.
<p>Step 3 Create an account for each administrator.</p>	<ol style="list-style-type: none"> 1. Select Device > Administrators and then click Add. 2. Enter a user Name and Password for the administrator, or create an Authentication Profile to use for validating an administrative user's credentials to an external authentication server. See Step 4 for details on setting up an authentication profile. 3. Select the Role to assign to this administrator. You can either select one of the predefined Dynamic roles or a custom Role Based profile if you created one in Step 1. 4. (Optional) Select a Password Profile. 5. Click OK to save the account.

Create a Local Administrator (Continued)	
<p>Step 4 (Optional) Set up authentication to an external server—LDAP, RADIUS or Kerberos.</p> <p>The server profile specifies how the firewall can connect to the authentication service you plan to use.</p>	<ol style="list-style-type: none"> 1. Select Device > Authentication Profile and then click Add. 2. Enter a user Name to identify the authentication profile. 3. Define the conditions for locking out the administrative user. <ol style="list-style-type: none"> a. Enter the Lockout Time. This is the number of minutes that a user is locked out upon reaching the maximum number of failed attempts (0-60 minutes; default 0). 0 means that the lockout is in effect until it is manually unlocked. b. Enter the Failed Attempts count. This is the number of failed login attempts that are allowed before the account is locked out (1-10; default 0). By default, the failed attempt count is 0 and the user is not locked out despite repeated failure to authenticate. 4. Specify the users and groups that are explicitly allowed to authenticate. By adding an Allow List to an authentication profile, you can limit access to specific users in a user group/directory. <ul style="list-style-type: none"> • Select the All check box to allow all users. • Click Add and enter the first few characters of a name in the field to list all the users and user groups that start with those characters. Repeat to add as many users/user groups as required. 5. In the Authentication drop-down, select the type of authentication you plan to use on your network. If you plan to use local database authentication, you must create the local database. Select Device > Local User Database and add the users and groups to be authenticated. 6. For access to an external authentication server (not local database), select the appropriate server profile in the Server Profile drop-down. To create a new server profile, click the link next to New and continue with configuring access to the LDAP, RADIUS or Kerberos server. 7. Click OK.
<p>Step 5 Commit your changes.</p>	<ol style="list-style-type: none"> 1. Click Commit.

Reference: Web Interface Administrator Access

While the privileges associated with dynamic administrator roles are fixed, you can configure privileges at a granular level for custom administrator roles. Configuring privileges at a granular level ensures that lower level administrators cannot access certain information. You can create custom roles for firewall administrators (see [Create an Administrative Account](#)), Panorama administrators, or Device Group and Template administrators (refer to the [Panorama Administrator's Guide](#)). The following topics describe the privileges you can configure for custom administrator roles.


- ▲ [Web Interface Access Privileges](#)
- ▲ [Panorama Web Interface Access](#)

Web Interface Access Privileges

If you want to prevent a role-based administrator from accessing specific tabs on the web interface, you can disable the tab and the administrator will not even see it when logging in using the associated role-based administrative account. For example, you could create an Admin Role Profile for your operations staff that provides access to the **Device** and **Network** tabs only and a separate profile for your security administrators that provides access to the **Object**, **Policy**, and **Monitor** tabs.

The following table describes the tab-level access privileges you can assign to the admin role profile. It also provides cross-references to additional tables that detail granular privileges within a tab. For specific information on how to set the admin role profile to protect end user privacy, see [Define User Privacy Settings in the Admin Role Profile](#).

Access Level	Description	Enable	Read Only	Disable
Dashboard	Controls access to the Dashboard tab. If you disable this privilege, the administrator will not see the tab and will not have access to any of the Dashboard widgets.	Yes	No	Yes
ACC	Controls access to the Application Command Center (ACC). If you disable this privilege, the ACC tab will not display in the web interface. Keep in mind that if you want to protect the privacy of your users while still providing access to the ACC, you can disable the Privacy > Show Full Ip Addresses option and/or the Show User Names In Logs And Reports option.	Yes	No	Yes
Monitor	Controls access to the Monitor tab. If you disable this privilege, the administrator will not see the Monitor tab and will not have access to any of the logs, packet captures, session information, reports or to App Scope. For more granular control over what monitoring information the admin can see, leave the Monitor option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Monitor Tab .	Yes	No	Yes

Access Level	Description	Enable	Read Only	Disable
Policies	Controls access to the Policies tab. If you disable this privilege, the administrator will not see the Policies tab and will not have access to any policy information. For more granular control over what policy information the admin can see, for example to enable access to a specific type of policy or to enable read-only access to policy information, leave the Policies option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Policy Tab .	Yes	No	Yes
Objects	Controls access to the Objects tab. If you disable this privilege, the administrator will not see the Objects tab and will not have access to any objects, security profiles, log forwarding profiles, decryption profiles, or schedules. For more granular control over what objects the admin can see, leave the Objects option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Objects Tab .	Yes	No	Yes
Network	Controls access to the Network tab. If you disable this privilege, the administrator will not see the Network tab and will not have access to any interface, zone, VLAN, virtual wire, virtual router, IPsec tunnel, DHCP, DNS Proxy, GlobalProtect, or QoS configuration information or to the network profiles. For more granular control over what objects the admin can see, leave the Network option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Network Tab .	Yes	No	Yes
Device	Controls access to the Device tab. If you disable this privilege, the administrator will not see the Device tab and will not have access to any device-wide configuration information, such as User-ID, High Availability, server profile or certificate configuration information. For more granular control over what objects the admin can see, leave the Objects option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Device Tab .  You cannot enable access to the Admin Roles or Administrators nodes for a role-based administrator even if you enable full access to the Device tab.	Yes	No	Yes

Provide Granular Access to the Monitor Tab


In some cases you might want to enable the administrator to view some but not all areas of the Monitor tab. For example, you may want to restrict operations admins to the configuration and system logs only, because they do not contain sensitive user data. Although this section of the admin role definition specifies what areas of the Monitor tab the admin can see, you can also couple privileges in this section with privacy privileges, such as disabling the ability to see user names in logs and reports. One thing to keep in mind, however, is that any system-generated reports will still show user names and IP addresses even if you disable that functionality in the role. For this reason, if you do not want the admin to see any of the private user information, you should disable access to the specific reports as detailed in the following table.

The following table lists the **Monitor** tab access levels and the administrator roles for which they are available.

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Monitor	Enables or disables access to the Monitor tab. If disabled, the admin will not see this tab or any of the associated logs or reports.	Firewall: Yes 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 admin to see. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the logs, you can disable the Privacy > Show Full Ip Addresses option and/or the Show User Names In Logs And Reports option.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Traffic	Specifies whether the admin can see the traffic logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Threat	Specifies whether the admin can see the threat logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
URL Filtering	Specifies whether the admin can see the URL filtering logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
WildFire Submissions	Specifies whether the admin can see the WildFire logs. These logs are only available if you have a WildFire subscription.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Data Filtering	Specifies whether the admin can see the data filtering logs.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
HIP Match	Specifies whether the admin can see the HIP Match logs. HIP Match logs are only available if you have a GlobalProtect portal license and gateway subscription.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Configuration	Specifies whether the admin can see the configuration logs.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
System	Specifies whether the admin can see the system logs.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
Alarms	Specifies whether the admin can see system-generated alarms.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Packet Capture	Specifies whether the admin can see packet captures (pcaps) from the Monitor tab. Keep in mind that packet captures are raw flow data and as such may contain user IP addresses. Disabling the Show Full IP Addresses privileges will not obfuscate the IP address in the pcap and you should therefore disable the Packet Capture privilege if you are concerned about user privacy.	Firewall: Yes Panorama: No Device Group/Template: No	Yes	Yes	Yes
App Scope	Specifies whether the admin can see the App Scope visibility and analysis tools. Enabling App Scope enables access to all of the App Scope charts.	Firewall: Yes Panorama: Yes Device Group/Template: Yes	Yes	No	Yes
Session Browser	Specifies whether the admin can browse and filter current running sessions on the firewall. Keep in mind that the session browser shows raw flow data and as such may contain user IP addresses. Disabling the Show Full IP Addresses privileges will not obfuscate the IP address in the session browser and you should therefore disable the Session Browser privilege if you are concerned about user privacy.	Firewall: Yes Panorama: No Device Group/Template: No	Yes	No	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Botnet	Specifies whether the admin can generate and view botnet analysis reports or view botnet reports in read-only mode. Disabling the Show Full IP Addresses privileges will not obfuscate the IP address in scheduled botnet reports and you should therefore disable the Botnet privilege if you are concerned about user privacy.	Firewall: Yes 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 admin to see. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the reports, you can disable the Privacy > Show Full Ip Addresses option and/or the Show User Names In Logs And Reports option.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
Manage PDF Summary	Specifies whether the admin can view, add or delete PDF summary report definitions. With read-only access, the admin can see PDF summary report definitions, but not add or delete them. If you disable this option, the admin can neither view the report definitions nor add/delete them.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
PDF Summary Reports	Specifies whether the admin can see the generated PDF Summary reports in Monitor > Reports . If you disable this option, the PDF Summary Reports category will not display in the Reports node.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
User Activity Report	Specifies whether the admin can view, add or delete User Activity report definitions and download the reports. With read-only access, the admin can see User Activity report definitions, but not add, delete, or download them. If you disable this option, the admin cannot see this category of PDF report.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Report Groups	Specifies whether the admin can view, add or delete report group definitions. With read-only access, the admin can see report group definitions, but not add or delete them. If you disable this option, the admin cannot see this category of PDF report.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Email Scheduler	Specifies whether the admin can schedule report groups for email. Because the generated reports that get emailed may contain sensitive user data that is not removed by disabling the Privacy > Show Full Ip Addresses option and/or the Show User Names In Logs And Reports options and because they may also show log data to which the admin does not have access, you should disable the Email Scheduler option if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Manage Custom Reports	Enables or disables access to all Custom report functionality. You can also leave this privilege enabled and then disable specific custom report categories that you do not want the admin to be able to access. Keep in mind that if you want to protect the privacy of your users while still providing access to one or more of the reports, you can disable the Privacy > Show Full Ip Addresses option and/or the Show User Names In Logs And Reports option.  Reports that are scheduled to run rather than run on demand will show IP address and user information. In this case, be sure to restrict access to the corresponding report areas. In addition, the custom report feature does not restrict the ability to generate reports that contain log data contained in logs that are excluded from the admin role.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
Application Statistics	Specifies whether the admin can create a custom report that includes data from the application statistics database.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
Data Filtering Log	Specifies whether the admin can create a custom report that includes data from the data filtering log.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
Threat Log	Specifies whether the admin can create a custom report that includes data from the threat log.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Threat Summary	Specifies whether the admin can create a custom report that includes data from the threat summary database.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
Traffic Log	Specifies whether the admin can create a custom report that includes data from the traffic log.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
Traffic Summary	Specifies whether the admin can create a custom report that includes data from the traffic summary database.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
Url Log	Specifies whether the admin can create a custom report that includes data from the URL filtering log.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
Hipmatch	Specifies whether the admin can create a custom report that includes data from the HIP match log.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
View Scheduled Custom Reports	Specifies whether the admin can view a custom report that has been scheduled to generate.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
View Predefined Application Reports	Specifies whether the admin can view Application Reports. Privacy privileges do not impact reports available on the Monitor > Reports node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
View Predefined Threat Reports	Specifies whether the admin can view Threat Reports. Privacy privileges do not impact reports available on the Monitor > Reports node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes
View Predefined URL Filtering Reports	Specifies whether the admin can view URL Filtering Reports. Privacy privileges do not impact reports available on the Monitor > Reports node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
View Predefined Traffic Reports	Specifies whether the admin can view Traffic Reports. Privacy privileges do not impact reports available on the Monitor > Reports node and you should therefore disable access to the reports if you have user privacy requirements.	Firewall: Yes Panorama: Yes Device Group/Template: No	Yes	No	Yes

Provide Granular Access to the Policy Tab

If you enable the Policy option in the admin role profile, you can then enable, disable, or provide read-only access to specific nodes within the tab as necessary for the admin role you are defining. By enabling access to a specific policy type, you enable the ability to view, add, or delete policy rules. By enabling read-only access to a specific policy, you enable the admin to view the corresponding policy rule base, but not add or delete rules. Disabling access to a specific type of policy prevents the admin from seeing the policy rule base.

Because policy that is based on specific users (by user name or IP address) must be explicitly defined, privacy settings that disable the ability to see full IP addresses or user names do not apply to the Policy tab. Therefore, you should only allow access to the Policy tab to administrators that are excluded from user privacy restrictions.

Access Level	Description	Enable	Read Only	Disable
Security	Enable this privilege to allow the admin to view, add, and/or delete security policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the security policy rule base, disable this privilege.	Yes	Yes	Yes
NAT	Enable this privilege to allow the admin to view, add, and/or delete NAT policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the NAT policy rule base, disable this privilege.	Yes	Yes	Yes
QoS	Enable this privilege to allow the admin to view, add, and/or delete QoS policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the QoS policy rule base, disable this privilege.	Yes	Yes	Yes
Policy Based Forwarding	Enable this privilege to allow the admin to view, add, and/or delete Policy-Based Forwarding (PBF) policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the PBF policy rule base, disable this privilege.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
Decryption	Enable this privilege to allow the admin to view, add, and/or delete decryption policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the decryption policy rule base, disable this privilege.	Yes	Yes	Yes
Application Override	Enable this privilege to allow the admin to view, add, and/or delete application override policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the application override policy rule base, disable this privilege.	Yes	Yes	Yes
Captive Portal	Enable this privilege to allow the admin to view, add, and/or delete captive portal policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the captive portal policy rule base, disable this privilege.	Yes	Yes	Yes
DoS Protection	Enable this privilege to allow the admin to view, add, and/or delete DoS protection policy rules. Set the privilege to read only if you want the admin to be able to see the rules, but not modify them. To prevent the admin from seeing the DoS protection policy rule base, disable this privilege.	Yes	Yes	Yes

Provide Granular Access to the Objects Tab

An *object* is a container that groups specific policy filter values—such as IP addresses, URLs, applications, or services—for simplified rule definition. For example, an address object might contain specific IP address definitions for the web and application servers in your DMZ zone.

When deciding whether to allow access to the objects tab as a whole, determine whether the admin will have policy definition responsibilities. If not, the admin probably does not need access to the tab. If, however, the admin will need to create policy, you can enable access to the tab and then provide granular access privileges at the node level.

By enabling access to a specific node, you give the admin the privilege to view, add, and delete the corresponding object type. Giving read-only access allows the admin to view the already defined objects, but not create or delete any. Disabling a node prevents the admin from seeing the node in the web interface.

Access Level	Description	Enable	Read Only	Disable
Addresses	Specifies whether the admin can view, add, or delete address objects for use in security policy.	Yes	Yes	Yes
Address Groups	Specifies whether the admin can view, add, or delete address group objects for use in security policy.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
Regions	Specifies whether the admin can view, add, or delete regions objects for use in security, decryption, or DoS policy.	Yes	Yes	Yes
Applications	Specifies whether the admin can view, add, or delete application objects for use in policy.	Yes	Yes	Yes
Application Groups	Specifies whether the admin can view, add, or delete application group objects for use in policy.	Yes	Yes	Yes
Application Filters	Specifies whether the admin can view, add, or delete application filters for simplification of repeated searches.	Yes	Yes	Yes
Services	Specifies whether the admin can view, add, or delete service objects for use in creating policies that limit the port numbers an application can use.	Yes	Yes	Yes
Service Groups	Specifies whether the admin can view, add, or delete service group objects for use in security policy.	Yes	Yes	Yes
Tags (Panorama only)	Specifies whether the admin can view, add, or delete tags that have been defined on the device.	Yes	Yes	Yes
GlobalProtect	Specifies whether the admin can view, add, or delete HIP objects and profiles. You can restrict access to both types of objects at the GlobalProtect level, or provide more granular control by enabling the GlobalProtect privilege and restricting HIP Object or HIP Profile access.	Yes	No	Yes
HIP Objects	Specifies whether the admin can view, add, or delete HIP objects, which are used to define HIP profiles. HIP Objects also generate HIP Match logs.	Yes	Yes	Yes
HIP Profiles	Specifies whether the admin can view, add, or delete HIP Profiles for use in security policy and/or for generating HIP Match logs.	Yes	Yes	Yes
Dynamic Block Lists	Specifies whether the admin can view, add, or delete dynamic block lists for use in security policy.	Yes	Yes	Yes
Custom Objects	Specifies whether the admin can see the custom spyware and vulnerability signatures. You can restrict access to either enable or disable access to all custom signatures at this level, or provide more granular control by enabling the Custom Objects privilege and then restricting access to each type of signature.	Yes	No	Yes
Data Patterns	Specifies whether the admin can view, add, or delete custom data pattern signatures for use in creating custom vulnerability protection profiles.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
Spyware	Specifies whether the admin can view, add, or delete custom spyware signatures for use in creating custom vulnerability protection profiles.	Yes	Yes	Yes
Vulnerability	Specifies whether the admin can view, add, or delete custom vulnerability signatures for use in creating custom vulnerability protection profiles.	Yes	Yes	Yes
URL Category	Specifies whether the admin can view, add, or delete custom URL categories for use in policy.	Yes	Yes	Yes
Security Profiles	Specifies whether the admin can see security profiles. You can restrict access to either enable or disable access to all security profiles at this level, or provide more granular control by enabling the Security Profiles privilege and then restricting access to each type of profile.	Yes	No	Yes
Antivirus	Specifies whether the admin can view, add, or delete antivirus profiles.	Yes	Yes	Yes
Anti-Spyware	Specifies whether the admin can view, add, or delete anti-spyware profiles.	Yes	Yes	Yes
Vulnerability Protection	Specifies whether the admin can view, add, or delete vulnerability protection profiles.	Yes	Yes	Yes
URL Filtering	Specifies whether the admin can view, add, or delete URL filtering profiles.	Yes	Yes	Yes
File Blocking	Specifies whether the admin can view, add, or delete file blocking profiles.	Yes	Yes	Yes
Data Filtering	Specifies whether the admin can view, add, or delete data filtering profiles.	Yes	Yes	Yes
DoS Protection	Specifies whether the admin can view, add, or delete DoS protection profiles.	Yes	Yes	Yes
Security Profile Groups	Specifies whether the admin can view, add, or delete security profile groups.	Yes	Yes	Yes
Log Forwarding	Specifies whether the admin can view, add, or delete log forwarding profiles.	Yes	Yes	Yes
Decryption Profile	Specifies whether the admin can view, add, or delete decryption profiles.	Yes	Yes	Yes
Schedules	Specifies whether the admin can view, add, or delete schedules for limiting a security policy to a specific date and/or time range.	Yes	Yes	Yes

Provide Granular Access to the Network Tab

When deciding whether to allow access to the **Network** tab as a whole, determine whether the admin will have network administration responsibilities, including GlobalProtect administration. If not, the admin probably does not need access to the tab.

You can also define access to the **Network** tab at the node level. By enabling access to a specific node, you give the admin the privilege to view, add, and delete the corresponding network configurations. Giving read-only access allows the admin to view the already-defined configuration, but not create or delete any. Disabling a node prevents the admin from seeing the node in the web interface.


Access Level	Description	Enable	Read Only	Disable
Interfaces	Specifies whether the admin can view, add, or delete interface configurations.	Yes	Yes	Yes
Zones	Specifies whether the admin can view, add, or delete zones.	Yes	Yes	Yes
VLANs	Specifies whether the admin can view, add, or delete VLANs.	Yes	Yes	Yes
Virtual Wires	Specifies whether the admin can view, add, or delete virtual wires.	Yes	Yes	Yes
Virtual Routers	Specifies whether the admin can view, add, modify or delete virtual routers.	Yes	Yes	Yes
IPSec Tunnels	Specifies whether the admin can view, add, modify, or delete IPSec Tunnel configurations.	Yes	Yes	Yes
DHCP	Specifies whether the admin can view, add, modify, or delete DHCP server and DHCP relay configurations.	Yes	Yes	Yes
DNS Proxy	Specifies whether the admin can view, add, modify, or delete DNS proxy configurations.	Yes	Yes	Yes
GlobalProtect	Specifies whether the admin can view, add, modify GlobalProtect portal and gateway configurations. You can disable access to the GlobalProtect functions entirely, or you can enable the GlobalProtect privilege and then restrict the role to either the portal or gateway configuration areas.	Yes	No	Yes
Portals	Specifies whether the admin can view, add, modify, or delete GlobalProtect portal configurations.	Yes	Yes	Yes
Gateways	Specifies whether the admin can view, add, modify, or delete GlobalProtect gateway configurations.	Yes	Yes	Yes
MDM	Specifies whether the admin can view add, modify, or delete GlobalProtect MDM server configurations.	Yes	Yes	Yes
QoS		Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
Network Profiles	Sets the default state to enable or disable for all of the Network settings described below.	Yes	No	Yes
IKE Gateways	Controls access to the Network Profiles > IKE Gateways node. If you disable this privilege, the administrator will not see the IKE Gateways node or define gateways that include the configuration information necessary to perform IKE protocol negotiation with peer gateway. 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 profile configuration but cannot add or edit a configuration.	Yes	Yes	Yes
Interface Mgmt	Controls access to the Network Profiles > Interface Mgmt node. If you disable this privilege, the administrator will not see the Network Profiles > Interface Mgmt node or be able to specify the protocols that are used to manage the firewall. If the privilege state is set to read only, you can view the currently configured Interface management profile configuration but cannot add or edit a configuration.	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
Zone Protection	<p>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.</p> <p>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.</p>	Yes	Yes	Yes
QoS Profile	<p>Controls access to the Network Profiles > QoS node. If you disable this privilege, the administrator will not see the Network Profiles > QoS node or be able to configure a profile QoS profile that determines how QoS traffic classes are treated.</p> <p>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.</p>	Yes	Yes	Yes

Provide Granular Access to the Device Tab

Access Level	Description	Enable	Read Only	Disable
Setup	<p>Controls access to the Setup node. If you disable this privilege, the administrator will not see the Setup node or have access to device-wide setup configuration information, such as Management, Operations, Service, Content-ID, Wildfire or Session setup information.</p> <p>If the privilege state is set to read only, you can view the current configuration but cannot make any changes.</p>	Yes	Yes	Yes
Config Audit	<p>Controls access to the Config Audit node. If you disable this privilege, the administrator will not see the Config Audit node or have access to any device-wide configuration information.</p>	Yes	No	Yes
Admin Roles	<p>Controls access to the Admin Roles node. This function can only be allowed for read only access.</p> <p>If you disable this privilege, the administrator will not see the Admin Roles node or have access to any device-wide information concerning admin roles configuration.</p> <p>If you set this privilege to read only, you can view the configuration information for all admin roles configured on the device.</p>	No	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
Administrators	<p>Controls access to the Administrators node. This function can only be allowed for read only access.</p> <p>If you disable this privilege, the administrator will not see the Administrators node or have access to information about their own admin account.</p> <p>If you set this privilege to read only, the administrator can view the configuration information for their own admin account. They will not see any information about other admin accounts configured on the device.</p>	No	Yes	Yes
Virtual Systems	<p>Controls access to the Virtual Systems node. If you disable this privilege, the administrator will not see or be able to configure virtual systems.</p> <p>If the privilege state is set to read only, you can view the currently configured virtual systems but cannot add or edit a configuration.</p>	Yes	Yes	Yes
Shared Gateways	<p>Controls access to the Shared Gateways node. Shared gateways allow virtual systems to share a common interface for external communications.</p> <p>If you disable this privilege, the administrator will not see or be able to configure shared gateways.</p> <p>If the privilege state is set to read only, you can view the currently configured shared gateways but cannot add or edit a configuration.</p>	Yes	Yes	Yes
User Identification	<p>Controls access to the User Identification node. If you disable this privilege, the administrator will not see the User Identification node or have access to device-wide User Identification configuration information, such as User Mapping, User-ID Agents, Service, Terminal Services Agents, Group Mappings Settings or Captive Portal Settings.</p> <p>If you set this privilege to read only, the administrator can view configuration information for the device but is not allowed to perform any configuration procedures.</p>	Yes	Yes	Yes
VM Information Source	<p>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.</p> <p>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.</p> <p> This privilege is not available to Device Group and Template administrators.</p>	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
High Availability	<p>Controls access to the High Availability node. If you disable this privilege, the administrator will not see the High Availability node or have access to device-wide high availability configuration information such as General setup information or Link and Path Monitoring.</p> <p>If you set this privilege to read only, the administrator can view High Availability configuration information for the device but is not allowed to perform any configuration procedures.</p>	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	<p>Controls access to the Certificates node. If you disable this privilege, the administrator will not see the Certificates node or be able to configure or access information regarding Device Certificates or Default Trusted Certificate Authorities.</p> <p>If you set this privilege to read only, the administrator can view Certificate configuration information for the device but is not allowed to perform any configuration procedures.</p>	Yes	Yes	Yes
Certificate Profile	<p>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.</p> <p>If you set this privilege to read only, the administrator can view Certificate Profiles that are currently configured for the device but is not allowed to create or edit a certificate profile.</p>	Yes	Yes	Yes
OCSP Responder	<p>Controls access to the OCSP Responder node. If you disable this privilege, the administrator will not see the OCSP Responder node or be able to define a server that will be used to verify the revocation status of certificates issues by the PAN-OS device.</p> <p>If you set this privilege to read only, the administrator can view the OCSP Responder configuration for the device but is not allowed to create or edit an OCSP responder configuration.</p>	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
Response Pages	<p>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.</p> <p>If you set this privilege to read only, the administrator can view the Response Page configuration for the device but is not allowed to create or edit a response page configuration.</p>	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	<p>Controls access to the Log Settings > System node. If you disable this privilege, the administrator will not see the Log Settings > System node or be able to specify the severity levels of the system log entries that are logged remotely with Panorama and sent as SNMP traps, syslog messages, and/or email notifications.</p> <p>If you set this privilege to read only, the administrator can view the Log Settings > System configuration for the device but is not allowed to create or edit a configuration.</p>	Yes	Yes	Yes
Config	<p>Controls access to the Log Settings > Config node. If you disable this privilege, the administrator will not see the Log Settings > Config node or be able to specify the configuration log entries that are logged remotely with Panorama, and sent as syslog messages and/or email notification.</p> <p>If you set this privilege to read only, the administrator can view the Log Settings > Config configuration for the device but is not allowed to create or edit a configuration.</p>	Yes	Yes	Yes
HIP Match	<p>Controls access to the Log Settings > HIP Match node. If you disable this privilege, the administrator will not see the Log Settings > HIP Match node or be able to specify the Host Information Profile (HIP) match log settings that are used to provide information on security policies that apply to GlobalProtect clients</p> <p>If you set this privilege to read only, the administrator can view the Log Settings > HIP configuration for the device but is not allowed to create or edit a configuration.</p>	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
Alarms	<p>Controls access to the Log Settings > Alarms node. If you disable this privilege, the administrator will not see the Log Settings > Alarms node or be able to configure notifications that are generated when a security rule (or group of rules) has been hit repeatedly in a set period of time.</p> <p>If you set this privilege to read only, the administrator can view the Log Settings > Alarms configuration for the device but is not allowed to create or edit a configuration.</p>	Yes	Yes	Yes
Manage Logs	<p>Controls access to the Log Settings > Manage Logs node. If you disable this privilege, the administrator will not see the Log Settings > Manage Logs node or be able to clear the indicated logs.</p> <p>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.</p>	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	<p>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.</p> <p>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.</p>	Yes	Yes	Yes
Syslog	<p>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.</p> <p>If you set this privilege to read only, the administrator can view the Server Profiles > Syslog information but cannot specify syslog servers.</p>	Yes	Yes	Yes
Email	<p>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</p> <p>If you set this privilege to read only, the administrator can view the Server Profiles > Email information but cannot configure and email profile.</p>	Yes	Yes	Yes



Access Level	Description	Enable	Read Only	Disable
Netflow	<p>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.</p> <p>If you set this privilege to read only, the administrator can view the Server Profiles > Netflow information but cannot define a Netflow profile.</p>	Yes	Yes	Yes
RADIUS	<p>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.</p> <p>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.</p>	Yes	Yes	Yes
LDAP	<p>Controls access to the Server Profiles > LDAP node. If you disable this privilege, the administrator will not see the Server Profiles > LDAP node or be able to configure settings for the LDAP servers to use for authentication by way of authentication profiles.</p> <p>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.</p>	Yes	Yes	Yes
Kerberos	<p>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.</p> <p>If you set this privilege to read only, the administrator can view the Server Profiles > Kerberos information but cannot configure settings for Kerberos servers.</p>	Yes	Yes	Yes
Local User Database	Sets the default state to enable or disable for all of the Local User Database settings described below.	Yes	No	Yes
Users	<p>Controls access to the Local User Database > Users node. If you disable this privilege, the administrator will not see the Local User Database > Users node or set up a local database on the firewall to store authentication information for remote access users, device administrators, and captive portal users.</p> <p>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.</p>	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
User Groups	<p>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.</p> <p>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.</p>	Yes	Yes	Yes
Authentication Profile	<p>Controls access to the Authentication Profile node. If you disable this privilege, the administrator will not see the Authentication Profile node or be able to create or edit authentication profiles that specify local database, RADIUS, LDAP, or Kerberos settings that can be assigned to administrator accounts.</p> <p>If you set this privilege to read only, the administrator can view the Authentication Profile information but cannot create or edit an authentication profile.</p>	Yes	Yes	Yes
Authentication Sequence	<p>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.</p> <p>If you set this privilege to read only, the administrator can view the Authentication Profile information but cannot create or edit an authentication sequence.</p>	Yes	Yes	Yes
Access Domain	<p>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.</p> <p>If you set this privilege to read only, the administrator can view the Access Domain information but cannot create or edit an access domain.</p>	Yes	Yes	Yes
Scheduled Log Export	<p>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 to schedule exports of logs and save them to a File Transfer Protocol (FTP) server in CSV format or use Secure Copy (SCP) to securely transfer data between the device and a remote host.</p> <p>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.</p>	Yes	No	Yes

Access Level	Description	Enable	Read Only	Disable
Software	<p>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.</p> <p>If you set this privilege to read only, the administrator can view the Software information but cannot download or install software.</p>	Yes	Yes	Yes
GlobalProtect Client	<p>Controls access to the GlobalProtect Client node. If you disable this privilege, the administrator will not see the GlobalProtect Client node or view available GlobalProtect releases, download the code or activate the GlobalProtect agent.</p> <p>If you set this privilege to read only, the administrator can view the available GlobalProtect Client releases but cannot download or install the agent software.</p>	Yes	Yes	Yes
Dynamic Updates	<p>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.</p> <p>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.</p>	Yes	Yes	Yes
Licenses	<p>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.</p> <p>If you set this privilege to read only, the administrator can view the installed Licenses, but cannot perform license management functions.</p>	Yes	Yes	Yes
Support	<p>Controls access to the Support node. If you disable this privilege, the administrator will not see the Support node or be able to access product and security alerts from Palo Alto Networks or generate tech support or stats dump files.</p> <p>If you set this privilege to read only, the administrator can view the Support node and access product and security alerts but cannot generate tech support or stats dump files.</p>	Yes	Yes	Yes

Access Level	Description	Enable	Read Only	Disable
Master Key and Diagnostics	<p>Controls access to the Master Key and Diagnostics node. If you disable this privilege, the administrator will not see the Master Key and Diagnostics node or be able to specify a master key to encrypt private keys on the firewall.</p> <p>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.</p>	Yes	Yes	Yes

Define User Privacy Settings in the Admin Role Profile

Access Level	Description	Enable	Read Only	Disable
Privacy	Sets the default state to enable or disable for all of the privacy settings described below.	Yes	N/A	Yes
Show Full IP addresses	<p>When set to disable, full IP addresses obtained by traffic running through the Palo Alto firewall are not shown in logs or reports. In place of the IP addresses that are normally displayed, the relevant subnet is displayed.</p> <p> 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.</p>	Yes	N/A	Yes
Show User Names in Logs and Reports	<p>When set to disable, user names obtained by traffic running through the Palo Alto Networks firewall are not shown in logs or reports. Columns where the user names would normally be displayed are empty.</p> <p> Scheduled reports that are displayed in the interface through Monitor > Reports or reports that are sent via the email scheduler will still display user names. Because of this exception, we recommend that the following settings within the Monitor tab be set to disable: Custom Reports, Application Reports, Threat Reports, URL Filtering Reports, Traffic Reports and Email Scheduler.</p>	Yes	N/A	Yes

Access Level	Description	Enable	Read Only	Disable
View Pcap Files	When set to disable, packet capture files that are normally available within the Traffic, Threat and Data Filtering logs are not displayed.	Yes	N/A	Yes

Restrict Admin Access to Commit Functions

Restrict User Access Using the Commit Setting

Access Level	Description	Enable	Read Only	Disable
Commit	When set to disable, an admin cannot commit any changes to a configuration.	Yes	N/A	Yes

Provide Granular Access to Global Settings


Restrict User Access Using the Global Settings


Access Level	Description	Enable	Read Only	Disable
Global	Sets the default state to enable or disable for all of the global settings described below. In effect, this setting is only for System Alarms at this time.	Yes	N/A	Yes
System Alarms	When set to disable, an admin cannot view or acknowledge alarms that are generated.	Yes	N/A	Yes

Provide Granular Access to the Panorama Tab


The following table lists the **Panorama** tab access levels and the custom Panorama administrator roles for which they are available. Firewall administrators cannot access any of these privileges.



Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Setup	Specifies whether the administrator can view or edit Panorama setup information, such as Management, Operations, Services, WildFire, or HSM . If you set the privilege to read only, the administrator can see the information but cannot edit it. If you disable this privilege, the administrator cannot see or edit the information.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes



Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Templates	<p>Specifies whether the administrator can view, edit, add, or delete templates.</p> <p>If you set the privilege to read only, the administrator can see template configurations but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage template configurations.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: Yes</p>	<p>Yes</p> <p>(No for Device Group and Template admins)</p>	Yes	Yes
Config Audit	<p>Specifies whether the administrator can run Panorama configuration audits. If you disable this privilege, the administrator cannot run Panorama configuration audits.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: No</p>	Yes	No	Yes
Managed Devices	<p>Specifies whether the administrator can view, add, edit, tag, or delete firewalls as managed devices, and install software or content updates on them.</p> <p>If you set this privilege to read only, the administrator can see managed firewalls but cannot add, delete, tag, or install updates on them.</p> <p>If you disable this privilege, the administrator cannot view, add, edit, tag, delete, or install updates on managed firewalls.</p> <p> This privilege applies only to the Panorama > Managed Devices page. An administrator with Device Deployment privileges can still use the Panorama > Device Deployment pages to install updates on managed firewalls.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: Yes</p>	<p>Yes</p> <p>(No for Device Group and Template roles)</p>	Yes	Yes
Device Groups	<p>Specifies whether the administrator can view, edit, add, or delete device groups.</p> <p>If you set this privilege to read only, the administrator can see device group configurations but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage device group configurations.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: Yes</p>	<p>Yes</p> <p>(No for Device Group and Template roles)</p>	Yes	Yes



Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Managed Collectors	<p>Specifies whether the administrator can view, edit, add, or delete Managed Collectors.</p> <p>If you set this privilege to read only, the administrator can see Managed Collector configurations but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot view, edit, add, or delete Managed Collector configurations.</p> <p> This privilege applies only to the Panorama > Managed Collectors page. An administrator with Device Deployment privileges can still use the Panorama > Device Deployment pages to install updates on Log Collectors.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Collector Groups	<p>Specifies whether the administrator can view, edit, add, or delete Collector Groups.</p> <p>If you set this privilege to read only, the administrator can see Collector Groups but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage Collector Groups.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Admin Roles	<p>Specifies whether the administrator can view Panorama administrator roles.</p> <p>You cannot enable full access to this function: just read only access. (Only Panorama administrators with a dynamic role can add, edit, or delete custom Panorama roles.) With read only access, the administrator can see Panorama administrator role configurations but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage Panorama administrator roles.</p>	Panorama: Yes Device Group/Template: No	No	Yes	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Administrators	<p>Specifies whether the administrator can view Panorama administrator account details.</p> <p>You cannot enable full access to this function: just read only access. (Only Panorama administrators with a dynamic role can add, edit, or delete Panorama administrators.) With read only access, the administrator can see information about his or her own account but no other Panorama administrator accounts.</p> <p>If you disable this privilege, the administrator cannot see information about any Panorama administrator account, including his or her own.</p>	Panorama: Yes Device Group/Template: No	No	Yes	Yes
High Availability	<p>Specifies whether the administrator can view and manage high availability (HA) settings for the Panorama management server.</p> <p>If you set this privilege to read only, the administrator can view HA configuration information for the Panorama management server but cannot manage the configuration.</p> <p>If you disable this privilege, the administrator cannot see or manage HA configuration settings for the Panorama management server.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
VMware Service Manager	<p>Specifies whether the administrator can view and edit VMware Service Manager settings.</p> <p>If you set this privilege to read only, the administrator can see the settings but cannot perform any related configuration or operational procedures.</p> <p>If you disable this privilege, the administrator cannot see the settings or perform any related configuration or operational procedures.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Certificate Management	<p>Sets the default state, enabled or disabled, for all of the Panorama certificate management privileges.</p>	Panorama: Yes Device Group/Template: No	Yes	No	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Certificates	<p>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.</p> <p>If you set this privilege to read only, the administrator can see Panorama certificates but cannot manage the certificates or HA keys.</p> <p>If you disable this privilege, the administrator cannot see or manage Panorama certificates or HA keys.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Certificate Profile	<p>Specifies whether the administrator can view, add, edit, delete or clone Panorama certificate profiles.</p> <p>If you set this privilege to read only, the administrator can see Panorama certificate profiles but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage Panorama certificate profiles.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Log Settings	<p>Sets the default state, enabled or disabled, for all the log setting privileges.</p>	Panorama: Yes Device Group/Template: No	Yes	No	Yes
System	<p>Specifies whether the administrator can see and configure the settings that control the forwarding of system logs to external services (Syslog, email, or SNMP trap servers).</p> <p>If you set this privilege to read only, the administrator can see the system log forwarding settings but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage the settings.</p> <p> On a Panorama M-100 appliance, this privilege pertains only to system logs that Panorama generates for itself. On a Panorama virtual appliance, this privilege applies to system logs that Panorama generates and to system logs that Panorama collects from firewalls.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes


Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Config	<p>Specifies whether the administrator can see and configure the settings that control the forwarding of configuration logs to external services (Syslog, email, or SNMP trap servers).</p> <p>If you set this privilege to read only, the administrator can see the configuration log forwarding settings but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage the settings.</p> <p> On a Panorama M-100 appliance, this privilege pertains only to configuration logs that Panorama generates for itself. On a Panorama virtual appliance, this privilege applies to configuration logs that Panorama generates and to configuration logs that Panorama collects from firewalls.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: No</p>	Yes	Yes	Yes
HIP Match	<p>Specifies whether the administrator can see and configure the settings that control the forwarding of HIP match logs from a Panorama virtual appliance to external services (Syslog, email, or SNMP trap servers).</p> <p> The Panorama > Collector Groups page controls the forwarding of HIP match logs from a Panorama M-100 appliance. The Device > Log Settings > HIP Match page controls the forwarding of HIP match logs directly from firewalls to external services (without aggregation on Panorama).</p> <p>If you set this privilege to read only, the administrator can see the forwarding settings of HIP match logs but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage the settings.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: No</p>	Yes	Yes	Yes


Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Traffic	<p>Specifies whether the administrator can see and configure the settings that control the forwarding of traffic logs from a Panorama virtual appliance to external services (Syslog, email, or SNMP trap servers).</p> <p> The Panorama > Collector Groups page controls the forwarding of traffic logs from a Panorama M-100 appliance. The Objects > Log Forwarding page controls the forwarding of traffic logs directly from firewalls to external services (without aggregation on Panorama).</p> <p>If you set this privilege to read only, the administrator can see the forwarding settings of traffic logs but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage the settings.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Threat	<p>Specifies whether the administrator can see and configure the settings that control the forwarding of threat logs from a Panorama virtual appliance to external services (Syslog, email, or SNMP trap servers).</p> <p> The Panorama > Collector Groups page controls the forwarding of threat logs from a Panorama M-100 appliance. The Objects > Log Forwarding page controls the forwarding of threat logs directly from firewalls to external services (without aggregation on Panorama).</p> <p>If you set this privilege to read only, the administrator can see the forwarding settings of threat logs but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage the settings.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes


Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Wildfire	<p>Specifies whether the administrator can see and configure the settings that control the forwarding of WildFire logs from a Panorama virtual appliance to external services (Syslog, email, or SNMP trap servers).</p> <p> The Panorama > Collector Groups page controls the forwarding of WildFire logs from a Panorama M-100 appliance. The Objects > Log Forwarding page controls the forwarding of WildFire logs directly from firewalls to external services (without aggregation on Panorama).</p> <p>If you set this privilege to read only, the administrator can see the forwarding settings of WildFire logs but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage the settings.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Server Profiles	<p>Sets the default state, enabled or disabled, for all the server profile privileges.</p> <p> These privileges pertain only to the server profiles that are used for forwarding logs that Panorama generates or collects from firewalls and the server profiles that are used for authenticating Panorama administrators. The Device > Server Profiles pages control the server profiles that are used for forwarding logs directly from firewalls to external services (without aggregation on Panorama) and for authenticating firewall administrators.</p>	Panorama: Yes Device Group/Template: No	Yes	No	Yes
SNMP Trap	<p>Specifies whether the administrator can see and configure SNMP trap server profiles.</p> <p>If you set this privilege to read only, the administrator can see SNMP trap server profiles but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage SNMP trap server profiles.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Syslog	<p>Specifies whether the administrator can see and configure Syslog server profiles.</p> <p>If you set this privilege to read only, the administrator can see Syslog server profiles but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage Syslog server profiles.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: No</p>	Yes	Yes	Yes
Email	<p>Specifies whether the administrator can see and configure email server profiles.</p> <p>If you set this privilege to read only, the administrator can see email server profiles but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage email server profiles.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: No</p>	Yes	Yes	Yes
RADIUS	<p>Specifies whether the administrator can see and configure the RADIUS server profiles that are used to authenticate Panorama administrators.</p> <p>If you set this privilege to read only, the administrator can see the RADIUS server profiles but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage the RADIUS server profiles.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: No</p>	Yes	Yes	Yes
LDAP	<p>Specifies whether the administrator can see and configure the LDAP server profiles that are used to authenticate Panorama administrators.</p> <p>If you set this privilege to read only, the administrator can see the LDAP server profiles but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage the LDAP server profiles.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: No</p>	Yes	Yes	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Kerberos	<p>Specifies whether the administrator can see and configure the Kerberos server profiles that are used to authenticate Panorama administrators.</p> <p>If you set this privilege to read only, the administrator can see the Kerberos server profiles but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage the Kerberos server profiles.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Authentication Profile	<p>Specifies whether the administrator can view, add, edit, delete, or clone authentication profiles for Panorama administrators.</p> <p>If you set this privilege to read only, the administrator can view Panorama authentication profiles but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage Panorama authentication profiles.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Authentication Sequence	<p>Specifies whether the administrator can view, add, edit, delete, or clone authentication sequences for Panorama administrators.</p> <p>If you set this privilege to read only, the administrator can view Panorama authentication sequences but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage Panorama authentication sequences.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Access Domain	<p>Specifies whether the administrator can view, add, edit, delete, or clone access domains for Panorama administrators.</p> <p>If you set this privilege to read only, the administrator can view Panorama access domains but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage Panorama access domains.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Scheduled Config Export	<p>Specifies whether the administrator can view, add, edit, delete, or clone scheduled Panorama configuration exports.</p> <p>If you set this privilege to read only, the administrator can view the scheduled exports but cannot manage them.</p> <p>If you disable this privilege, the administrator cannot see or manage the scheduled exports.</p>	Panorama: Yes Device Group/Template: No	Yes	No	Yes
Software	<p>Specifies whether the administrator can: view information about Panorama software updates; download, upload, or install the updates; and view the associated release notes.</p> <p>If you set this privilege to read only, the administrator can view information about Panorama software updates and view the associated release notes but cannot perform any related operations.</p> <p>If you disable this privilege, the administrator cannot see Panorama software updates, see the associated release notes, or perform any related operations.</p> <p> This privilege pertains only to software installed on a Panorama management server. The Panorama > Device Deployment > Software page controls access to PAN-OS software deployed on firewalls and Panorama software deployed on dedicated Log Collectors.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Dynamic Updates	<p>Specifies whether the administrator can: view information about Panorama content updates (for example, WildFire updates); download, upload, install, or revert the updates; and view the associated release notes.</p> <p>If you set this privilege to read only, the administrator can view information about Panorama content updates and view the associated release notes but cannot perform any related operations.</p> <p>If you disable this privilege, the administrator cannot see Panorama content updates, see the associated release notes, or perform any related operations.</p> <p> This privilege pertains only to content updates installed on a Panorama management server. The Panorama > Device Deployment > Dynamic Updates page controls access to content updates deployed on firewalls and dedicated Log Collectors.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes
Support	<p>Specifies whether the administrator can: view Panorama support license information, product alerts, and security alerts; activate a support license, generate Tech Support files, and manage cases</p> <p>If you set this privilege to read only, the administrator can view Panorama support information, product alerts, and security alerts, but cannot activate a support license, generate Tech Support files, or manage cases.</p> <p>If you disable this privilege, the administrator cannot: see Panorama support information, product alerts, or security alerts; activate a support license, generate Tech Support files, or manage cases.</p>	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Device Deployment	<p>Sets the default state, enabled or disabled, for all the device deployment privileges.</p>  <p>These privilege pertain only to software and content updates that Panorama administrators deploy on firewalls and dedicated Log Collectors. The Panorama > Software and Panorama > Dynamic Updates pages control the software and content updates installed on a Panorama management server.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: Yes</p>	Yes	No	Yes
Software	<p>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.</p> <p>If you set this privilege to read only, the administrator can see information about the software updates and view the associated release notes but cannot deploy the updates to firewalls or dedicated Log Collectors.</p> <p>If you disable this privilege, the administrator cannot see information about the software updates, see the associated release notes, or deploy the updates to firewalls or dedicated Log Collectors.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: Yes</p>	Yes	Yes	Yes
SSL VPN Client	<p>Specifies whether the administrator can: view information about SSL VPN client software updates on firewalls; download, upload, or activate the updates; and view the associated release notes.</p> <p>If you set this privilege to read only, the administrator can see information about SSL VPN client software updates and view the associated release notes but cannot activate the updates on firewalls.</p> <p>If you disable this privilege, the administrator cannot see information about SSL VPN client software updates, see the associated release notes, or activate the updates on firewalls.</p>	<p>Panorama: Yes</p> <p>Device Group/Template: Yes</p>	Yes	Yes	Yes

Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
GlobalProtect Client	<p>Specifies whether the administrator can: view information about GlobalProtect client software updates on firewalls; download, upload, or activate the updates; and view the associated release notes.</p> <p>If you set this privilege to read only, the administrator can see information about GlobalProtect client software updates and view the associated release notes but cannot activate the updates on firewalls.</p> <p>If you disable this privilege, the administrator cannot see information about GlobalProtect client software updates, see the associated release notes, or activate the updates on firewalls.</p>	Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes
Dynamic Updates	<p>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.</p> <p>If you set this privilege to read only, the administrator can see information about the content updates and view the associated release notes but cannot deploy the updates to firewalls or dedicated Log Collectors.</p> <p>If you disable this privilege, the administrator cannot see information about the content updates, see the associated release notes, or deploy the updates to firewalls or dedicated Log Collectors.</p>	Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes
Licenses	<p>Specifies whether the administrator can view, refresh, and activate firewall licenses.</p> <p>If you set this privilege to read only, the administrator can view firewall licenses but cannot refresh or activate those licenses.</p> <p>If you disable this privilege, the administrator cannot view, refresh, or activate firewall licenses.</p>	Panorama: Yes Device Group/Template: Yes	Yes	Yes	Yes


Access Level	Description	Admin Role Availability	Enable	Read Only	Disable
Master Key and Diagnostics	Specifies whether the administrator can view and configure a master key by which to encrypt private keys on Panorama. If you set this privilege to read only, the administrator can view the Panorama master key configuration but cannot change it. If you disable this privilege, the administrator cannot see or edit the Panorama master key configuration.	Panorama: Yes Device Group/Template: No	Yes	Yes	Yes

Panorama Web Interface Access

The custom Panorama administrator roles allow you to define access to the options on Panorama and the ability to only allow access to Device Groups and Templates (**Policies, Objects, Network, Device** tabs).

The admin roles you can create are **Panorama** and **Device Group and Template**. You cannot assign CLI access privileges to a **Device Group and Template** admin role. If you assign superuser privileges for the CLI to a **Panorama** admin role, administrators with that role can access all features regardless of the web interface privileges you assign.

Access Level	Description	Enable	Read Only	Disable
Dashboard	Controls access to the Dashboard tab. If you disable this privilege, the administrator will not see the tab and will not have access to any of the Dashboard widgets.	Yes	No	Yes
ACC	Controls access to the Application Command Center (ACC). If you disable this privilege, the ACC tab will not display in the web interface. Keep in mind that if you want to protect the privacy of your users while still providing access to the ACC, you can disable the Privacy > Show Full Ip Addresses option and/or the Show User Names In Logs And Reports option.	Yes	No	Yes
Monitor	Controls access to the Monitor tab. If you disable this privilege, the administrator will not see the Monitor tab and will not have access to any of the logs, packet captures, session information, reports or to App Scope. For more granular control over what monitoring information the admin can see, leave the Monitor option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Monitor Tab .	Yes	No	Yes

Access Level	Description	Enable	Read Only	Disable
Policies	Controls access to the Policies tab. If you disable this privilege, the administrator will not see the Policies tab and will not have access to any policy information. For more granular control over what policy information the admin can see, for example to enable access to a specific type of policy or to enable read-only access to policy information, leave the Policies option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Policy Tab .	Yes	No	Yes
Objects	Controls access to the Objects tab. If you disable this privilege, the administrator will not see the Objects tab and will not have access to any objects, security profiles, log forwarding profiles, decryption profiles, or schedules. For more granular control over what objects the admin can see, leave the Objects option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Objects Tab .	Yes	No	Yes
Network	Controls access to the Network tab. If you disable this privilege, the administrator will not see the Network tab and will not have access to any interface, zone, VLAN, virtual wire, virtual router, IPsec tunnel, DHCP, DNS Proxy, GlobalProtect, or QoS configuration information or to the network profiles. For more granular control over what objects the admin can see, leave the Network option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Network Tab .	Yes	No	Yes
Device	Controls access to the Device tab. If you disable this privilege, the administrator will not see the Device tab and will not have access to any device-wide configuration information, such as User-ID, High Availability, server profile or certificate configuration information. For more granular control over what objects the admin can see, leave the Device option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Device Tab .  You cannot enable access to the Admin Roles or Administrators nodes for a role-based administrator even if you enable full access to the Device tab.	Yes	No	Yes


Access Level	Description	Enable	Read Only	Disable
Panorama	<p>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.</p> <p>For more granular control over what objects the admin can see, leave the Panorama option enabled and then enable or disable specific nodes on the tab as described in Provide Granular Access to the Panorama Tab.</p>	Yes	No	Yes

Reference: Port Numbers Used by Palo Alto Networks Devices

The following tables list the ports that Palo Alto Networks devices use to communicate with each other, or with other services on the network.

- ▲ [Ports Used for Management Functions](#)
- ▲ [Ports Used for HA](#)
- ▲ [Ports Used for Panorama](#)
- ▲ [Ports Used for User-ID](#)

Ports Used for Management Functions

Destination Port	Protocol	Description
22	TCP	Used for communication from a client system to the firewall CLI interface.
80	TCP	The port the firewall listens on for Online Certificate Status Protocol (OCSP) updates when acting as an OCSP responder.
123	UDP	Port the firewall uses for NTP updates.
443	TCP	Used for communication from a client system to the firewall web interface. This is also the port the firewall and/or the User-ID agent listens on for VM Information source updates. For monitoring an AWS environment, this is the only port that is used. For monitoring a VMware vCenter/ESXi environment, the listening port defaults to 443, but it is configurable.
162	UDP	Port the firewall, Panorama, or the log collector use to send SNMP traps to a SNMP trap receiver.  This port does not need to be open on the Palo Alto Networks device. The SNMP trap receiver must be configured to listen on this port.
161	UDP	Port the firewall listens on for SNMP polling requests from the NMS.
514	TCP	Port the firewall, Panorama, or the log collector use to send logs to a Syslog server and the ports the User-ID agent (on the firewall or on a Windows server) listens on for authentication Syslog messages for use with User-ID .
514	UDP	
6514	SSL	
2055	UDP	Default port the firewall uses to send NetFlow records, but this is configurable.
5008	TCP	Port the GlobalProtect Mobile Security Manager listens on for HIP requests from the GlobalProtect gateways . 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.

Destination Port	Protocol	Description
6081	TCP	Ports used for Captive Portal .
6082	TCP	

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	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.
28260	TCP	
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	IP	Used for the HA2 link to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between firewalls in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active device (Active/Passive) or active-primary (Active/Active) to the passive device (Active/Passive) or active-secondary (Active/Active). The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default. 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.
29281	UDP	

Ports Used for Panorama

Destination Port	Protocol	Description
22	TCP	Used for communication from a client system to the Panorama CLI interface.
443	TCP	Used for communication from a client system to the Panorama web interface.
3978	TCP	Used for communication between Panorama and managed devices and/or managed log collectors as well as for communication between log collectors in a collector group as follows: <ul style="list-style-type: none"> • For communication between Panorama and managed devices, this is a bi-directional connection on which the managed firewalls forward logs to Panorama and for Panorama to push configuration changes to the managed devices. Context switching commands are sent over the same connection. • Log collectors use this destination port to forward logs to Panorama. • For communication between the default log collector on a Panorama in Panorama mode, and for communicating with log collectors in a DLC architecture mode.
28769 (5.1 and later)	TCP	Used for the HA connectivity and synchronization between Panorama HA peers using clear text communication. Communication can be initiated by either peer.
28260 (5.0 and later)	TCP	
49160 (5.0 and earlier)	TCP	
28	TCP	Used for the HA connectivity and synchronization between Panorama HA peers using encrypted communication (SSH over TCP). Communication can be initiated by either peer.
28270 (6.0 and later)	TCP	Used for communication between log collectors in a collector group for log distribution.
49190 (5.1 and earlier)		
2049	TCP	Used by the Panorama virtual appliance to write logs to the NFS datastore.


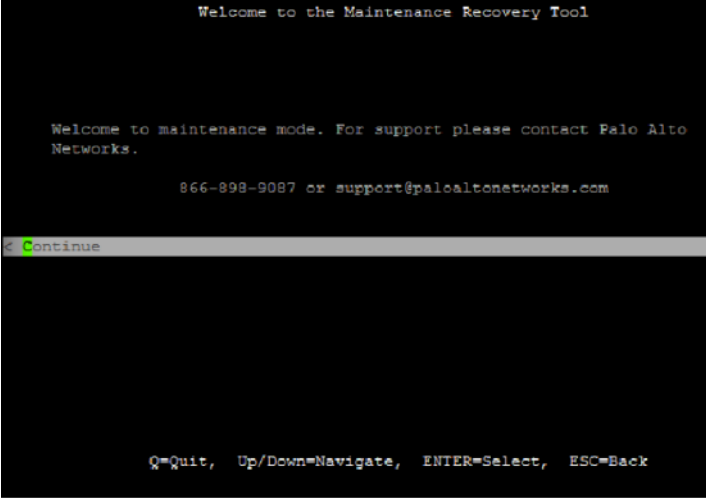
Ports Used for User-ID

User-ID is a feature that enables mapping of user IP addresses to usernames and group memberships, enabling user- or group-based policy and visibility into user activity on your network (for example, to be able to quickly track down a user who may be the victim of a threat). To perform this mapping, the firewall, the User-ID agent (either installed on a Windows-based system or the PAN-OS integrated agent running on the firewall), and/or the Terminal Services agent must be able to connect to directory services on your network to perform [Group Mapping](#) and [User Mapping](#). Additionally, if the agents are running on systems external to the firewall, they must be able to connect to the firewall to communicate the IP address to username mappings to the firewall. The following table lists the communication requirements for User-ID along with the port numbers required to establish connections.

Destination Port	Protocol	Description
389	TCP	Port the firewall uses to connect to the LDAP server (plaintext or StartTLS) in order to Map Users to Groups .
636	TCP	Port the firewall uses to connect to the LDAP server (LDAP over SSL) in order to Map Users to Groups . Used for LDAP over SSL connections.
514 514 6514	TCP UDP SSL	Ports the User-ID agent (on the firewall or on a Windows server) listens on for authentication Syslog messages for use with User-ID .
5007	TCP	Port the firewall listens on for user mapping information from the User-ID or Terminal Server agent. The agent sends the IP address and username mapping along with a timestamp whenever it learns of a new or updated mapping. In addition, it connects to the firewall at regular intervals to refresh known mappings.
5006	TCP	Port the User-ID agent listens on for User-ID XML API requests. The source for this communication is typically the system running a script that invokes the API.
88	UDP	Port the User-ID agent uses to authenticate to a Kerberos server.
1812	UDP	Port the User-ID agent uses to authenticate to a RADIUS server.
135	TCP	Port the User-ID agent uses to establish TCP-based WMI connections with the Microsoft Remote Procedure Call (RPC) Endpoint Mapper. The Endpoint Mapper then assigns the agent a randomly assigned port in the 49152-65535 port range. The agent uses this connection to make RPC queries for Exchange Server or AD server security logs, session tables. This is also the port used to access Terminal Services. The User-ID agent also uses this port to connect to client systems to perform WMI probes .
139	TCP	Port the User-ID agent uses to establish TCP-based NetBIOS connections to the AD server so that it can send RPC queries for security logs and session information. The User-ID agent also uses this port to connect to client systems for NetBIOS probing (supported on the Windows-based User-ID agent only).
445	TCP	Port the User-ID agent uses to connect to the Active Directory (AD) using TCP-based SMB connections to the AD server for access to user logon information (print spooler and Net Logon).

Reset the Firewall to Factory Default Settings

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

Reset the Firewall to Factory Default Settings	
<p>Step 1 Set up a console connection to the firewall.</p>	<ol style="list-style-type: none"> 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: <code>debug system maintenance-mode</code> The firewall will reboot in the maintenance mode.
<p>Step 2 Reset the system to factory default settings.</p>	<ol style="list-style-type: none"> 1. When the device 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.



Certificate Management

The following topics describe the different keys and certificates that Palo Alto Networks devices use, and how to obtain and manage them:

- ▲ [Keys and Certificates](#)
- ▲ [Certificate Revocation](#)
- ▲ [Certificate Deployment](#)
- ▲ [Set Up Verification for Certificate Revocation Status](#)
- ▲ [Configure the Master Key](#)
- ▲ [Obtain Certificates](#)
- ▲ [Configure a Certificate Profile](#)
- ▲ [Configure the Key Size for SSL Forward Proxy Server Certificates](#)
- ▲ [Revoke and Renew Certificates](#)
- ▲ [Secure Keys with a Hardware Security Module](#)

Keys and Certificates

To ensure trust between parties in a secure communication session, Palo Alto Networks devices use digital certificates. Each certificate contains a cryptographic key to encrypt plaintext or decrypt 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 devices use certificates in the following applications:

- User authentication for Captive Portal, GlobalProtect, Mobile Security Manager, and firewall/Panorama web interface access.
- Device authentication for GlobalProtect VPN (remote user-to-site or large scale).
- Device authentication for IPSec site-to-site VPN with Internet Key Exchange (IKE).
- Decrypting inbound and outbound SSL traffic. A firewall decrypts the traffic to apply security policies and rules, then re-encrypts it before forwarding the traffic to the final destination. For outbound traffic, the firewall acts as a forward proxy server, establishing an SSL/TLS connection to the destination server. To secure a connection between itself and the client, the firewall uses a *signing certificate* to automatically generate a copy of the destination server certificate.

The following table describes the keys and certificates that Palo Alto Networks devices use. As a best practice, use different keys/certificates for each usage.

Table: Palo Alto Networks Device Keys/Certificates

Key/Certificate Usage	Description
Administrative Access	Secure access to device administration interfaces (HTTPS access to the web interface) requires a server certificate for the MGT interface (or a designated interface on the dataplane if the device does not use MGT) and, optionally, a certificate to authenticate the administrator.
Captive Portal	In deployments where Captive Portal identifies users who access HTTPS resources, designate a server certificate for the Captive Portal interface. If you configure Captive Portal to use certificates (instead of, or in addition to, username/password credentials) for user identification, designate a user certificate also. For more information on Captive Portal, see Map IP Addresses to User Names Using Captive Portal .
Forward Trust	For outbound SSL/TLS traffic, if a firewall acting as a forward proxy trusts the CA that signed the certificate of the destination server, the firewall uses the forward trust CA certificate to generate a copy of the destination server certificate to present to the client. To set the key size, see Configure the Key Size for SSL Forward Proxy Server Certificates . For added security, store the key on a hardware security module (for details, see Secure Keys with a Hardware Security Module).
Forward Untrust	For outbound SSL/TLS traffic, if a firewall acting as a forward proxy does not trust the CA that signed the certificate of the destination server, the firewall uses the forward untrust CA certificate to generate a copy of the destination server certificate to present to the client.

Key/Certificate Usage	Description
SSL Inbound Inspection	The keys that decrypt inbound SSL/TLS traffic for inspection and policy enforcement. For this application, import onto the firewall a private key for each server that is subject to SSL/TLS inbound inspection. See Configure SSL Inbound Inspection .
SSL Exclude Certificate	Certificates for servers to exclude from SSL/TLS decryption. For example, if you enable SSL decryption but your network includes servers for which the firewall should not decrypt traffic (for example, web services for your HR systems), import the corresponding certificates onto the firewall and configure them as SSL Exclude Certificates. See Configure Decryption Exceptions .
GlobalProtect	All interaction among GlobalProtect components occurs over SSL/TLS connections. Therefore, as part of the GlobalProtect deployment, deploy server certificates for all GlobalProtect portals, gateways, and Mobile Security Managers. Optionally, deploy certificates for authenticating users also. Note that the GlobalProtect Large Scale VPN (LSVPN) feature requires a CA signing certificate.
Site-to-Site VPNs (IKE)	In a site-to-site IPSec VPN deployment, peer devices use Internet Key Exchange (IKE) gateways to establish a secure channel. IKE gateways use certificates or preshared keys to authenticate the peers to each other. You configure and assign the certificates or keys when defining an IKE gateway on a firewall. See Site-to-Site VPN Overview .
Master Key	The firewall uses a master key to encrypt all private keys and passwords. If your network requires a secure location for storing private keys, you can use an encryption (wrapping) key stored on a hardware security module (HSM) to encrypt the master key. For details, see Encrypt a Master Key Using an HSM .
Secure Syslog	The certificate to enable secure connections between the firewall and a syslog server. See Configure the Firewall to Authenticate to the Syslog Server .
Trusted Root CA	The designation for a root certificate issued by a CA that the firewall trusts. The firewall can use a self-signed root CA certificate to automatically issue certificates for other applications (for example, SSL Forward Proxy). Also, if a firewall must establish secure connections with other firewalls, the root CA that issues their certificates must be in the list of trusted root CAs on the firewall.

Certificate Revocation

Palo Alto Networks devices use digital certificates to ensure trust between parties in a secure communication session. Configuring a device to check the revocation status of certificates provides additional security. A party that presents a revoked certificate is not trustworthy. When a certificate is part of a chain, the device checks the status of every certificate in the chain except the root CA certificate, for which the device cannot verify revocation status.

Various circumstances can invalidate a certificate before the expiration date. Some examples are a change of name, change of association between subject and certificate authority (for example, an employee terminates employment), and compromise (known or suspected) of the private key. Under such circumstances, the certificate authority that issued the certificate must revoke it.

Palo Alto Networks devices support the following methods for verifying certificate revocation status. If you configure both, the devices first try the OCSP method; if the OCSP server is unavailable, the devices use the CRL method.

- ▲ [Certificate Revocation List \(CRL\)](#)
- ▲ [Online Certificate Status Protocol \(OCSP\)](#)



In PAN-OS, certificate revocation status verification is an optional feature. It is a best practice to enable it for certificate profiles, which define user and device authentication for Captive Portal, GlobalProtect, site-to-site IPsec VPN, and firewall/Panorama web interface access.

Certificate Revocation List (CRL)

Each certificate authority (CA) periodically issues a certificate revocation list (CRL) to a public repository. The CRL identifies revoked certificates by serial number. After the CA revokes a certificate, the next CRL update will include the serial number of that certificate.

The Palo Alto Networks firewall downloads and caches the last-issued CRL for every CA listed in the trusted CA list of the firewall. Caching only applies to validated certificates; if a firewall never validated a certificate, the firewall cache does not store the CRL for the issuing CA. Also, the cache only stores a CRL until it expires.

The firewall supports CRLs only in Distinguished Encoding Rules (DER) format. If the firewall downloads a CRL in any other format—for example, Privacy Enhanced Mail (PEM) format—any revocation verification process that uses that CRL will fail when a user performs an activity that triggers the process (for example, sending outbound SSL data). The firewall will generate a system log for the verification failure. If the verification was for an SSL certificate, the firewall will also display the SSL Certificate Errors Notify response page to the user.

To use CRLs for verifying the revocation status of certificates used for the decryption of inbound and outbound SSL/TLS traffic, see [Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption](#).

To use CRLs for verifying the revocation status of certificates that authenticate users and devices, configure a certificate profile and assign it to the interfaces that are specific to the application: Captive Portal, GlobalProtect (remote user-to-site or large scale), site-to-site IPsec VPN, or firewall/Panorama web interface access. For details, see [Configure Revocation Status Verification of Certificates Used for User/Device Authentication](#).

Online Certificate Status Protocol (OCSP)

When establishing an SSL/TLS session, clients can use Online Certificate Status Protocol (OCSP) to check the revocation status of the authentication certificate. The authenticating client sends a request containing the serial number of the certificate to the OCSP responder (server). The responder searches the database of the certificate authority (CA) that issued the certificate and returns a response containing the status (*good*, *revoked* or *unknown*) to the client. The advantage of the OCSP method is that it can verify status in real-time, instead of depending on the issue frequency (hourly, daily, or weekly) of CRLs.

The Palo Alto Networks firewall downloads and caches OCSP status information for every CA listed in the trusted CA list of the firewall. Caching only applies to validated certificates; if a firewall never validated a certificate, the firewall cache does not store the OCSP information for the issuing CA. If your enterprise has its own public key infrastructure (PKI), you can configure the firewall as an OCSP responder (see [Configure an OCSP Responder](#)).

To use OCSP for verifying the revocation status of certificates when the firewall functions as an SSL forward proxy, perform the steps under [Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption](#).

The following applications use certificates to authenticate users and/or devices: Captive Portal, GlobalProtect (remote user-to-site or large scale), site-to-site IPsec VPN, and firewall/Panorama web interface access. To use OCSP for verifying the revocation status of the certificates:

- Configure an OCSP responder.
- Enable the HTTP OCSP service on the firewall.
- Create or obtain a certificate for each application.
- Configure a certificate profile for each application.
- Assign the certificate profile to the relevant application.

To cover situations where the OCSP responder is unavailable, configure CRL as a fall-back method. For details, see [Configure Revocation Status Verification of Certificates Used for User/Device Authentication](#).

Certificate Deployment

The basic approaches to deploy certificates for Palo Alto Networks devices are:

- **Obtain certificates from a trusted third-party CA**—The benefit of obtaining a certificate from a trusted third-party certificate authority (CA) such as VeriSign or GoDaddy is that end clients will already trust the certificate because common browsers include root CA certificates from well-known CAs in their trusted root certificate stores. Therefore, for applications that require end clients to establish secure connections with a Palo Alto Network device, purchase a certificate from a CA that the end clients trust to avoid having to pre-deploy root CA certificates to the end clients. (Some such applications are a GlobalProtect portal or GlobalProtect Mobile Security Manager.) However, note that most third-party CAs cannot issue signing certificates. Therefore, this type of certificate is not appropriate for applications (for example, SSL/TLS decryption and large-scale VPN) that require the firewall to issue certificates. See [Obtain a Certificate from an External CA](#).
- **Obtain certificates from an enterprise CA**—Enterprises that have their own internal CA can use it to issue certificates for firewall applications and import them onto the firewall. The benefit is that end clients probably already trust the enterprise CA. You can either generate the needed certificates and import them onto the firewall, or generate a certificate signing request (CSR) on the firewall and send it to the enterprise CA for signing. The benefit of this method is that the private key does not leave the firewall. An enterprise CA can also issue a signing certificate, which the firewall uses to automatically generate certificates (for example, for GlobalProtect large-scale VPN or sites requiring SSL/TLS decryption). See [Import a Certificate and Private Key](#).
- **Generate self-signed certificates**—You can [Create a Self-Signed Root CA Certificate](#) on the firewall and use it to automatically issue certificates for other firewall applications. Note that if you use this method to generate certificates for an application that requires an end client to trust the certificate, end users will see a certificate error because the root CA certificate is not in their trusted root certificate store. To prevent this, deploy the self-signed root CA certificate to all end user systems. You can deploy the certificates manually or use a centralized deployment method such as an Active Directory Group Policy Object (GPO).

Set Up Verification for Certificate Revocation Status

To verify the revocation status of certificates, the firewall uses Online Certificate Status Protocol (OCSP) and/or certificate revocation lists (CRLs). For details on these methods, see [Certificate Revocation](#). If you configure both methods, the firewall first tries OCSP and only falls back to the CRL method if the OCSP responder is unavailable. If your enterprise has its own public key infrastructure (PKI), you can configure the firewall to function as the OCSP responder.

The following topics describe how to configure the firewall to verify certificate revocation status:

- ▲ [Configure an OCSP Responder](#)
- ▲ [Configure Revocation Status Verification of Certificates Used for User/Device Authentication](#)
- ▲ [Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption](#)

Configure an OCSP Responder

To use Online Certificate Status Protocol (OCSP) for verifying the revocation status of certificates, you must configure the firewall to access an OCSP responder (server). The entity that manages the OCSP responder can be a third-party certificate authority (CA) or, if your enterprise has its own public key infrastructure (PKI), the firewall itself. For details on OCSP, see [Certificate Revocation](#).

Configure an OCSP Responder	
<p>Step 1 Define an OCSP responder.</p>	<ol style="list-style-type: none"> 1. In a firewall, select Device > Certificate Management > OCSP Responder and click Add. In Panorama, select Device > Certificate Management > OCSP Responder, select a Template 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 supports multiple virtual systems, the dialog displays a Location drop-down. Select the virtual system where the responder will be available or select Shared to enable availability on all the virtual systems. 4. In the Host Name field, enter the host name (recommended) or IP address of the OCSP responder. From this value, PAN-OS automatically derives a URL and adds it to the certificate being verified. If you configure the firewall itself as an OCSP responder, the host name must resolve to an IP address in the interface that the firewall uses for OCSP services (specified in Step 3). 5. Click OK.

Configure an OCSP Responder	
<p>Step 2 Enable OCSP communication on the firewall.</p>	<ol style="list-style-type: none"> 1. In a firewall, select Device > Setup > Management. In Panorama, select Device > Setup > Management and select a Template. 2. In the Management Interface Settings section, edit to select the HTTP OCSP check box, then click OK.
<p>Step 3 Optionally, to configure the firewall itself as an OCSP responder, add an Interface Management Profile to the interface used for OCSP services.</p>	<ol style="list-style-type: none"> 1. Select Network > Network Profiles > Interface Mgmt. 2. Click Add to create a new profile or click the name of an existing profile. 3. Select the HTTP OCSP check box and click OK. 4. Select Network > Interfaces and click the name of the interface that the firewall will use for OCSP services. The OCSP Host Name specified in Step 1 must resolve to an IP address in this interface. 5. Select Advanced > Other info and select the Interface Management Profile you configured. 6. Click OK and Commit.

Configure Revocation Status Verification of Certificates Used for User/Device Authentication

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

Configure Revocation Status Verification of Certificates Used for User/Device Authentication	
<p>Step 1 Configure a Certificate Profile for each application.</p>	<p>Assign one or more root CA certificates to the profile and select how the firewall verifies certificate revocation status. The common name (FQDN or IP address) of a certificate must match an interface to which you apply the profile in Step 2.</p> <p>For details on the certificates that various applications use, see Keys and Certificates</p>
<p>Step 2 Assign the certificate profiles to the relevant applications.</p>	<p>The steps to assign a certificate profile depend on the application that requires it.</p>

Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption

The firewall decrypts inbound and outbound SSL/TLS traffic to apply security policies and rules, then re-encrypts the traffic before forwarding it. (For details, see [SSL Inbound Inspection](#) and [SSL Forward Proxy](#).) You can configure the firewall to verify the revocation status of certificates used for decryption as follows.



Enabling revocation status verification for SSL/TLS decryption certificates will add time to the process of establishing the session. The first attempt to access a site might fail if the verification does not finish before the session times out. For these reasons, verification is disabled by default.

Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption

<p>Step 1 Access the Decryption Certificate Revocation Settings page.</p>	<p>In a firewall, select Device > Setup > Session and, in the Session Features section, select Decryption Certificate Revocation Settings.</p> <p>In Panorama, select Device > Setup > Session, select a Template and, in the Session Features section, select Decryption Certificate Revocation Settings.</p>
<p>Step 2 Define the service-specific timeout intervals for revocation status requests.</p>	<p>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.</p> <ol style="list-style-type: none"> 1. In the CRL section, select the Enable check box and enter the Receive Timeout. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from the CRL service. 2. In the OCSP section, select the Enable check box and enter the Receive Timeout. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from the OCSP responder. <p>Depending on the Certificate Status Timeout value you specify in Step 3, the firewall might register a timeout before either or both of the Receive Timeout intervals pass.</p>

Configure Revocation Status Verification of Certificates Used for SSL/TLS Decryption (Continued)

<p>Step 3 Define the total timeout interval for revocation status requests.</p>	<p>Enter the Certificate Status Timeout. This is the interval (1-60 seconds) after which the firewall stops waiting for a response from any certificate status service and applies the session-blocking logic you optionally define in Step 4. The Certificate Status Timeout relates to the OCSP/CRL Receive Timeout as follows:</p> <ul style="list-style-type: none"> • 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.
<p>Step 4 Define the blocking behavior for <i>unknown</i> certificate status or a revocation status request timeout.</p>	<p>If you want the firewall to block SSL/TLS sessions when the OCSP or CRL service returns a certificate revocation status of <i>unknown</i>, select the Block Session With Unknown Certificate Status check box. Otherwise, the firewall proceeds with the session.</p> <p>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.</p>
<p>Step 5 Save and apply your entries.</p>	<p>Click OK and Commit.</p>

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).

If you deploy firewalls or Panorama in a high availability (HA) configuration, use the same master key on both HA peers. Otherwise, HA synchronization will not work properly.



If you use Panorama, use the same master key on Panorama and all managed firewalls. Otherwise, Panorama cannot push configurations to the firewalls.



As a best practice, configure a new master key instead of using the default, periodically change the key, and store the key in a safe location. For extra security you can also [Encrypt a Master Key Using an HSM](#).

The only way to restore the default master key is to [Reset the Firewall to Factory Default Settings](#).

Configure a Master Key

1. Select **Device > Master Key and Diagnostics** and edit the Master Key settings.
 2. Enter the **Current Master Key** if one exists.
 3. Define a new **New Master Key** and then **Confirm New Master Key**. The key must contain exactly 16 characters.
 4. To specify the master key **Life Time**, enter the number of **Days** and/or **Hours** after which the key will expire.
 You must configure a new master key before the current key expires. If the master key expires, the firewall or Panorama automatically reboots in Maintenance mode. You must then [Reset the Firewall to Factory Default Settings](#).
 5. Enter a **Time for Reminder** that specifies the number of **Days** and **Hours** before the master key expires when the firewall generates an expiration alarm. The firewall automatically opens the System Alarms dialog to display the alarm.
 To ensure the expiration alarm displays, select **Device > Log Settings**, edit the Alarm Settings, and **Enable Alarms**.
 6. (Optional) Select whether to use an **HSM** to encrypt the master key. For details, see [Encrypt a Master Key Using an HSM](#).
 7. Click **OK** and **Commit**.
-

Obtain Certificates

- ▲ [Create a Self-Signed Root CA Certificate](#)
- ▲ [Generate a Certificate on the Firewall](#)
- ▲ [Import a Certificate and Private Key](#)
- ▲ [Obtain a Certificate from an External CA](#)

Create a Self-Signed Root CA Certificate

A self-signed root certificate authority (CA) certificate is the top-most certificate in a certificate chain. A firewall can use this certificate to automatically issue certificates for other uses. For example, the firewall issues certificates for SSL/TLS decryption and for satellite devices in a GlobalProtect large-scale VPN.

When establishing a secure connection with the firewall, the remote client must trust the root CA that issued the certificate. Otherwise, the client browser will display a warning that the certificate is invalid and might (depending on security settings) block the connection. To prevent this, after generating the self-signed root CA certificate, import it into the client systems.

Generate a Self-signed Root CA Certificate

1. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.
In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**.
 2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select a virtual system for the certificate. To make the certificate available to all virtual systems, select the shared option described in [Step 6](#).
 3. Click **Generate**.
 4. Enter a **Certificate Name**, such as *GlobalProtect_CA*. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.
 5. In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.
 6. To make the certificate available to all virtual systems, select the **Shared** check box. This check box appears only if the device supports multiple virtual systems.
 7. Leave the **Signed By** field blank to designate the certificate as self-signed.
 8. Select the **Certificate Authority** check box.
 9. Do *not* select an **OCSP Responder**. Certificate revocation status verification does not apply to root CA certificates.
 10. Click **Generate** and **Commit**.
-

Generate a Certificate on the Firewall


The firewall uses certificates to authenticate clients, servers, users and devices in several applications, including SSL/TLS decryption, Captive Portal, GlobalProtect, site-to-site IPSec VPN, and firewall/Panorama web interface access. Generate certificates for each usage. For details on application-specific certificates, see [Keys and Certificates](#)


To generate a certificate, you must first create or import a root CA certificate to sign it. For details, see [Create a Self-Signed Root CA Certificate](#) or [Import a Certificate and Private Key](#).

To use Online Certificate Status Protocol (OCSP) for verifying certificate revocation status, [Configure an OCSP Responder](#) *before* generating the certificate. For details on status verification, see [Certificate Revocation](#)

Generate a Certificate on the Firewall

1. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.
In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**.
 2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select a virtual system for the certificate. To make the certificate available to all virtual systems, select the shared option described in [Step 6](#).
 3. Click **Generate**.
 4. Enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.
 5. In the **Common Name** field, enter the FQDN (recommended) or IP address of the interface where you will configure the service that will use this certificate.
 6. To make the certificate available to all virtual systems, select the **Shared** check box. This check box appears only if the device supports multiple virtual systems.
 7. In the **Signed By** field, select the root CA certificate that will issue the certificate.
 8. If applicable, select an **OCSP Responder**.
 9. (Optional) Define the **Cryptographic Settings** as necessary to create a certificate that will work with the devices that must authenticate to it. The default and recommended key size (**Number of Bits**) is 2048 bits. The default and recommended encryption algorithm (**Digest**) is SHA256.
 10. (Optional) **Add** the **Certificate Attributes** to uniquely identify the firewall and the service that will use the certificate.

 If you add a **Host Name** (DNS name) attribute, it is a best practice for it to match the **Common Name**. The host name populates the Subject Alternative Name field of the certificate.
 11. Click **Generate** and, in the Device Certificates tab, click the certificate **Name**.

 Regardless of the time zone on the firewall, it always displays the corresponding Greenwich Mean Time (GMT) for certificate validity and expiration dates/times.
 12. Select the check boxes that correspond to the intended use of the certificate on the firewall. For example, if the firewall will use this certificate to authenticate user access to its web interface, select the **Certificate for Secure Web GUI** check box.
 13. Click **OK** and **Commit**.
-

Import a Certificate and Private Key

If your enterprise has its own public key infrastructure (PKI), you can import a certificate and private key into the firewall from your enterprise certificate authority (CA). Enterprise CA certificates (unlike most certificates purchased from a trusted, third-party CA) can automatically issue CA certificates for applications such as SSL/TLS decryption or large-scale VPN.



Instead of importing a self-signed root CA certificate into all the client systems, it is a best practice to import a certificate from the enterprise CA because the clients will already have a trust relationship with the enterprise CA, which simplifies the deployment. If the certificate you will import is part of a certificate chain, it is a best practice to import the entire chain.


Import a Certificate and Private Key

1. From the enterprise CA, export the certificate and private key that the firewall will use for authentication. When exporting a private key, you must enter a passphrase to encrypt the key for transport. Ensure the management system can access the certificate and key files. When importing the key onto the firewall, you must enter the same passphrase to decrypt it.
2. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**. In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**.
3. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select a virtual system for the certificate. To make the certificate available to all virtual systems, select the shared option described in [Step 6](#).
4. Click **Import**.
5. Enter a **Certificate Name**. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores.
6. To make the certificate available to all virtual systems, select the **Shared** check box. This check box appears only if the device supports multiple virtual systems.
7. Enter the path and name of the **Certificate File** received from the CA, or **Browse** to find the file.
8. Select a **File Format**:
 - **Encrypted Private Key and Certificate (PKCS12)**—This is the default and most common format, in which the key and certificate are in a single container (**Certificate File**). If a hardware security module (HSM) will store the private key for this certificate, select the **Private key resides on Hardware Security Module** check box.
 - **Base64 Encoded Certificate (PEM)**—You must import the key separately from the certificate. If a hardware security module (HSM) stores the private key for this certificate, select the **Private key resides on Hardware Security Module** check box and skip Step 9. Otherwise, select the **Import Private Key** check box, enter the **Key File** or **Browse** to it, then perform Step 9.
9. Enter and re-enter (confirm) the **Passphrase** used to encrypt the private key.
10. Click **OK**. The Device Certificates tab displays the imported certificate.

Obtain a Certificate from an External CA

The advantage of obtaining a certificate from an external certificate authority (CA) is that the private key does not leave the firewall. To obtain a certificate from an external CA, generate a certificate signing request (CSR) and submit it to the CA. After the CA issues a certificate with the specified attributes, import it onto the firewall. The CA can be a well-known, public CA or an enterprise CA.

To use Online Certificate Status Protocol (OCSP) for verifying the revocation status of the certificate, [Configure an OCSP Responder](#) *before* generating the CSR.

Obtain a Certificate from an External CA	
<p>Step 1 Request the certificate from an external CA.</p>	<ol style="list-style-type: none"> In a firewall, select Device > Certificate Management > Certificates > Device Certificates. In Panorama, select Device > Certificate Management > Certificates > Device Certificates and select a Template. If the device supports multiple virtual systems, the tab displays a Location drop-down. Select a virtual system for the certificate. To make the certificate available to all virtual systems, select the shared option described in sub-step 6. Click Generate. Enter a Certificate Name. The name is case-sensitive and can have up to 31 characters. It must be unique and use only letters, numbers, hyphens, and underscores. 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. To make the certificate available to all virtual systems, select the Shared check box. This check box appears only if the device supports multiple virtual systems. In the Signed By field, select External Authority (CSR). If applicable, select an OCSP Responder. (Optional) Add the Certificate Attributes to uniquely identify the firewall and the service that will use the certificate.  If you add a Host Name attribute, it is a best practice for it to match the Common Name (this is mandatory for GlobalProtect). The host name populates the Subject Alternative Name field of the certificate. Click Generate. The Device Certificates tab displays the CSR with a Status of <i>pending</i>.
<p>Step 2 Submit the CSR to the CA.</p>	<ol style="list-style-type: none"> Select the CSR and click Export to save the .csr file to a local computer. Upload the .csr file to the CA.

Obtain a Certificate from an External CA	
Step 3 Import the certificate.	<ol style="list-style-type: none">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 in Step 1-4.3. 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 <i>valid</i>.
Step 4 Configure the certificate.	<ol style="list-style-type: none">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 authenticate administrators who access the web interface, select the Certificate for Secure Web GUI check box.3. Click OK and Commit.

Configure a Certificate Profile

Certificate profiles define user and device authentication for Captive Portal, GlobalProtect, site-to-site IPsec VPN, Mobile Security Manager, and firewall/Panorama web interface access. The profiles specify which certificates to use, how to verify certificate revocation status, and how that status constrains access. Configure a certificate profile for each application.



It is a best practice to enable Online Certificate Status Protocol (OCSP) and/or Certificate Revocation List (CRL) status verification for certificate profiles. For details on these methods, see [Certificate Revocation](#).

Configure a Certificate Profile	
<p>Step 1 Obtain the certificate authority (CA) certificates you will assign.</p>	<p>Perform one of the following steps to obtain the CA certificates you will assign to the profile. You must assign at least one.</p> <ul style="list-style-type: none"> • Generate a Certificate on the Firewall. • Export a certificate from your enterprise CA and then import it onto the firewall (see Step 3).
<p>Step 2 Identify the certificate profile.</p>	<ol style="list-style-type: none"> 1. In a firewall, select Device > Certificate Management > Certificates Profile and click Add. In Panorama, select Device > Certificate Management > Certificates Profile, select a Template and click Add. 2. Enter a Name to identify the profile. The name is case-sensitive. It must be unique and use only letters, numbers, spaces, hyphens, and underscores. It can have up to 31 characters. 3. If the firewall supports multiple virtual systems, the dialog displays a Location drop-down. Select the virtual system where the profile will be available or select Shared to enable availability on all the virtual systems.

Configure a Certificate Profile

Step 3 Assign one or more certificates.

Perform the following steps for each certificate:

1. In the CA Certificates table, click **Add**.
2. Select a **CA Certificate** from [Step 1](#), or click **Import** and perform the following sub-steps:
 - a. Enter a **Certificate Name**.
 - b. Enter the path and name of the **Certificate File** you exported from your enterprise CA, or **Browse** to find the file.
 - c. Click **OK**.
3. Optionally, if the firewall uses OCSP to verify certificate revocation status, configure the following fields to override the default behavior. For most deployments, these fields do not apply.
 - By default, the firewall uses the OCSP responder URL that you set in the procedure [Configure an OCSP Responder](#). To override that setting, enter a **Default OCSP URL** (starting with *http://* or *https://*).
 - By default, the firewall uses the certificate selected in the **CA Certificate** field to validate OCSP responses. To use a different certificate for validation, select it in the **OCSP Verify CA Certificate** field.
4. Click **OK**. The CA Certificates table displays the assigned certificate.

Configure a Certificate Profile	
<p>Step 4 Define the methods for verifying certificate revocation status and the associated blocking behavior.</p>	<ol style="list-style-type: none"> 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: <ul style="list-style-type: none"> • 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 <i>unknown</i>, select the Block session if certificate status is unknown check box. Otherwise, the firewall proceeds with the session. 5. If you want the firewall to block sessions after it registers an OCSP or CRL request timeout, select the Block session if certificate status cannot be retrieved within timeout check box. Otherwise, the firewall proceeds with the session.
<p>Step 5 Save and apply your entries.</p>	<p>Click OK and Commit.</p>

Configure the Key Size for SSL Forward Proxy Server Certificates

When responding to a client in an [SSL Forward Proxy](#) session, the firewall creates a copy of the certificate that the destination server presents and uses the copy to establish a connection with the client. By default, the firewall generates certificates with the same key size as the certificate that the destination server presented. However, you can change the key size for the firewall-generated certificate as follows:

Configure the Key Size for SSL Forward Proxy Server Certificates

Step 1 Select **Device > Setup > Session** and, in the Decryption Settings section, click **Forward Proxy Server Certificate Settings**.

Step 2 Select a **Key Size**:

- **Defined by destination host**—The firewall determines the key size for the certificates it generates to establish SSL proxy sessions with clients based on the key size of the destination server certificate. If the destination server uses a 1,024-bit RSA key, the firewall generates a certificate with that key size and an SHA-1 hashing algorithm. If the destination server uses a key size larger than 1,024 bits (for example, 2,048 bits or 4,096 bits), the firewall generates a certificate that uses a 2,048-bit RSA key and SHA-256 algorithm. This is the default setting.
- **1024-bit RSA**—The firewall generates certificates that use a 1,024-bit RSA key and SHA-1 hashing algorithm regardless of the key size of the destination server certificates. As of December 31, 2013, public certificate authorities (CAs) and popular browsers have limited support for X.509 certificates that use keys of fewer than 2,048 bits. In the future, depending on security settings, when presented with such keys the browser might warn the user or block the SSL/TLS session entirely.
- **2048-bit RSA**—The firewall generates certificates that use a 2,048-bit RSA key and SHA-256 hashing algorithm regardless of the key size of the destination server certificates. Public CAs and popular browsers support 2,048-bit keys, which provide better security than the 1,024-bit keys.



Changing the key size setting clears the current certificate cache.

Step 3 Click **OK** and **Commit**.

Revoke and Renew Certificates

▲ [Revoke a Certificate](#)

▲ [Renew a Certificate](#)

Revoke a Certificate

Various circumstances can invalidate a certificate before the expiration date. Some examples are a change of name, change of association between subject and certificate authority (for example, an employee terminates employment), and compromise (known or suspected) of the private key. Under such circumstances, the certificate authority (CA) that issued the certificate must revoke it. The following task describes how to revoke a certificate for which the firewall is the CA.

Revoke a Certificate

1. Select **Device > Certificate Management > Certificates > Device Certificates**.
2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select the virtual system to which the certificate belongs.
3. Select the certificate to revoke.
4. Click **Revoke**. PAN-OS immediately sets the status of the certificate to revoked and adds the serial number to the Online Certificate Status Protocol (OCSP) responder cache or certificate revocation list (CRL). You need not perform a commit.

Renew a Certificate

If a certificate expires, or soon will, you can reset the validity period. If an external certificate authority (CA) signed the certificate and the firewall uses the Online Certificate Status Protocol (OCSP) to verify certificate revocation status, the firewall uses the OCSP responder information to update the certificate status (see [Configure an OCSP Responder](#)). If the firewall is the CA that issued the certificate, the firewall replaces it with a new certificate that has a different serial number but the same attributes as the old certificate.

Renew a Certificate

1. In a firewall, select **Device > Certificate Management > Certificates > Device Certificates**.
In Panorama, select **Device > Certificate Management > Certificates > Device Certificates** and select a **Template**.
2. If the device supports multiple virtual systems, the tab displays a **Location** drop-down. Select the virtual system to which the certificate belongs.
3. Select a certificate to renew and click **Renew**.
4. Enter a **New Expiration Interval** (in days).
5. Click **OK** and **Commit**.

Secure Keys with a Hardware Security Module

A hardware security module (HSM) is a physical device that manages digital keys. An HSM provides secure storage and generation of digital keys. It provides both logical and physical protection of these materials from non-authorized use and potential adversaries.

HSM clients integrated with Palo Alto Networks devices enable enhanced security for the private keys used in SSL/TLS decryption (both SSL forward proxy and SSL inbound inspection). In addition, you can use the HSM to encrypt device master keys.

The following topics describe how to integrate an HSM with your Palo Alto Networks devices:

- ▲ [Set up Connectivity with an HSM](#)
- ▲ [Encrypt a Master Key Using an HSM](#)
- ▲ [Store Private Keys on an HSM](#)
- ▲ [Manage the HSM Deployment](#)

Set up Connectivity with an HSM

HSM clients are integrated with PA-3000 Series, PA-4000 Series, PA-5000 Series, PA-7050, and VM-Series firewalls and on Panorama (virtual appliance and M-100 appliance) for use with the following HSMs:

- SafeNet Network 5.2.1
- Thales nShield Connect 11.62 or later



The HSM server version must be compatible with these client versions. Refer to the HSM vendor documentation for the client-server version compatibility matrix.

The following topics describe how to set up connectivity between the firewall/Panorama and one of the supported HSMs:

- ▲ [Set Up Connectivity with a SafeNet Network HSM](#)
- ▲ [Set Up Connectivity with a Thales nShield Connect HSM](#)

Set Up Connectivity with a SafeNet Network HSM

To set up connectivity between the Palo Alto Networks device and a SafeNet Network HSM, you must specify the address of the HSM server and the password for connecting to it in the firewall configuration. In addition, you must register the firewall with the HSM server. Prior to beginning the configuration, make sure you have created a partition for the Palo Alto Networks devices on the HSM server.





HSM configuration is not synced between high availability firewall peers. Consequently, you must configure the HSM module separately on each of the peers.

In Active-Passive HA deployments, you must manually perform one failover to configure and authenticate each HA peer individually to the HSM. After this manual failover has been performed, user interaction is not required for the failover function.

Set up a Connectivity with a SafeNet Network HSM

<p>Step 1 Configure the firewall to communicate with the SafeNet Network HSM.</p>	<ol style="list-style-type: none"> 1. Log in to the firewall web interface and select Device > Setup > HSM. 2. Edit the Hardware Security Module Provider section and select Safenet Luna SA (SafeNet Network) as the Provider Configured. 3. Click Add and enter a Module Name. This can be any ASCII string up to 31 characters in length. 4. Enter the IPv4 address of the HSM module as the Server Address. If you are configuring a high availability HSM configuration, enter module names and IP addresses for the additional HSM devices. 5. (Optional) If configuring a high availability HSM configuration, select the High Availability check box and add the following: a value for Auto Recovery Retry and a High Availability Group Name. If two HSM servers are configured, you should configure high availability. Otherwise the second HSM server is not used. 6. Click OK and Commit.
--	--

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

<p>Step 2 (Optional) Configure a service route to enable the firewall to connect to the HSM.</p> <p>By default, the firewall uses the Management Interface to communicate with the HSM. To use a different interface, you must configure a service route.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Services. 2. Select Service Route Configuration from the Services Features area. 3. Select Customize from the Service Route Configuration area. 4. Select the IPv4 tab. 5. Select HSM from the Service column. 6. Select an interface to use for HSM from the Source Interface drop-down. <ul style="list-style-type: none">  If you select a dataplane connected port for HSM, issuing the <code>clear session all</code> CLI command will clear all existing HSM sessions, causing all HSM states to be brought down and then up. During the several seconds required for HSM to recover, all SSL/TLS operations will fail. 7. Click OK and Commit.
<p>Step 3 Configure the firewall to authenticate to the HSM.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > HSM. 2. Select Setup Hardware Security Module in the Hardware Security Operations area. 3. Select the HSM Server Name from the drop-down. 4. Enter the Administrator Password to authenticate the firewall to the HSM. 5. Click OK. The firewall attempts to perform an authentication with the HSM and displays a status message. 6. Click OK.
<p>Step 4 Register the firewall (the HSM client) with the HSM and assign it to a partition on the HSM.</p> <p> If the HSM already has a firewall with the same <code><cl-name></code> registered, you must remove the duplicate registration using the following command before registration will succeed:</p> <pre>client delete -client <cl-name></pre> <p>where <code><cl-name></code> is the name of the client (firewall) registration you want to delete.</p>	<ol style="list-style-type: none"> 1. Log in to the HSM from a remote system. 2. Register the firewall using the following command: <code>client register -c <cl-name> -ip <fw-ip-addr></code> where <code><cl-name></code> is a name that you assign to the firewall for use on the HSM and <code><fw-ip-addr></code> is the IP address of the firewall that is being configured as an HSM client. 3. Assign a partition to the firewall using the following command: <code>client assignpartition -c <cl-name> -p <partition-name></code> where <code><cl-name></code> is the name assigned to the firewall in the <code>client register</code> command and <code><partition-name></code> is the name of a previously configured partition that you want to assign to the firewall.

Set up a Connectivity with a SafeNet Network HSM (Continued)	
<p>Step 5 Configure the firewall to connect to the HSM partition.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > HSM. 2. Click the Refresh icon. 3. Select the Setup HSM Partition in the Hardware Security Operations area. 4. Enter the Partition Password to authenticate the firewall to the partition on the HSM. 5. Click OK.
<p>Step 6 (Optional) Configure an additional HSM for high availability (HA).</p>	<ol style="list-style-type: none"> 1. Follow Step 1 through Step 5 to add an additional HSM for high availability (HA). This process adds a new HSM to the existing HA group. 2. If you remove an HSM from your configuration, repeat Step 5. This will remove the deleted HSM from the HA group.
<p>Step 7 Verify connectivity with the HSM.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > HSM. 2. Check the Status of the HSM connection: Green—HSM is authenticated and connected. Red—HSM was not authenticated or network connectivity to the HSM is down. 3. View the following columns in Hardware Security Module Status area to determine authentication status: Serial Number—The serial number of the HSM partition if the HSM was successfully authenticated. Partition—The partition name on the HSM that was assigned on the firewall. Module State—The current operating state of the HSM. It always has the value Authenticated if the HSM is displayed in this table.


Set Up Connectivity with a Thales nShield Connect HSM

The following workflow describes how to configure the firewall to communicate with a Thales nShield Connect HSM. This configuration requires that you set up a remote filesystem (RFS) to use as a *hub* to sync key data for all firewalls in your organization that are using the HSM.



HSM configuration is not synced between high availability firewall peers. Consequently, you must configure the HSM module separately on each of the peers.

If the high availability firewall configuration is in Active-Passive mode, you must manually perform one failover to configure and authenticate each HA peer individually to the HSM. After this manual failover has been performed, user interaction is not required for the failover function.

Set up Connectivity with a Thales nShield Connect HSM	
<p>Step 1 Configure the Thales nShield Connect server as the firewall's HSM provider.</p>	<ol style="list-style-type: none"> 1. From the firewall web interface, select Device > Setup > HSM and edit the Hardware Security Module Provider section. 2. Select Thales Nshield Connect as the Provider Configured. 3. Click Add and enter a Module Name. This can be any ASCII string up to 31 characters in length. 4. Enter the IPv4 address as the Server Address of the HSM module. If you are configuring a high availability HSM configuration, enter module names and IP addresses for the additional HSM devices. 5. Enter the IPv4 address of the Remote Filesystem Address. 6. Click OK and Commit.
<p>Step 2 (Optional) Configure a service route to enable the firewall to connect to the HSM.</p> <p>By default, the firewall uses the Management Interface to communicate with the HSM. To use a different interface, you must configure a service route.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Services. 2. Select Service Route Configuration from the Services Features area. 3. Select Customize from the Service Route Configuration area. 4. Select the IPv4 tab. 5. Select HSM from the Service column. 6. Select an interface to use for HSM from the Source Interface drop-down.  If you select a dataplane connected port for HSM, issuing the <code>clear session</code> all CLI command will clear all existing HSM sessions, causing all HSM states to be brought down and then up. During the several seconds required for HSM to recover, all SSL/TLS operations will fail. 7. Click OK and Commit.
<p>Step 3 Register the firewall (the HSM client) with the HSM server.</p> <p>This step briefly describes the procedure for using the front panel interface of the Thales nShield Connect HSM. For more details, consult the Thales documentation.</p>	<ol style="list-style-type: none"> 1. Log in to the front panel display of the Thales nShield Connect HSM unit. 2. On the unit front panel, use the right-hand navigation button to select System > System configuration > Client config > New client. 3. Enter the IP address of the firewall. 4. Select System > System configuration > Client config > Remote file system and enter the IP address of the client computer where you set up the remote file system.

Set up Connectivity with a Thales nShield Connect HSM (Continued)

<p>Step 4 Set up the remote filesystem to accept connections from the firewall.</p>	<ol style="list-style-type: none"> Log in to the remote filesystem (RFS) from a Linux client. Obtain the electronic serial number (ESN) and the hash of the K_{NETI} key. The K_{NETI} key authenticates the module to clients: <code>anonkneti <ip-address></code> where <code><ip-address></code> is the IP address of the HSM. The following is an example: <code>anonkneti 192.0.2.1</code> <code>B1E2-2D4C-E6A2 5a2e5107e70d525615a903f6391ad72b1c03352c</code> In this example, B1E2-2D4C-E6A2 is the ESM and <code>5a2e5107e70d525615a903f6391ad72b1c03352c</code> is the hash of the K_{NETI} key. Use the following command from a superuser account to perform the remote filesystem setup: <code>rfs-setup --force <ip-address> <ESN> <hash-Kneti-key></code> where <code><ip-address></code> is the IP address of the HSM, <code><ESN></code> is the electronic serial number (ESN) and <code><hash-Kneti-key></code> is the hash of the K_{NETI} key. The following example uses the values obtained in this procedure: <code>rfs-setup --force <192.0.2.1> <B1E2-2D4C-E6A2></code> <code><5a2e5107e70d525615a903f6391ad72b1c03352c></code> Use the following command to permit client submit on the Remote Filesystem: <code>rfs-setup --gang-client --write-noauth <FW-IPaddress></code> where <code><FW-IPaddress></code> is the IP address of the firewall.
<p>Step 5 Configure the firewall to authenticate to the HSM.</p>	<ol style="list-style-type: none"> From the firewall web interface, select Device > Setup > HSM. Select Setup Hardware Security Module in the Hardware Security Operations area. Click OK. The firewall attempts to perform an authentication with the HSM and displays a status message. Click OK.
<p>Step 6 Synchronize the firewall with the remote filesystem.</p>	<ol style="list-style-type: none"> Select the Device > Setup > HSM. Select Synchronize with Remote Filesystem in the Hardware Security Operations section.

Set up Connectivity with a Thales nShield Connect HSM (Continued)

<p>Step 7 Verify that the firewall can connect to the HSM.</p>	<ol style="list-style-type: none">1. Select Device > Setup > HSM.2. Check the Status indicator to verify that the firewall is connected to the HSM: Green—HSM is authenticated and connected. Red—HSM was not authenticated or network connectivity to the HSM is down.3. View the following columns in Hardware Security Module Status section to determine authentication status. Name: The name of the HSM attempting to be authenticated. IP address: The IP address of the HSM that was assigned on the firewall. Module State: The current operating state of the HSM: Authenticated or Not Authenticated.
---	--

Encrypt a Master Key Using an HSM

A master key is configured on a Palo Alto Networks firewall to encrypt all private keys and passwords. If you have security requirements to store your private keys in a secure location, you can encrypt the master key using an encryption key that is stored on an HSM. The firewall then requests the HSM to decrypt the master key whenever it is required to decrypt a password or private key on the firewall. Typically, the HSM is located in a highly secure location that is separate from the firewall for greater security.

The HSM encrypts the master key using a wrapping key. To maintain security, this encryption key must occasionally be changed. For this reason, a command is provided on the firewall to rotate the wrapping key which changes the master key encryption. The frequency of this wrapping key rotation depends on your application.



Master key encryption using an HSM is not supported on firewalls configured in FIPS or CC mode.

The following topics describe how to encrypt the master key initially and how to refresh the master key encryption:

- ▲ [Encrypt the Master Key](#)
- ▲ [Refresh the Master Key Encryption](#)

Encrypt the Master Key

If you have not previously encrypted the master key on a device, use the following procedure to encrypt it. Use this procedure for first time encryption of a key, or if you define a new master key and you want to encrypt it. If you want to refresh the encryption on a previously encrypted key, see [Refresh the Master Key Encryption](#).

Encrypt a Master Key Using an HSM

Step 1 Select **Device > Master Key and Diagnostics**.

Step 2 Specify the key that is currently used to encrypt all of the private keys and passwords on the firewall in the **Master Key** field.

Step 3 If changing the master key, enter the new master key and confirm.

Step 4 Select the **HSM** check box.

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, refresh the master key encryption on a regular basis by rotating the master key wrapping key on the HSM. This command is the same for both the SafeNet Network and Thales nShield Connect HSMs.

Refresh the Master Key Encryption

1. Use the following CLI command to rotate the wrapping key for the master key on an HSM:

```
> request hsm mkey-wrapping-key-rotation
```

If the master key is encrypted on the HSM, the CLI command will generate a new wrapping key on the HSM and encrypt the master key with the new wrapping key.


If the master key is not encrypted on the HSM, the CLI command will generate new wrapping key on the HSM for future use.

The old wrapping key is not deleted by this command.

Store Private Keys on an HSM

For added security, you can use an HSM to secure the private keys used in SSL/TLS decryption for:

- **SSL Forward Proxy**—The HSM can store the private key of the Forward Trust certificate that signs certificates in SSL/TLS forward proxy operations. The firewall will then send the certificates that it generates during such operations to the HSM for signing before forwarding the certificates to the client.
- **SSL Inbound Inspection**—The private keys for the internal servers for which you are doing SSL/TLS inbound inspection can be stored on the HSM.

Store Private Keys on an HSM	
<p>Step 1 On the HSM, import or generate the certificate and private key used in your decryption deployment.</p>	For instructions on importing or generating a certificate and private key on the HSM, refer to your HSM documentation.
<p>Step 2 (Thales nShield Connect only) Synchronize the key data from the Thales nShield remote file system to the firewall.</p> <p> Synchronization with the SafeNet Network HSM is automatic.</p>	<ol style="list-style-type: none"> 1. From the firewall web interface, select Device > Setup > HSM. 2. Select Synchronize with Remote Filesystem in the Hardware Security Operations section.
<p>Step 3 Import the certificate that corresponds to the HSM-stored key onto the firewall.</p>	<ol style="list-style-type: none"> 1. From the firewall web interface, select Device > Certificate Management > Certificates > Device Certificates. 2. Click Import. 3. Enter the Certificate Name. 4. Browse to the Certificate File on the HSM. 5. Select a File Format. 6. Select Private Key resides on Hardware Security Module. 7. Click OK and Commit.
<p>Step 4 (Forward trust certificates only) Enable the certificate for use in SSL/TLS Forward Proxy.</p>	<ol style="list-style-type: none"> 1. Locate the certificate you imported in Step 3. 2. Select Forward Trust Certificate. 3. Click OK and Commit.
<p>Step 5 Verify that you successfully imported the certificate onto the firewall.</p>	<p>Locate the certificate you imported in Step 3 and check the icon in the Key column:</p> <ul style="list-style-type: none"> • Lock icon—The private key for the certificate is on the HSM. • Error icon—The private key is not on the HSM or the HSM is not properly authenticated or connected.

Manage the HSM Deployment

Manage HSM	
<ul style="list-style-type: none">View the HSM configuration settings.	Select Device > Setup > HSM .
<ul style="list-style-type: none">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.
<ul style="list-style-type: none">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.
<ul style="list-style-type: none">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 configuration in which two firewalls are placed in a group and their configuration is synchronized to prevent a single point to failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up the firewalls in HP pairs provides redundancy and allows you to ensure business continuity.

The Palo Alto Networks firewalls support stateful active/passive or active/active high availability with session and configuration synchronization. Some models of the firewall, such as the VM-Series firewall and the PA-200, only support [HA lite](#) without session synchronization capability. The following topics provide more information about high availability and how to configure it in your environment.

- ▲ [HA Overview](#)
- ▲ [HA Concepts](#)
- ▲ [Set Up Active/Passive HA](#)
- ▲ [HA Resources](#)

HA Overview

On Palo Alto Networks firewalls, you can set up two firewalls as an HA pair. HA allows you to minimize downtime by making sure that an alternate firewall is available in the event that the peer firewall fails. The firewalls in an HA pair use dedicated or in-band HA ports on the firewall to synchronize data—network, object, and policy configurations—and to maintain state information. Firewall-specific configuration such as management interface IP address or administrator profiles, HA specific configuration, log data, and the Application Command Center (ACC) information is not shared between peers. For a consolidated application and log view across the HA pair, you must use Panorama, the Palo Alto Networks centralized management system.

When a failure occurs on one firewall in the pair and the other firewall takes over the task of securing traffic, the event is called a *failover*. The conditions that trigger a failover are:

- One or more of the monitored interfaces fail. ([Link Monitoring](#))
- One or more of the destinations specified on the firewall cannot be reached. ([Path Monitoring](#))
- The firewall does not respond to heartbeat polls. ([Heartbeat Polling and Hello messages](#))

After you understand the [HA Concepts](#), continue to [Set Up Active/Passive HA](#).

HA Concepts

The following topics provide conceptual information about how HA works on a Palo Alto Networks firewall:

- ▲ [HA Modes](#)
- ▲ [HA Links and Backup Links](#)
- ▲ [Device Priority and Preemption](#)
- ▲ [Failover Triggers](#)
- ▲ [HA Timers](#)

HA Modes

You can set up the firewalls for HA in two modes:

- **Active/Passive**— One firewall actively manages traffic while the other is synchronized and ready to transition to the active state, should a failure occur. In this configuration, both firewalls share the same configuration settings, and one actively manages traffic until a path, link, system, or network failure occurs. When the active firewall fails, the passive firewall takes over seamlessly and enforces the same policies to maintain network security. Active/passive HA is supported in the virtual wire, Layer 2 and Layer 3 deployments. For information on setting up your firewall in an active/passive configuration, see [Configure Active/Passive HA](#).



The PA-200 and the VM-Series firewalls support a lite version of active/passive HA. HA lite provides configuration synchronization and some runtime data synchronization such as IPsec security associations. It does not support any session synchronization, and therefore, HA Lite does not offer stateful failover.

- **Active/Active**— Both the firewalls in the pair are active and processing traffic, and work synchronously to handle session setup and session ownership. The active/active deployment is supported in virtual wire and Layer 3 deployments, and is only recommended for networks with asymmetric routing. For information on setting up the firewalls in an active/active configuration, refer to the [Active/Active High Availability Tech Note](#).

HA Links and Backup Links

The firewalls in an HA pair use HA links to synchronize data and maintain state information. Some models of the firewall have dedicated HA ports—Control link (HA1) and Data link (HA2), while others require you to use the in-band ports as HA links.

On firewalls with dedicated HA ports such as the PA-3000 Series, PA-4000 Series, PA-5000 Series, and PA-7050 firewalls (see [HA Ports on the PA-7050 Firewall](#)), use the dedicated HA ports to manage communication and synchronization between the firewalls. For firewalls without dedicated HA ports such as the PA-200, PA-500, and PA-2000 Series firewalls, as a best practice use the management port for the HA1 link to allow for a direct connection between the management planes on the firewalls, and an in-band port for the HA2 link.



The HA1 and HA2 links provide synchronization for functions that reside on the management plane. Using the dedicated HA interfaces on the management plane is more efficient than using the in-band ports as this eliminates the need to pass the synchronization packets over the dataplane.

- **Control Link:** The HA1 link is used to exchange hellos, heartbeats, and HA state information, and management plane sync for routing, and User-ID information. This link is also used to synchronize configuration changes on either the active or passive firewall with its peer. The HA1 link is a Layer 3 link and requires an IP address.

Ports used for HA1: TCP port 28769 and 28260 for clear text communication; port 28 for encrypted communication (SSH over TCP).

- **Data Link:** The HA2 link is used to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between firewalls in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active firewall to the passive firewall. The HA2 link is a Layer 2 link, and it uses ether type 0x7261 by default.

Ports used for HA2: The HA data link can be configured to use either IP (protocol number 99) or UDP (port 29281) as the transport, and thereby allow the HA data link to span subnets.

Additionally, an HA3 link is used in Active/Active HA deployments. When there is an asymmetric route, the HA3 link is used for forwarding packets to the HA peer that owns the session. The HA3 link is a Layer 2 link and it does not support Layer 3 addressing or encryption.

- **Backup Links:** Provide redundancy for the HA1 and the HA2 links. In-band ports are used as backup links for both HA1 and HA2. Consider the following guidelines when configuring backup HA links:
 - The IP addresses of the primary and backup HA links must not overlap each other.
 - HA backup links must be on a different subnet from the primary HA links.
 - HA1-backup and HA2-backup ports must be configured on separate physical ports. The HA1-backup link uses port 28770 and 28260.



Palo Alto Networks recommends enabling heartbeat backup (uses port 28771 on the MGT interface) if you use an in-band port for the HA1 or the HA1 backup links.

- **Packet-Forwarding Link:** In addition to the HA1 and HA2 links, an active/active deployment also requires a dedicated HA3 link. The firewalls use this link for forwarding packets to the peer during session setup and asymmetric traffic flow. The HA3 link is a Layer 2 link that uses MAC-in-MAC encapsulation; it does not support Layer 3 addressing or encryption. You can configure aggregate interfaces on the PA-3000 Series, PA-4000 Series, PA-5000 Series and PA-7000 Series firewalls as an HA3 link. The aggregate interfaces can also provide redundancy for the HA3 link; you cannot configure backup links for the HA3 link.

HA Ports on the PA-7050 Firewall

For HA connectivity on the PA-7050, use the following table for details on which ports on the Switch Management Card (SMC) are mandated and where ports on the Network Processing Card (NPC) are suitable.

HA Links and Backup Links	Ports on the SMC	Description
Control Link	HA1-A Speed: Ethernet 10/100/1000	Used for HA control and synchronization. Connect this port directly from the HA1-A port on the first firewall to the HA1-A on the second firewall in the pair, or connect them together through a switch or router. HA1 cannot be configured on NPC data ports or the MGT port.
Control Link Backup	HA1-B Speed: Ethernet 10/100/1000 port	Used for HA control and synchronization as a backup for HA1-A. Connect this port directly from the HA1-B port on the first firewall to the HA1-B on the second firewall in the pair, or connect them together through a switch or router. HA1 Backup cannot be configured on NPC data ports or the MGT port.
Data Link Data Link Backup	HSCI-A HSCI-B	The High Speed Chassis Interconnect (HSCI) ports are layer-1 Quad Port SFP+ (QSFP+) interfaces which are used to connect two PA-7000 Series firewalls in an HA configuration. Each port is comprised of four 10 gigabit channels multiplexed for a combined speed of 40 gigabits. The traffic carried on the HSCI ports is raw layer-1, which is not routable or switchable; therefore the HSCI ports must be connected directly to each other. The HSCI-A on the first chassis connects directly to HSCI-A on the second chassis and HSCI-B on the first chassis connects to HSCI-B on the second chassis. This will provide full 80 gigabit transfer rates. In software, both ports (HSCI-A and HSCI-B) are treated as one HA interface. Palo Alto Networks recommends using the dedicated HSCI ports for the HA2 link. The HA3 link, required for packet forwarding in an active/active deployment, must use the HSCI port; the HA3 traffic cannot be configured on data ports. If the firewalls are deployed in: <ul style="list-style-type: none"> • an active/active configuration, the HA3 link may use only the HSCI ports. The HA2 link and HA2 backup links can use the HSCI ports or data ports on the NPC. • an active/passive configuration, you can configure a data port on the NPC for the HA2 link or the HA2 backup link, if needed.

For an overview of the Modules and Interface cards on the PA-7050 firewall, refer to the [PA-7050 Hardware Reference Guide](#).

Device Priority and Preemption

The firewalls in an HA pair can be assigned a *device priority* value to indicate a preference for which firewall should assume the active role and manage traffic. If you need to use a specific firewall in the HA pair for actively securing traffic, you must enable the preemptive behavior on both the firewalls and assign a device priority value for each firewall. The firewall with the lower numerical value, and therefore *higher priority*, is designated as active

and manages all traffic on the network. The other firewall is in a passive state, and synchronizes configuration and state information with the active firewall so that it is ready to transition to an active state should a failure occur.

By default, preemption is disabled on the firewalls and must be enabled on both firewalls. When enabled, the preemptive behavior allows the firewall with the *higher priority* (lower numerical value) to resume as active after it recovers from a failure. When preemption occurs, the event is logged in the system logs.

Failover Triggers

When a failure occurs on one firewall and the peer takes over the task of securing traffic, the event is called a *failover*. A failover is triggered when a monitored metric on a firewall in an HA pair fails. The metrics that are monitored for detecting a device failure are:

- **Heartbeat Polling and Hello messages**

The firewalls use hello message and heartbeats to verify that the peer firewall is responsive and operational. Hello messages are sent from one peer to the other at the configured *Hello Interval* to verify the state of the firewall. The heartbeat is an ICMP ping to the HA peer over the control link, and the peer responds to the ping to establish that the firewalls are connected and responsive. By default, the interval for the heartbeat is 1000 milliseconds. For details on the HA timers that trigger a failover, see [HA Timers](#).

- **Link Monitoring**

The physical interfaces to be monitored are grouped into a link group and their state (link up or link down) is monitored. A link group can contain one or more physical interfaces. A firewall failure is triggered when any or all of the interfaces in the group fail. The default behavior is failure of any one link in the link group will cause the firewall to change the HA state to non-functional to indicate a failure of a monitored object.

- **Path Monitoring**

Monitors the full path through the network to mission-critical IP addresses. ICMP pings are used to verify reachability of the IP address. The default interval for pings is 200ms. An IP address is considered unreachable when 10 consecutive pings (the default value) fail, and a firewall failure is triggered when any or all of the IP addresses monitored become unreachable. The default behavior is any one of the IP addresses becoming unreachable will cause the firewall to change the HA state to non-functional to indicate a failure of a monitored object.

In addition to the failover triggers listed above, a failover also occurs when the administrator places the firewall in a suspended state or if preemption occurs.

On the PA-3000 Series, PA-5000 Series, and PA-7050 firewalls, a failover can occur when an internal health check fails. This health check is not configurable and is enabled to verify the operational status for all the components within the firewall.

HA Timers

High Availability (HA) timers are used to detect a firewall failure and trigger a failover. To reduce the complexity in configuring HA timers, you can select from three profiles that have been added: **Recommended**, **Aggressive** and **Advanced**. These profiles auto-populate the optimum HA timer values for the specific firewall platform to enable a speedier HA deployment.

Use the **Recommended** profile for typical failover timer settings and the **Aggressive** profile for faster failover timer settings. The **Advanced** profile allows you to customize the timer values to suit your network requirements.

The following table describes each timer included in the profiles and the current preset values across the different hardware models; these values are for current reference only and can change in a subsequent release.

Recommended/Aggressive HA Timer Values by Platform

Timers	Description	PA-7050 PA-5000 Series PA-4000 Series PA-3000 Series	PA-2000 Series PA-500 Series PA-200 Series VM-Series	Panorama VM M-100
Monitor fail hold up time	Interval during which the firewall will remain active following a path monitor or link monitor failure. This setting is recommended to avoid an HA failover due to the occasional flapping of neighboring devices.	0/0	0/0	0/0
Preemption hold time	Time a passive or active-secondary firewall will wait before taking over as the active or active-primary firewall.	1/1	1/1	1/1
Heartbeat interval	Frequency at which the HA peers exchange heartbeat messages in the form of an ICMP ping.	1000/1000	2000/1000	2000/1000
Promotion hold time	Time that the passive firewall (in active/passive mode) or the active-secondary firewall (in active/active mode) will wait before taking over as the active or active-primary firewall after communications with the HA peer have been lost. This hold time will begin only after the peer failure declaration has been made.	2000/500	2000/500	2000/500

Timers	Description	PA-7050 PA-5000 Series PA-4000 Series PA-3000 Series	PA-2000 Series PA-500 Series PA-200 Series VM-Series	Panorama VM M-100
Additional master hold up time	This time interval is applied to the same event as Monitor Fail Hold Up Time (range 0-60000 ms, default 500 ms). The additional time interval is applied only to the active firewall in active/passive mode and to the active-primary firewall in active/active mode. This timer is recommended to avoid a failover when both firewalls experience the same link/path monitor failure simultaneously.	500/500	500/500	7000/5000
Hello interval	Interval in milliseconds between the hello packets that are sent to verify that the HA functionality on the other firewall is operational. The range is 8000-60000 ms with a default of 8000 ms for all platforms.	8000/8000	8000/8000	8000/8000
Maximum no. of flaps	A flap is counted when the firewall leaves the active state within 15 minutes after it last left the active state. This value indicates the maximum number of flaps that are permitted before the firewall is determined to be suspended and the passive firewall takes over (range 0-16, default 3).	3/3	3/3	Not Applicable

Set Up Active/Passive HA

- ▲ Prerequisites for Active/Passive HA
- ▲ Configuration Guidelines for Active/Passive HA
- ▲ Configure Active/Passive HA
- ▲ Define HA Failover Conditions
- ▲ Verify Failover

Prerequisites for Active/Passive HA

To set up high availability on your Palo Alto Networks firewalls, you need a pair of firewalls that meet the following requirements:

- ❑ **The same model**—Both the firewalls in the pair must be of the same hardware model or virtual machine model.
- ❑ **The same PAN-OS version**—Both the firewalls must be running the same PAN-OS version and must each be up-to-date on the application, URL, and threat databases. They must also both have the same multiple virtual systems capability (single or multi vsys).
- ❑ **The same type of interfaces**—Dedicated HA links, or a combination of the management port and in-band ports that are set to *interface type* HA.
 - Determine the IP address for the HA1 (control) connection between the HA peers. The HA1 IP address for both peers must be on the same subnet if they are directly connected or are connected to the same switch.

For firewall without dedicated HA ports, you can use the management port for the control connection. Using the management port provides a direct communication link between the management planes on both firewalls. However, because the management ports will not be directly cabled between the firewalls, make sure that you have a route that connects these two interfaces across your network.
 - If you use Layer 3 as the transport method for the HA2 (data) connection, determine the IP address for the HA2 link. Use *Layer 3 only* if the HA2 connection must communicate over a routed network. The IP subnet for the HA2 links must not overlap with that of the HA1 links or with any other subnet assigned to the data ports on the firewall.
- ❑ **The same set of licenses**—Licenses are unique to each firewall and cannot be shared between the firewalls. Therefore, you must license both firewalls identically. If both firewalls do not have an identical set of licenses, they cannot synchronize configuration information and maintain parity for a seamless failover.



If you have an existing firewall and you want to add a new firewall for HA purposes and the new firewall has an existing configuration, it is recommended that you [Reset the Firewall to Factory Default Settings](#) on the new firewall. This will ensure that the new firewall has a clean configuration. After HA is configured, you will then sync the configuration on the primary firewall to the newly introduced firewall with the clean config.

Configuration Guidelines for Active/Passive HA

To set up an active (PeerA) passive (PeerB) pair in HA, you must configure some options identically on both firewalls and some independently (non-matching) on each firewall. These HA settings are not synchronized between the firewalls. For details on what is/is not synchronized, refer to [HA Synchronization](#).

To proceed with the instructions on configuring the firewalls in HA, see [Configure Active/Passive HA](#).

The following table lists the settings that you must configure identically on both firewalls:

Identical Configuration Settings on PeerA and PeerB

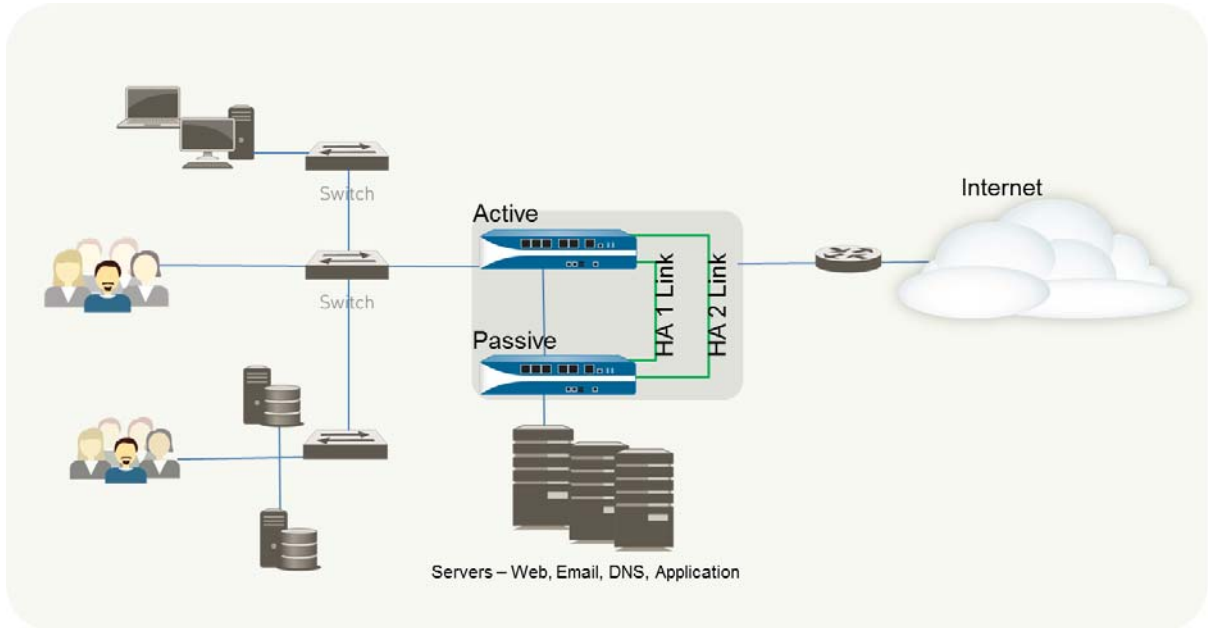
- HA must be enabled on both firewalls.
- Both firewall must have the same Group ID value. The Group ID value is used to create a virtual MAC address for all the configured interfaces. The format of the virtual MAC is 00-1B-17:00: xx: yy where
00-1B-17: vendor ID; 00: fixed; xx: HA group ID; yy: interface ID.
When a new active firewall takes over, Gratuitous ARPs are sent from each of the connected interfaces of the new active member to inform the connected Layer 2 switches of the virtual MAC address' new location.
- If using in-band ports, the interfaces for the HA1 and HA2 links must be set to type HA.
- The HA mode must be set to Active Passive.
- If required, preemption must be enabled on both firewalls. The device priority value, however, must not be identical.
- If required, encryption on the HA1 link (for communication between the HA peers) must be configured 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:
 - 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 settings that must be configured independently on each firewall:

Independent Configuration Settings	PeerA	PeerB
Control Link	<p>IP address of the HA1 link configured on this firewall (PeerA).</p> <p>For firewalls without dedicated HA ports, use the management port IP address for the control link.</p>	<p>IP address of the HA1 link configured on this firewall (PeerB).</p>
<p>Data Link</p> <p>The data link information is synchronized between the firewalls after HA is enabled and the control link is established between the firewalls.</p>	<p>By default, the HA2 link uses Ethernet/Layer 2.</p> <p>If using a Layer 3 connection, configure the IP address for the data link on this firewall (PeerA).</p>	<p>By default, the HA2 link uses Ethernet/Layer 2.</p> <p>If using a Layer 3 connection, configure the IP address for the data link on this firewall (PeerB).</p>
<p>Device Priority (required, if preemption is enabled)</p>	<p>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.</p>	<p>If PeerB is passive, set the device priority value to a number larger than that on PeerA. For example, set the value to 110.</p>
<p>Link Monitoring— Monitor one or more physical interfaces that handle vital traffic on this firewall and define the failure condition.</p>	<p>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.</p>	<p>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.</p>
<p>Path Monitoring— Monitor one or more destination IP addresses that the firewall can use ICMP pings to ascertain responsiveness.</p>	<p>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.</p> <p>Make sure that the node/device that you are monitoring is not likely to be unresponsive, especially when it comes under load, as this could cause a path monitoring failure and trigger a failover.</p>	<p>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.</p>

Configure Active/Passive HA

The following procedure shows how to configure a pair of firewalls in an active/passive deployment as depicted in the following example topology.



Connect and Configure the Firewalls

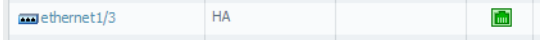
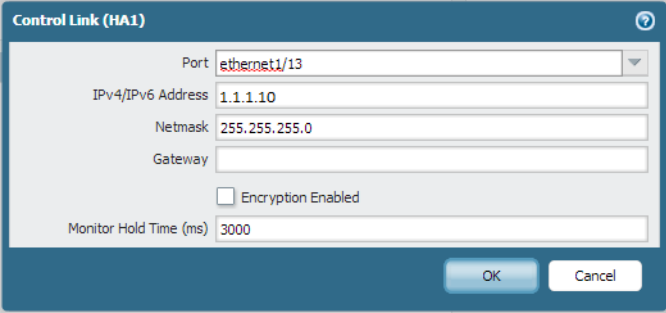
Step 1 Connect the HA ports to set up a physical connection between the firewalls.

- For firewalls with dedicated HA ports, use an Ethernet cable to connect the dedicated HA1 ports and the HA2 ports on the HA firewalls. Use a crossover cable if the firewalls are directly connected to each other.
 - 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.

Pick a firewall in the pair and complete the following steps:

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.

Connect and Configure the Firewalls (Continued)	
<p>Step 3 If the firewall does not have dedicated HA ports, set up the data ports to function as HA ports.</p> <p>For firewalls with dedicated HA ports continue to the next step.</p>	<ol style="list-style-type: none"> 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 the interface type to HA.  4. Set the Link Speed and Link Duplex settings, as appropriate.
<p>Step 4 Set the HA mode and group ID.</p>	<ol style="list-style-type: none"> 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.
<p>Step 5 Set up the control link connection.</p> <p>This example shows an in-band port that is set to interface type HA.</p> <p>For firewalls that use the management port as the control link, the IP address information is automatically pre-populated.</p>	<ol style="list-style-type: none"> 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. 4. If the HA1 interfaces are on separate subnets, enter the IP address of the Gateway. Do not add a gateway address if the firewalls are directly connected. 
<p>Step 6 (Optional) Enable encryption for the control link connection.</p> <p>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.</p>	<ol style="list-style-type: none"> 1. Export the HA key from a firewall and import it into the peer firewall. <ol style="list-style-type: none"> a. Select Device > Certificate Management > Certificates. b. Select Export HA key. Save the HA key to a network location that the peer can access. c. On the peer, select Device > Certificate Management > Certificates, and select Import HA key to browse to the location that you saved the key and import it in to the peer. 2. Select Device > High Availability > General, edit the Control Link (HA1) section. 3. Select Encryption Enabled.

Connect and Configure the Firewalls (Continued)

Step 7 Set up the backup control link connection.

1. In **Device > High Availability > General**, edit the Control Link (HA1 Backup) section.
2. Select the HA1 backup interface and set the **IPv4/IPv6 Address** and **Netmask**.


The screenshot shows the 'Control Link (HA1 Backup)' configuration window. It has a title bar with a question mark icon. The main area contains four input fields: 'Port' (dropdown menu showing 'ethernet1/22'), 'IPv4/IPv6 Address' (text box with '2.2.2.2'), 'Netmask' (text box with '255.255.0.0'), and 'Gateway' (empty text box). At the bottom right, there are 'OK' and 'Cancel' buttons.




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**.

The screenshot shows the 'Data Link (HA2)' configuration window. It has a title bar with a question mark icon. The main area contains several options: a checked checkbox for 'Enable Session Synchronization', a 'Port' dropdown menu showing 'ethernet1/14', 'IPv4/IPv6 Address', 'Netmask', and 'Gateway' text boxes, a 'Transport' dropdown menu showing 'ethernet', and an 'HA2 Keep-alive' section with a checked checkbox. Under 'HA2 Keep-alive', there are radio buttons for 'Log Only' (selected) and 'Split Datapath', and a 'Threshold (ms)' text box with '10000'.

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 you only enable the option on one peer, only that firewall will send the keep-alive messages. The other firewall will be notified if a failure occurs.
7. Edit the Data Link (HA2 Backup) section, select the interface, and add the **IPv4/IPv6 Address** and **Netmask**.

Connect and Configure the Firewalls (Continued)	
<p>Step 9 Enable heartbeat backup if your control link uses a dedicated HA port or an in-band port.</p> <p>You do not need to enable heartbeat backup if you are using the management port for the control link.</p>	<ol style="list-style-type: none"> In Device > High Availability > General, edit the Election Settings. 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.
<p>Step 10 Set the device priority and enable preemption.</p> <p>This setting is only required if you want to make sure that a specific firewall is the preferred active firewall. For information, see Device Priority and Preemption.</p>	<ol style="list-style-type: none"> In Device > High Availability > General, edit the Election Settings. 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. Select Preemptive. You must enable preemptive on both the active firewall and the passive firewall.
<p>Step 11 (Optional) Modify the failover timers.</p> <p>By default, the HA timer profile is set to the Recommended profile and is suited for most HA deployments.</p>	<ol style="list-style-type: none"> In Device > High Availability > General, edit the Election Settings. 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.

Connect and Configure the Firewalls (Continued)

Step 12 (Optional, only configured on the passive firewall) Modify the link status of the HA ports on the passive firewall.



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**.



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.

The auto option decreases the amount of time it takes for the passive firewall to take over when a failover occurs.

Step 13 Enable HA.

1. In **Device > High Availability > General**, 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 and should always be enabled.
4. Enter the IP address assigned to the control link of the peer firewall 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 Save your configuration changes.

Click **Commit**.

Step 15 Complete [Step 2](#) through [Step 14](#) on the other firewall in the HA pair.

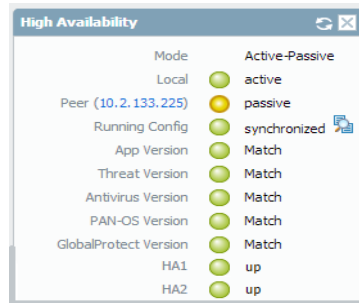
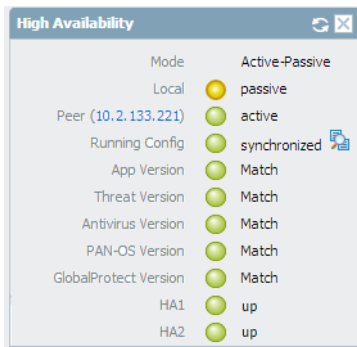
Connect and Configure the Firewalls (Continued)

Step 16 After you finish configuring both firewalls, verify that the firewalls are paired in active/passive HA.


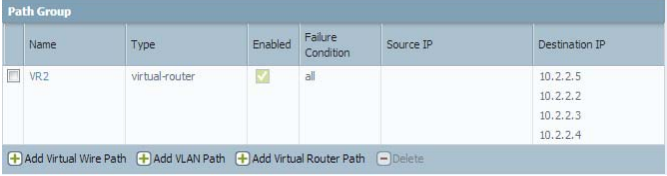
1. Access the **Dashboard** on both firewalls, and view the High Availability widget.
2. On the active firewall, click the **Sync to peer** link.
3. Confirm that the firewalls are paired and synced, as shown below:

On the passive firewall: the state of the local firewall should display **passive** and the Running Config should show as **synchronized**.

On the active firewall: The state of the local firewall should display **active** and the Running Config should show as **synchronized**.



Define HA Failover Conditions

Configure the Failover Triggers																				
<p>Step 1 To configure link monitoring, define the interfaces that you would like to monitor. A change in the link state of these interface will trigger a failover.</p>	<ol style="list-style-type: none"> 1. Select Device > High Availability > Link and Path Monitoring. 2. In the Link Group section, click Add. 3. Name the Link Group, Add the interfaces to monitor, and select the Failure Condition for the group. The Link group you define is added to the Link Group section. 																			
 <table border="1"> <thead> <tr> <th colspan="4">Link Group</th> </tr> <tr> <th>Name</th> <th>Enabled</th> <th>Group Failure Condition</th> <th>Interfaces</th> </tr> </thead> <tbody> <tr> <td>Main</td> <td><input checked="" type="checkbox"/></td> <td>any</td> <td>ethernet1/11 ethernet1/15 ethernet1/21</td> </tr> </tbody> </table>	Link Group				Name	Enabled	Group Failure Condition	Interfaces	Main	<input checked="" type="checkbox"/>	any	ethernet1/11 ethernet1/15 ethernet1/21	<p>Step 2 (Optional) Modify the failure condition for the Link Groups that you configured (in the preceding step) on the firewall.</p> <p>By default, the firewall will trigger a failover when any monitored link fails.</p>	<ol style="list-style-type: none"> 1. Select the Link Monitoring section. 2. Set the Failure Condition to All. The default setting is Any. 						
Link Group																				
Name	Enabled	Group Failure Condition	Interfaces																	
Main	<input checked="" type="checkbox"/>	any	ethernet1/11 ethernet1/15 ethernet1/21																	
<p>Step 3 To configure path monitoring, define the destination IP addresses that the firewall should ping to verify network connectivity.</p>	<ol style="list-style-type: none"> 1. In the Path Group section of the Device > High Availability > Link and Path Monitoring tab, pick the Add option for your set up: Virtual Wire, VLAN, or Virtual Router. 2. Select the appropriate item from the drop-down for the Name and Add the IP addresses (source and/or destination, as prompted) that you wish to monitor. Then select the Failure Condition for the group. The path group you define is added to the Path Group section. 																			
 <table border="1"> <thead> <tr> <th colspan="6">Path Group</th> </tr> <tr> <th>Name</th> <th>Type</th> <th>Enabled</th> <th>Failure Condition</th> <th>Source IP</th> <th>Destination IP</th> </tr> </thead> <tbody> <tr> <td>VR2</td> <td>virtual-router</td> <td><input checked="" type="checkbox"/></td> <td>all</td> <td></td> <td>10.2.2.5 10.2.2.2 10.2.2.3 10.2.2.4</td> </tr> </tbody> </table>	Path Group						Name	Type	Enabled	Failure Condition	Source IP	Destination IP	VR2	virtual-router	<input checked="" type="checkbox"/>	all		10.2.2.5 10.2.2.2 10.2.2.3 10.2.2.4	<p>Step 4 (Optional) Modify the failure condition for all Path Groups configured on the firewall.</p> <p>By default, the firewall will trigger a failover when any monitored path fails.</p>	<p>Set the Failure Condition to All. The default setting is Any.</p>
Path Group																				
Name	Type	Enabled	Failure Condition	Source IP	Destination IP															
VR2	virtual-router	<input checked="" type="checkbox"/>	all		10.2.2.5 10.2.2.2 10.2.2.3 10.2.2.4															
<p>Step 5 Save your changes.</p>	<p>Click Commit.</p>																			

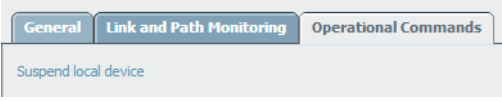
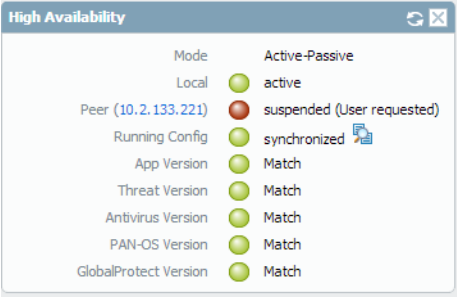
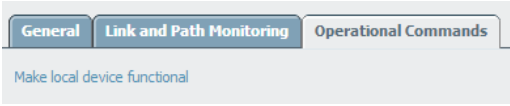
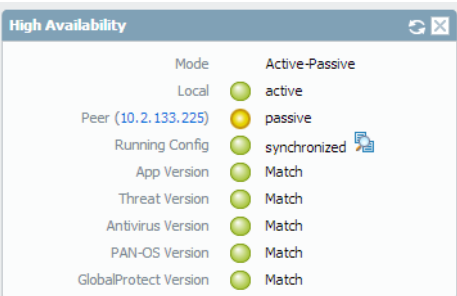


If you are using SNMPv3 to monitor the firewalls, note that the SNMPv3 Engine ID is unique to each firewall; the EngineID is not synchronized between the HA pair and, therefore, allows you to independently monitor each firewall in the HA pair. For information on setting up SNMP, see [Set Up SNMP Trap Destinations](#).

Because the EngineID is generated using the firewall's unique serial number, on the VM-Series firewall you must apply a valid license in order to obtain a unique EngineID for each firewall.

Verify Failover

To test that your HA configuration works properly, trigger a manual failover and verify that the firewalls transition states successfully.

Verify Failover	
<p>Step 1 Suspend the active firewall.</p>	<p>Click the Suspend local device link on the Device > High Availability > Operational Commands tab.</p> 
<p>Step 2 Verify that the passive firewall has taken over as active.</p>	<p>On the Dashboard, verify that the state of the passive firewall changes to active in the High Availability widget.</p> 
<p>Step 3 Restore the suspended firewall to a functional state. Wait for a couple minutes, and then verify that preemption has occurred, if preemptive is enabled.</p>	<ol style="list-style-type: none"> On the firewall you previously suspended, select the Make local device functional link on the Device > High Availability > Operational Commands tab.  <ol style="list-style-type: none"> 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. 

HA Resources

For more information on HA, refer to the following sources:

- [Active/Active HA](#)
- [High Availability Synchronization](#)
- [High Availability Failover Optimization](#)
- [Upgrading an HA pair](#)
- [Examples: Deploying HA](#)





Reports and Logging

The firewall provides reports and logs that are useful for monitoring activity on your network. You can monitor the logs and filter the information to generate reports with predefined or customized views. You can, for example, use the predefined templates to generate reports on a user's activity or analyze the reports and logs to interpret unusual behavior on your network and generate a custom report on the traffic pattern. The following topics describe how to view, manage, customize, and generate the reports and logs on the firewall:


- ▲ [Use the Dashboard](#)
- ▲ [Use the Application Command Center](#)
- ▲ [Use App-Scope](#)
- ▲ [Take Packet Captures](#)
- ▲ [Monitor the Firewall](#)
- ▲ [Forward Logs to External Services](#)
- ▲ [Monitor the Firewall Using SNMP](#)
- ▲ [Monitor the Firewall Using NetFlow](#)
- ▲ [NetFlow Templates](#)
- ▲ [Identify Firewall Interfaces in External Monitoring Systems](#)
- ▲ [Manage Reporting](#)
- ▲ [Syslog Field Descriptions](#)

Use the Dashboard

The **Dashboard** tab widgets show general device information, such as the software version, the operational status of each interface, resource utilization, and up to 10 of the most recent entries in the threat, configuration, and system logs. All of the available widgets are displayed by default, but each administrator can remove and add individual widgets, as needed.

Click the refresh icon  to update the Dashboard or an individual widget. To change the automatic refresh interval, select an interval from the drop-down (**1 min**, **2 mins**, **5 mins**, or **Manual**). To add a widget to the Dashboard, click the widget drop-down, select a category and then the widget name. To delete a widget, click  in the title bar.

The following table describes the Dashboard widgets.

Dashboard Charts	Descriptions
Top Applications	Displays the applications with the most sessions. The block size indicates the relative number of sessions (mouse-over the block to view the number), and the color indicates the security risk—from green (lowest) to red (highest). Click an application to view its application profile.
Top High Risk Applications	Similar to Top Applications, except that it displays the highest-risk applications with the most sessions.
General Information	Displays the device name, model, PAN-OS software version, the application, threat, and URL filtering definition versions, the current date and time, and the length of time since the last restart.
Interface Status	Indicates whether each interface is up (green), down (red), or in an unknown state (gray).
Threat Logs	Displays the threat ID, application, and date and time for the last 10 entries in the Threat log. The threat ID is a malware description or URL that violates the URL filtering profile.
Config Logs	Displays the administrator username, client (Web or CLI), and date and time for the last 10 entries in the Configuration log.
Data Filtering Logs	Displays the description and date and time for the last 60 minutes in the Data Filtering log.
URL Filtering Logs	Displays the description and date and time for the last 60 minutes in the URL Filtering log.
System Logs	Displays the description and date and time for the last 10 entries in the System log.  A <code>Config installed</code> entry indicates configuration changes were committed successfully.
System Resources	Displays the Management CPU usage, Data Plane usage, and the Session Count, which displays the number of sessions established through the firewall.
Logged In Admins	Displays the source IP address, session type (Web or CLI), and session start time for each administrator who is currently logged in.

Dashboard Charts	Descriptions
ACC Risk Factor	Displays the average risk factor (1 to 5) for the network traffic processed over the past week. Higher values indicate higher risk.
High Availability	If high availability (HA) is enabled, indicates the HA status of the local and peer device—green (active), yellow (passive), or black (other). For more information about HA, see High Availability .
Locks	Shows configuration locks taken by administrators.

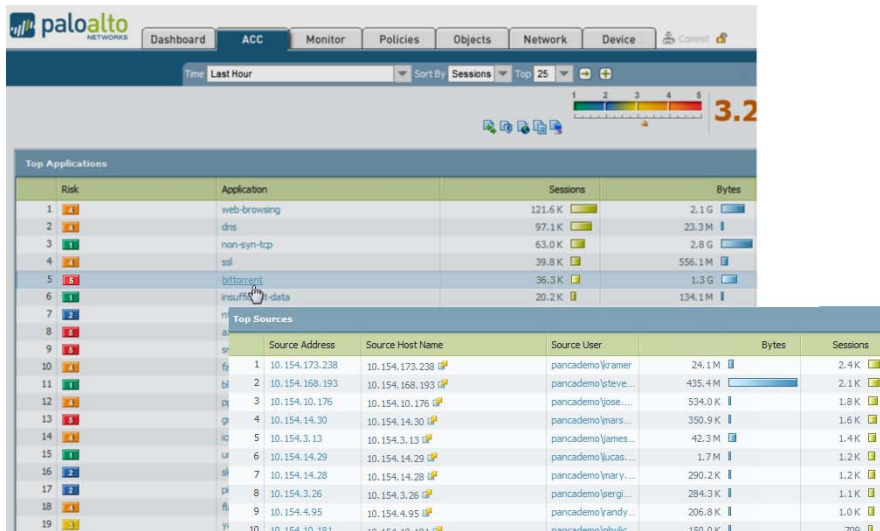
Use the Application Command Center

The **ACC** tab visually depicts trends and historic view of traffic on your network.

- ▲ [ACC Risk Level](#)
- ▲ [ACC Charts](#)
- ▲ [ACC Detail Pages](#)
- ▲ [Use the ACC](#)

ACC Risk Level

The **ACC** tab displays the overall risk level for all network traffic, the risk levels and number of threats detected for the most active and highest-risk applications on your network, and the number of threats detected from the busiest application categories and from all applications at each risk level. Use the ACC to view application data for the past hour, day, week, month, or any custom-defined time frame. **Risk** levels (1=lowest to 5=highest) indicate the application's relative security risk based on criteria such as whether the application can share files, is prone to misuse, or tries to evade firewalls.



ACC Charts

There are five charts displayed on the **Application Command Center (ACC)** tab:

- ▲ [Application](#)
- ▲ [URL Filtering](#)
- ▲ [Threat Prevention](#)
- ▲ [Data Filtering](#)
- ▲ [HIP Matches](#)

ACC Chart	Description
Application	<p>Displays application information grouped by the following attributes:</p> <ul style="list-style-type: none"> • Applications • High risk applications • Categories • Sub Categories • Technology • Risk <p>Each chart can include the number of sessions, bytes transmitted and received, number of threats, application category, application subcategories, application technology, and risk level, as applicable.</p>
URL Filtering	<p>Displays URL/category information grouped by the following attributes:</p> <ul style="list-style-type: none"> • URL Categories • URLs • Blocked URL Categories • Blocked URLs <p>Each chart can include the URL, URL category, repeat count (number of times access was attempted, as applicable).</p>
Threat Prevention	<p>Displays threat information grouped by the following attributes:</p> <ul style="list-style-type: none"> • Threats • Types • Spyware • Spyware Phone Home • Spyware Downloads • Vulnerability • Virus <p>Each chart can include the threat ID, count (number of occurrences), number of sessions, and subtype (such as vulnerability), as applicable.</p>

ACC Chart	Description
Data Filtering	Displays information on data filtered by the firewall grouped by the following attributes: <ul style="list-style-type: none">• Content/File Types• Types• File Names
HIP Matches	Displays the host information collected by the firewall grouped by: <ul style="list-style-type: none">• HIP Objects• HIP Profiles

ACC Detail Pages

To view additional details, click any of the links on the ACC charts. A details page opens to show information about the item at the top and additional lists for related items. For example, click on the web-browsing link on the Application chart opens the Application Information page for web-browsing:

Application Information

<p>Name: web-browsing</p> <p>Description: Web Browsing is using Hypertext Transfer Protocol (HTTP), which is a method used to transfer or convey information on the World Wide Web. Its original purpose was to provide a way to publish and retrieve HTML pages.</p> <p>Standard Ports: tcp/80</p> <p>Capable of File Transfer: yes</p> <p>Used by Malware: yes</p> <p>Excessive Bandwidth Use: no</p> <p>Evasive: no</p> <p>Tunnels Other Applications: yes</p> <p>Additional Information: Wikipedia Google Yahoo!</p>	<p>Category: general-internet</p> <p>Subcategory: internet-utility</p> <p>Technology: browser-based</p> <p>Risk: 4</p> <p>Widely Used: yes</p> <p>Has Known Vulnerabilities: yes</p> <p>Prone to Misuse: no</p> <p>Session Timeout (seconds):</p> <p>TCP Timeout (seconds):</p> <p>UDP Timeout (seconds):</p>
---	---

Top Applications

Risk	Application	Sessions	Bytes
1 4	web-browsing	559 <div style="width: 100%; height: 10px; background: linear-gradient(to right, #f4a460, #d9ead3);"></div>	11.1 M <div style="width: 100%; height: 10px; background: linear-gradient(to right, #5bc0de, #d9ead3);"></div>

Top Sources

	Source address	Source Host Name	Source User	Bytes	Sessions
1	10.16.0.54	10.16.0.54		9.9 M <div style="width: 100%; height: 10px; background: linear-gradient(to right, #5bc0de, #d9ead3);"></div>	491 <div style="width: 100%; height: 10px; background: linear-gradient(to right, #f4a460, #d9ead3);"></div>
2	10.16.1.1	10.16.1.1		829.3 K <div style="width: 100%; height: 10px; background: linear-gradient(to right, #5bc0de, #d9ead3);"></div>	45 <div style="width: 100%; height: 10px; background: linear-gradient(to right, #f4a460, #d9ead3);"></div>
3	10.16.0.33	10.16.0.33		220.2 K <div style="width: 100%; height: 10px; background: linear-gradient(to right, #5bc0de, #d9ead3);"></div>	13 <div style="width: 100%; height: 10px; background: linear-gradient(to right, #f4a460, #d9ead3);"></div>
4	10.16.0.34	10.16.0.34		25.0 K <div style="width: 100%; height: 10px; background: linear-gradient(to right, #5bc0de, #d9ead3);"></div>	7 <div style="width: 100%; height: 10px; background: linear-gradient(to right, #f4a460, #d9ead3);"></div>
5	10.16.0.34	10.16.0.34	paloaltonetwork\jpara...	121.4 K <div style="width: 100%; height: 10px; background: linear-gradient(to right, #5bc0de, #d9ead3);"></div>	3 <div style="width: 100%; height: 10px; background: linear-gradient(to right, #f4a460, #d9ead3);"></div>

Use the ACC

The following procedure describes how to use the **ACC** tab and how to customize your view:

Use the ACC

- Step 1** On the **ACC**, change one or more of the settings at the top of the page.
- Use the drop-down to select Applications, URL Categories, Threats, Content/File Types, and HIP Objects to view.
 - Select a virtual system, if virtual systems are defined.
 - Select a time period from the **Time** drop-down. The default is Last Hour.
 - Select a sorting method from the **Sort By** drop-down. You can sort the charts in descending order by number of sessions, bytes, or threats. The default is by number of sessions.
 - For the selected sorting method, select the top number of applications and application categories shown in each chart from the **Top** drop-down. Click the Submit icon to apply the selected settings.

- Step 2** To open log pages associated with the information on the page, use the log links in the upper-right corner of the page, as shown here. The context for the logs matches the information on the page.



- Step 3** To filter the list, click an item in one of the columns, this will add that item to the filter bar located above the log column names. After adding the desired filters, click the Apply Filter icon.

	Receive Time	Type	From Zone	To Zone	Source	Source User	Destinati
	09/05 16:44:14	end	I3-vlan-trust	I3-untrust	192.168.2.10		10.0.
	09/05 16:44:14	end	I3-vlan-trust	I3-untrust	192.168.2.10		10.0.

Use App-Scope

The App-Scope reports provide visibility and analysis tools to help pinpoint problematic behavior, helping you understand changes in application usage and user activity, users and applications that take up most of the network bandwidth, and identify network threats.

With the App-Scope reports, you can quickly see if any behavior is unusual or unexpected. Each report provides a dynamic, user-customizable window into the network; hovering the mouse over and clicking either the lines or bars on the charts opens detailed information about the specific application, application category, user, or source on the ACC. The App-Scope charts give you the ability to:

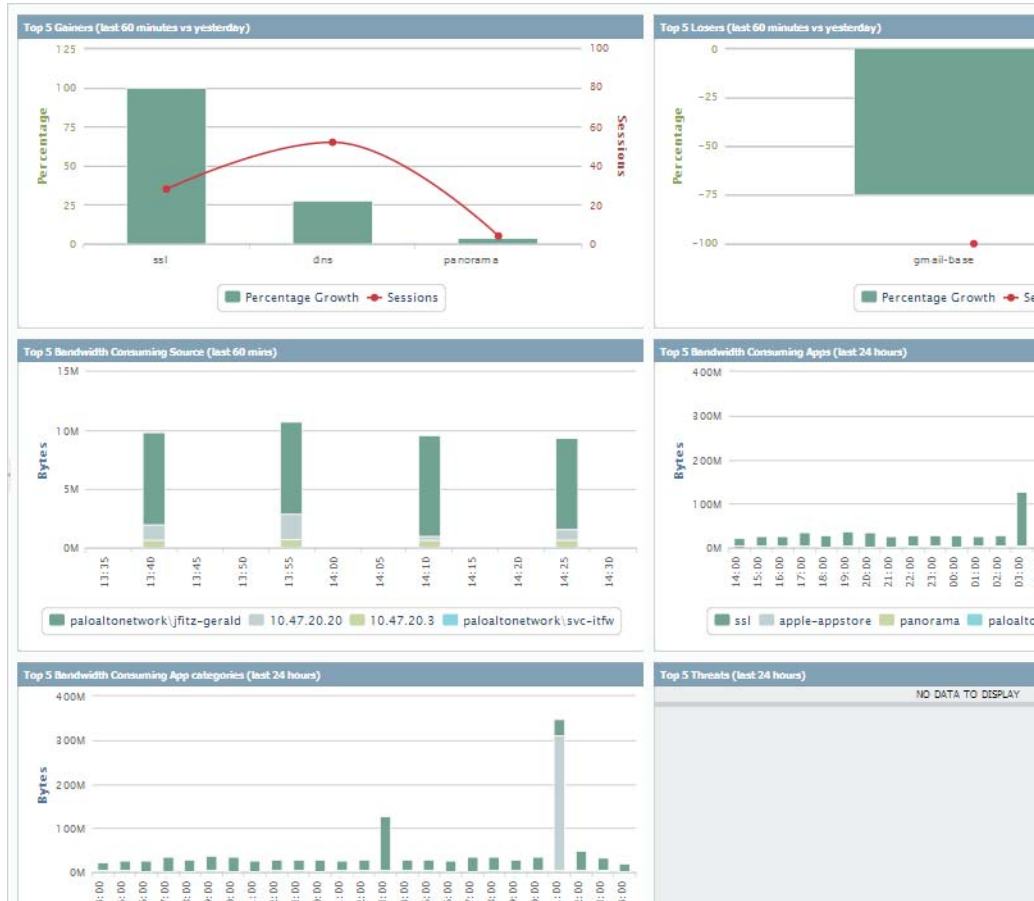
- Toggle the attributes in the legend to only view chart details that you want to review. The ability to include or exclude data from the chart allows you to change the scale and review details more closely.
- Click into an attribute in a bar chart and drill down to the related sessions in the ACC. Click into an Application name, Application Category, Threat Name, Threat Category, Source IP address or Destination IP address on any bar chart to filter on the attribute and view the related sessions in the ACC.
- Export a chart or map to PDF or as an image. For portability and offline viewing, you can Export charts and maps as PDFs or PNG images.

The following App-Scope reports are available:

- ▲ [Summary Report](#)
- ▲ [Change Monitor Report](#)
- ▲ [Threat Monitor Report](#)
- ▲ [Threat Map Report](#)
- ▲ [Network Monitor Report](#)
- ▲ [Traffic Map Report](#)

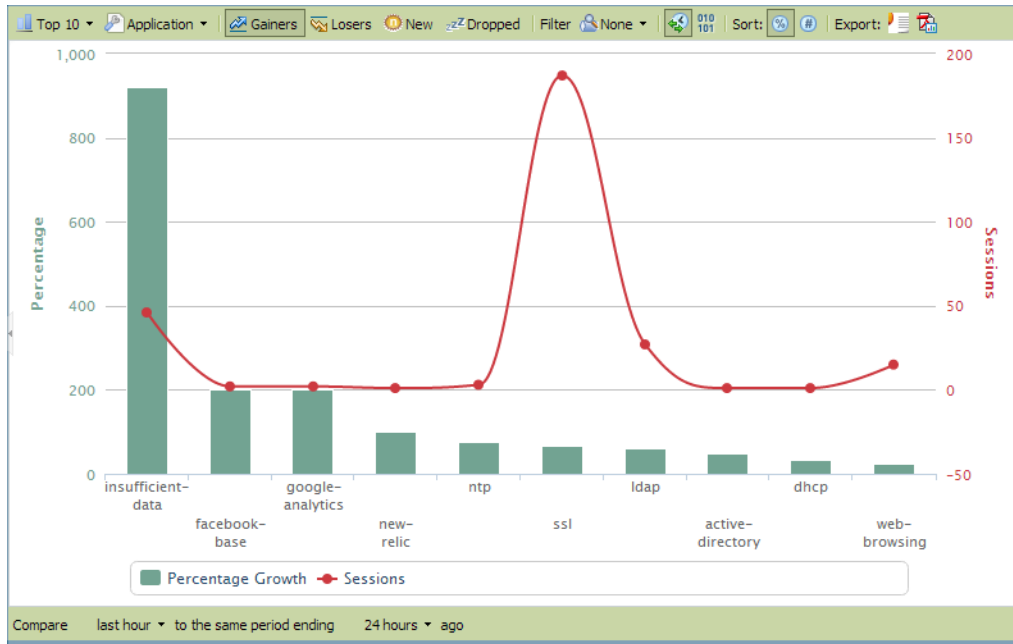
Summary Report

The App-Scope Summary report displays charts for the top five gainers, losers, and bandwidth consuming applications, application categories, users, and sources.



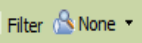
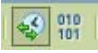
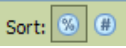
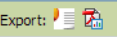
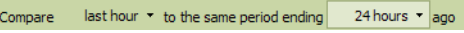
Change Monitor Report

The App-Scope Change Monitor report displays changes over a specified time period. For example, the following chart displays the top applications that gained in use over the last hour as compared with the last 24-hour period. The top applications are determined by session count and sorted by percent.




The Change Monitor Report contains the following buttons and options.

Button	Description
Top 10 ▾	Determines the number of records with the highest measurement included in the chart.
Application ▾	Determines the type of item reported: Application, Application Category, Source, or Destination.
Gainers	Displays measurements of items that have increased over the measured period.
Losers	Displays measurements of items that have decreased over the measured period.
New	Displays measurements of items that were added over the measured period.
Dropped	Displays measurements of items that were discontinued over the measured period.

Button	Description
	Applies a filter to display only the selected item. None displays all entries.
	Determines whether to display session or byte information.
	Determines whether to sort entries by percentage or raw growth.
	Exports the graph as a .png image or as a PDF.
	Specifies the period over which the change measurements are taken.

Threat Monitor Report

The App-Scope Threat Monitor report displays a count of the top threats over the selected time period. For example, the following figure shows the top 10 threat types over the last 6 hours.



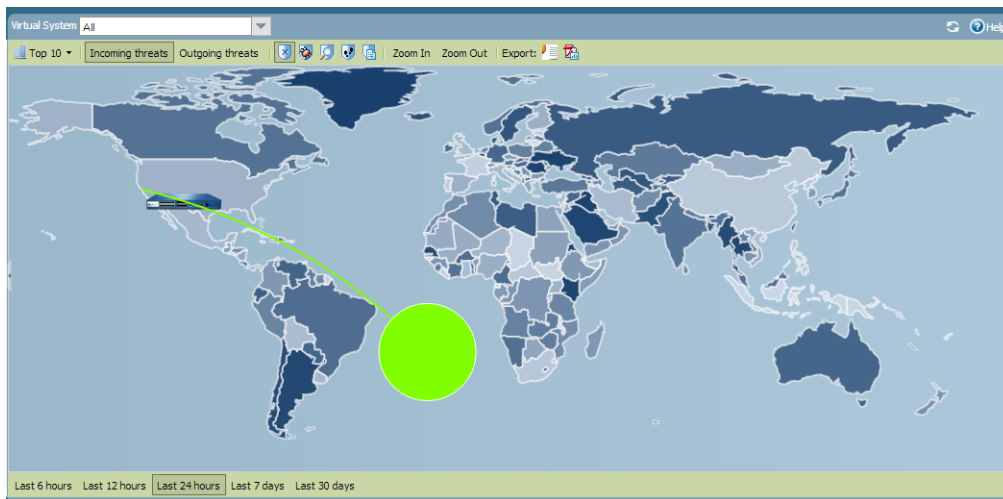
Each threat type is color-coded as indicated in the legend below the chart. The Threat Monitor report contains the following buttons and options.

Button	Description
	Determines the number of records with the highest measurement included in the chart.
	Determines the type of item measured: Threat, Threat Category, Source, or Destination.
	Applies a filter to display only the selected type of items.
	Determines whether the information is presented in a stacked column chart or a stacked area chart.
	Exports the graph as a .png image or as a PDF.
	Specifies the period over which the measurements are taken.

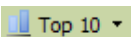
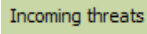
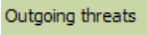

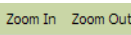
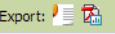
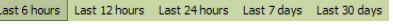
Threat Map Report

The App-Scope Threat Map report shows a geographical view of threats, including severity. Each threat type is color-coded as indicated in the legend below the chart.

The firewall uses geolocation for creating threat maps. The firewall is placed at the bottom of the threat map screen, if you have not specified the geolocation coordinates (**Device > Setup > Management**, General Settings section) on the firewall.




The Threat Map report contains the following buttons and options.

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

Network Monitor Report

The App-Scope Network Monitor report displays the bandwidth dedicated to different network functions over the specified period of time. Each network function is color-coded as indicated in the legend below the chart. For example, the image below shows application bandwidth for the past 7 days based on session information.

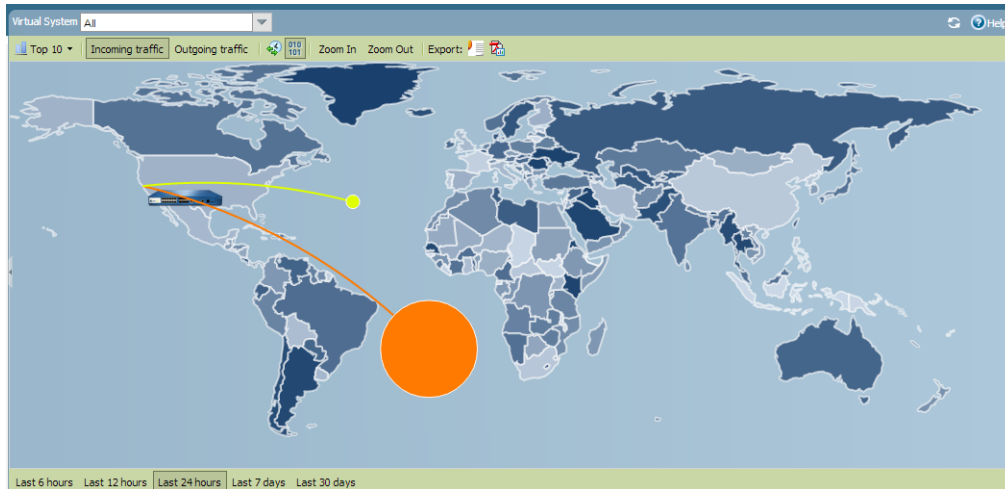


The Network Monitor report contains the following buttons and options.

Button	Description
	Determines the number of records with the highest measurement included in the chart.
	Determines the type of item reported: Application, Application Category, Source, or Destination.
	Applies a filter to display only the selected item. None displays all entries.
	Determines whether to display session or byte information.
	Exports the graph as a .png image or as a PDF.
	Determines whether the information is presented in a stacked column chart or a stacked area chart.
	Indicates the period over which the change measurements are taken.

Traffic Map Report

The App-Scope Traffic Map report shows a geographical view of traffic flows according to sessions or flows. The firewall uses geolocation for creating traffic maps. The firewall is placed at the bottom of the traffic map screen, if you have not specified the geolocation coordinates (**Device > Setup > Management**, General Settings section) on the firewall.



Each traffic type is color-coded as indicated in the legend below the chart. The Traffic Map report contains the following buttons and options.

Buttons	Description
	Determines the number of records with the highest measurement included in the chart.
	Displays incoming threats.
	Displays outgoing threats.
	Determines whether to display session or byte information.
	Zoom in and zoom out of the map.
	Exports the graph as a .png image or as a PDF.
	Indicates the period over which the change measurements are taken.

Take Packet Captures

PAN-OS supports packet capture for troubleshooting or detecting unknown applications. You can define filters such that only the packets that match the filters are captured. The packet captures are locally stored on the device and are available for download to your local computer.



Packet Capture is for troubleshooting only. This feature can cause the system performance to degrade and should be used only when necessary. Remember to disable the feature after you complete the packet capture.

The following table describes the packet capture settings on **Monitor > Packet Capture**.

Field	Description
Manage Filters	<p>Click Manage Filters, click Add to add a new filter, and specify the following information:</p> <ul style="list-style-type: none"> • Id—Enter or select an identifier for the filter. • Ingress Interface—Select the firewall interface. • Source—Specify the source IP address. • Destination—Specify the destination IP address. • Src Port—Specify the source port. • Dest Port—Specify the destination port. • Proto—Specify the protocol to filter. • Non-IP—Choose how to treat non-IP traffic (exclude all IP traffic, include all IP traffic, include only IP traffic, or do not include an IP filter). • IPv6—Select the check box to include IPv6 packets in the filter.
Filtering	Click to toggle the filtering selections on or off.
Pre-Parse Match	<p>Click to toggle the pre-parse match option on or off.</p> <p>The pre-parse-match option is added for advanced troubleshooting purposes. After a packet enters the ingress port, it proceeds through several processing steps before it is parsed for matches against pre-configured filters.</p> <p>It is possible for a packet, due to a failure, to not reach the filtering stage. This can occur, for example, if a route lookup fails.</p> <p>Set the pre-parse-match setting to ON to emulate a positive match for every packet entering the system. This allows the firewall to capture even the packets that do not reach the filtering process. If a packet is able to reach the filtering stage, it is then processed according to the filter configuration and discarded if it fails to meet filtering criteria.</p>

Field	Description
<p>Packet Capture</p>	<p>Click to toggle packet capturing on or off.</p> <p>For anti-spyware and vulnerability protection profiles, you can enable extended packet captures for rules and exceptions defined in the profile. This functionality allows the firewall to capture from 1 to 50 packets and provides more context when analyzing the threat logs.</p> <p>To define the extended packet capture length:</p> <ol style="list-style-type: none"> 1. Select Device > Setup > Content-ID. 2. Edit the Threat Detection Settings section to specify the Capture Length for the number of packets to capture. <div data-bbox="548 594 930 674" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p>Threat Detection Settings</p> <p>Extended Packet Capture Length (packets) 5</p> </div> <ol style="list-style-type: none"> 3. View the packet capture in Monitor > Logs > Threat. <p>Locate the threat log entry and click the green arrow (Packet Capture) icon in the corresponding row to view the capture.</p>
<p>Packet Capture Stage</p>	<p>Select Add and specify the following:</p> <ul style="list-style-type: none"> • Stage—Indicate the point at which to capture the packet: • drop—When packet processing encounters an error and the packet is to be dropped. • firewall—When the packet has a session match or a first packet with a session is successfully created. • receive—When the packet is received on the dataplane processor. • transmit—When the packet is to be transmitted on the dataplane processor. • File—Specify the capture file name. The file name should begin with a letter and can include letters, digits, periods, underscores, or hyphens. • Packet Count—Specify the number of packets after which capturing stops. • Byte Count—Specify the number of bytes after which capturing stops.
<p>Captured Files</p>	<p>Select Delete to remove a packet capture file from the list displaying captured files.</p>
<p>Clear All Settings</p>	<p>Select Clear All Settings to clear all packet capture settings.</p>

Monitor the Firewall

The following sections describe the methods you can use to monitor the firewall and provide basic setup instructions:

- ▲ [Monitor Applications and Threats](#)
- ▲ [Monitor Log Data](#)
- ▲ [Monitor the Dashboard](#)
- ▲ [View Reports](#)



You can also configure the firewall (excluding PA-4000 Series and PA-7050 firewalls) to export flow data to a NetFlow collector for analysis and reporting. To configure NetFlow Settings, refer to the [PAN-OS Web Interface Reference Guide](#).

Monitor Applications and Threats

All Palo Alto Networks next-generation firewalls come equipped with the [App-ID](#) technology, which identifies the applications traversing your network, irrespective of protocol, encryption, or evasive tactic. You can then [Use the Application Command Center](#) to monitor the applications. ACC graphically summarizes the log database to highlight the applications traversing your network, who is using them, and their potential security impact. ACC is dynamically updated, using the continuous traffic classification that App-ID performs; if an application changes ports or behavior, App-ID continues to see the traffic, displaying the results in ACC.

You can quickly investigate new, risky, or unfamiliar applications that appear in ACC with a single click that displays a description of the application, its key features, its behavioral characteristics, and who is using it. Additional visibility into URL categories, threats, and data provides a complete and well-rounded picture of network activity. With ACC, you can very quickly learn more about the traffic traversing the network and then translate that information into a more informed security policy.

URL	Category	Time
www.sherriffguadagno.com/favicon.ico	personal-sites-and-blogs	09/19 09:16:42
ichart.finance.yahoo.com/instrument/L0/N225/chart?range=1dℑsize=239x110	news-and-media	09/19 09:16:42
b.static.ak.fbcdn.net/src.php/z3OWB/pkg/81s3vfbn/en_US/141/165156/js/dk4xcv99hrsvsko0.pkg.js	content-delivery-networks	09/19 09:16:42
b.static.ak.fbcdn.net/src.php/z31SR/b55dq2qo/en_US/146748/css/UIProfileBox.css	content-delivery-networks	09/19 09:16:42
b.static.ak.fbcdn.net/src.php/z83FN/pkg/dvtt7mrg/en_US/141/169335/css/3w4176j09e2o88ow.pkg.css	content-delivery-networks	09/19 09:16:42
entertain.teenee.com/thaistar/vote.gif	personal-sites-	09/19 09:16:42

File Name	Name	Time
1537077.flv	FLV File	09/19 09:16:43
validateForm.js	Confidential	09/19 09:16:43
	MP4 Detected	09/19 09:16:37
i18n.js	Confidential	09/19 09:16:33
	Confidential	09/19 09:16:32
i18n.js	Confidential	09/19 09:16:32
Week 4 Summary.doc	Microsoft MSOFFICE	09/19 09:16:31
Week 4 Summary.doc	Microsoft Word	09/19 09:16:31
pa4.doc	Microsoft Word	09/19 09:16:28

ACC Risk Factor

Last 60 minutes: 1 2 3 4 5

3.4

Name	Severity	Time
FTP Login Failed	informational	09/19 09:16:43
MIT Kerberos kadmind RPC Library Vulnerability	high	09/19 09:16:42
Unix Authentication Stack Overflow	informational	09/19 09:16:42
HTTP Forbidden Error	informational	09/19 09:16:42
HTTP WWW-Authentication Failed	informational	09/19 09:16:41

Monitor Log Data

All Palo Alto Networks next-generation firewalls can generate log files that provide an audit trail of the activities and events on the firewall. There are separate logs for separate types of activities and events. For example, the Threat logs record all traffic that causes the firewall to generate a security alarm, URL Filtering logs record all traffic that matches a URL Filtering profile attached to a security policy, and Config logs record all changes to the firewall configuration.

You can either [Forward Logs to External Services](#) or you can view logs locally on the device as follows:

- ▲ [View the Log Files](#)
- ▲ [Filter Log Data](#)

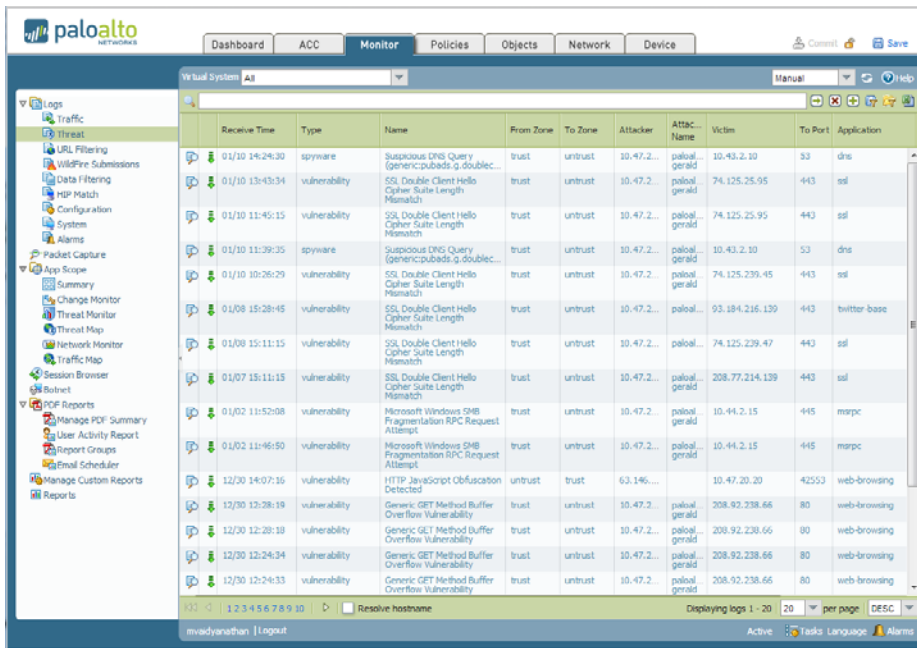
View the Log Files

The firewall maintains logs for WildFire, configurations, system, alarms, traffic flows, threats, URL filtering, data filtering, and Host Information Profile (HIP) matches. You can view the current logs at any time. To locate specific entries, you can apply filters to most of the log fields.

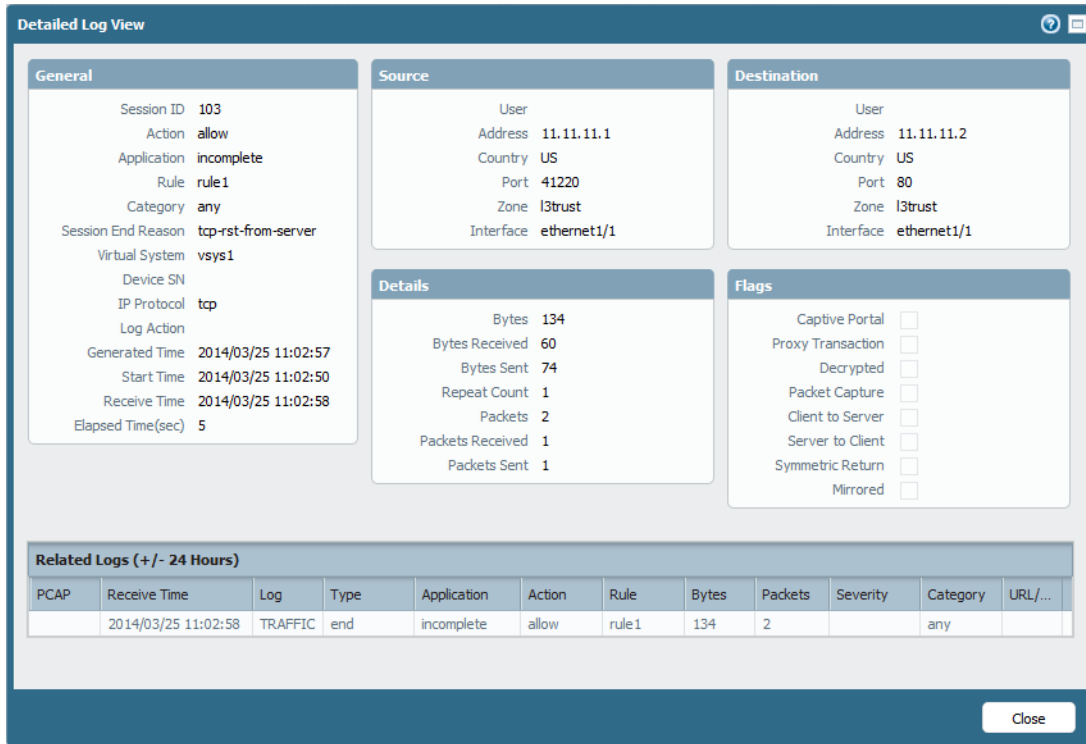


The firewall displays the information in logs so that role-based administration permissions are respected. When you display logs, only the information that you have permission to see is included. For information on administrator permissions, see [Administrative Roles](#).


By default all log files are generated and stored locally on the firewall. You can view these log files directly (**Monitor > Logs**):





To display additional details, click the spyglass icon for an entry.



The following table includes information on each log type:

Log Description Charts	Description
Traffic	<p>Displays an entry for the start and end of each session. Each entry includes the date and time, source and destination zones, addresses and ports, application name, security rule name applied to the flow, rule action (allow, deny, or drop), ingress and egress interface, number of bytes, and session end reason.</p> <p>Click  next to an entry to view additional details about the session, such as whether an ICMP entry aggregates multiple sessions between the same source and destination (the Count value will be greater than one).</p> <p>The Type column indicates whether the entry is for the start or end of the session, or whether the session was denied or dropped. A drop indicates that the security rule that blocked the traffic specified any application, while a deny indicates the rule identified a specific application.</p> <p>If traffic is dropped before the application is identified, such as when a rule drops all traffic for a specific service, the application is shown as not-applicable.</p>

Log Description Charts	Description
Threat	<p>Displays an entry when traffic matches a Security Profile (Antivirus, Anti-spyware, Vulnerability, URL Filtering, File Blocking, Data Filtering, or DoS Protection) that is attached to a security policy on the firewall. Each entry includes the date and time, a threat name or URL, the source and destination zones, addresses, and ports, the application name, and the alarm action (allow or block) and severity.</p> <p>Click  next to an entry to view additional details about the threat, such as whether the entry aggregates multiple threats of the same type between the same source and destination (the Count value will be greater than one).</p> <p>The Type column indicates the type of threat, such as “virus” or “spyware.” The Name column is the threat description or URL, and the Category column is the threat category (such as “keylogger”) or URL category.</p> <p>If local packet captures are enabled, click  next to an entry to access the captured packets. To enable local packet captures, see Take Packet Captures.</p>
URL Filtering	<p>Displays logs for all traffic that matches a URL Filtering profile attached to a security policy. For example, if policy blocks access to specific web sites and web site categories or if policy is configured to generate an alert when a web site is accessed. For information on defining URL filtering profiles, see URL Filtering.</p>
WildFire Submissions	<p>Displays logs for files that are uploaded and analyzed by the WildFire cloud; log data is sent back to the device after analysis, along with the analysis results.</p>
Data Filtering	<p>Displays logs for the security policies that help prevent sensitive information such as credit card or social security numbers from leaving the area protected by the firewall. See Set Up Data Filtering for information on defining data filtering profiles.</p> <p>This log also shows information for file-blocking profiles. For example, if you are blocking <i>.exe</i> files, the log will show the files that were blocked. If you forward files to WildFire, you will see the results of that action. In this case, if you are forwarding PE files to WildFire, for example, the log will show that the file was forwarded and will also show the status on whether or not it was uploaded to WildFire successfully.</p>
Configuration	<p>Displays an entry for each configuration change. Each entry includes the date and time, the administrator username, the IP address from where the change was made, the type of client (XML, Web or CLI), the type of command executed, whether the command succeeded or failed, the configuration path, and the values before and after the change.</p>
System	<p>Displays an entry for each system event. Each entry includes the date and time, the event severity, and an event description.</p>
HIP Match	<p>Displays traffic flows that match a HIP Object or HIP Profile that you have configured.</p>

Filter Log Data

Each log page has a filter area at the top of the page.



Use the filter area as follows:

- Click any of the underlined links in the log listing to add that item as a log filter option. For example, if you click the **Host** link in the log entry for 10.0.0.252 and **Web Browsing**, both items are added, and the search will find entries that match both (AND search).
- To define other search criteria, click **Add Log Filter**. Select the type of search (and/or), the attribute to include in the search, the matching operator, and the values for the match, if appropriate. Click **Add** to add the criterion to the filter area on the log page, and then click **Close** to close the pop-up window. Click **Apply Filter** to display the filtered list.



If the **Value** string matches an **Operator** (such as **has** or **in**), enclose the string in quotation marks to avoid a syntax error. For example, if you filter by destination country and use **IN** as a **Value** to specify INDIA, enter the filter as (`dstloc eq "IN"`).

You can combine filter expressions added on the log page with those you define in the Add Log Filter dialog. The filter field on the log page displays each filter as an entry.

If you add a **Receive Time** filter with the **Operator** set to **in** and the **Value** set to **Last 60 seconds**, some of the page links on the log viewer might not show results because the number of pages might grow or shrink due to the dynamic nature of the selected time.

- To clear filters and redisplay the unfiltered list, click **Clear Filter**.
- To save your selections as a new filter, click **Save Filter**, enter a name for the filter, and click **OK**.
- To export the current log listing (as shown on the page, including any applied filters) click **Save Filter**. Select whether to open the file or save it to disk, and select the check box if you want to always use the same option. Click **OK**.
- To export the current log listing in CSV Format, select the **Export to CSV** icon. By default, exporting the log listing to CSV format generates a CSV report with up to 2,000 rows of logs. To change the limit for rows displayed in CSV reports, use the **Max Rows in CSV Export** field on the **Log Export and Reporting** tab (select **Device > Setup > Management > Logging and Reporting Settings**).

To change the automatic refresh interval, select an interval from the drop-down (**1 min**, **30 seconds**, **10 seconds**, or **Manual**).

To change the number of log entries per page, select the number of rows from the **Rows** drop-down.

Log entries are retrieved in blocks of 10 pages. Use the paging controls at the bottom of the page to navigate through the log list. Select the **Resolve Hostname** check box to begin resolving external IP addresses to domain names.

Monitor the Dashboard

You can also monitor the local log data directly from the **Dashboard** by adding the associated widgets:

The screenshot displays the Palo Alto Networks Dashboard interface. At the top, there are navigation tabs for Dashboard, ACC, Monitor, Policies, Objects, Network, and Device. Below the navigation, there are settings for 'Layout: 3 Columns', 'Widgets', and 'Last updated: 09:16:44'. The dashboard is divided into four main sections:

- URL Filtering Logs:** A table listing various URLs, categories, and times.
- Data Logs:** A table listing file names, names, and times.
- ACC Risk Factor:** A gauge showing a risk score of 3.4 over the last 60 minutes.
- Threat Logs:** A table listing threat names, severities, and times.

URL	Category	Time
www.sheriffguadagno.com/favico	personal-sites-and-blogs	09/19 09:16:42
ichart.finance.yahoo.com/instrum	news-and-media	09/19 09:16:42
b.static.ak.fbcdn.net/frsrc.php	content-delivery-networks	09/19 09:16:42
b.static.ak.fbcdn.net/frsrc.php	content-delivery-networks	09/19 09:16:42
b.static.ak.fbcdn.net/frsrc.php	content-delivery-networks	09/19 09:16:42
b.static.ak.fbcdn.net/frsrc.php	content-delivery-networks	09/19 09:16:42
entertain.teenee.com/thaistar	personal-sites-and-blogs	09/19 09:16:42
www.jstor.org/templates	training-and-tools	09/19 09:16:42
statase.webtrendslive.com/	computer-and-internet-info	09/19 09:16:42
newsrss.bbc.co.uk/rss	news-and-media	09/19 09:16:42
www.jstor.org/bbb.csueastbay.e	reference-and-research	09/19 09:16:42

File Name	Name	Time
1537077.flv	FLV File	09/19 09:16:43
validateForm.js	Confidential	09/19 09:16:43
	MP4 Detected	09/19 09:16:37
i18n.js	Confidential	09/19 09:16:33
	Confidential	09/19 09:16:32
i18n.js	Confidential	09/19 09:16:32
Week 4 Summary.doc	Microsoft MSOFFICE	09/19 09:16:31
Week 4 Summary.doc	Microsoft Word	09/19 09:16:31
pa4.doc	Microsoft Word	09/19 09:16:28
pa4.doc	Microsoft MSOFFICE	09/19 09:16:28

Score
3.4

Name	Severity	Time
FTP Login Failed	informational	09/19 09:16:43
MIT Kerberos kadmind RPC Library Vulnerability	high	09/19 09:16:42
HTTP Forbidden Error	informational	09/19 09:16:42
HTTP WWW-Authentication Failed	informational	09/19 09:16:41
FTP Login Failed	informational	09/19 09:16:40
FTP: login brute force attempt	high	09/19 09:16:40
HTTP Forbidden Error	informational	09/19 09:16:38
FTP Login Failed	informational	09/19 09:16:36
Microsoft ASP.NET Remote Unauthenticated Denial of Service Vulnerability	low	09/19 09:16:34

Description	Time
User 'Imelvin' logged in via Web from 64.124.57.5 using https	09/19 09:15:39
User 'Imelvin' authenticated. From: 64.124.57.5.	09/19 09:15:39

View Reports

The firewall also uses the log data to generate reports (**Monitor > Reports**) that display the log data in a tabular or graphical format. See [About Reports](#) for more details on the predefined and custom reports available on the firewall.

Threat/Content Name	ID	Count
1 FTP Login Failed	40000	22.3 K
2 FTP: login brute force attempt	40001	12.3 K
3 HTTP Forbidden Error	34556	6.7 K
4 HTTP OPTIONS Method	30520	4.5 K
5 HTTP WWW-Authentication Failed	31708	2.2 K
6 Microsoft ASP.NET Remote Unauthenticated Denial of Service Vulnerability	32513	2.0 K
7 DNS ANY Request	34842	999
8 HTTP Non RFC-Compliant Response Found	32880	699
9 Adobe PDF File With Embedded Javascript	31971	446
10 Microsoft ASP.Net Information Leak Vulnerability	33435	377
11 Generic GET Method Buffer Overflow Vulnerability	34267	141
12 NetBIOS nbstat query	31707	135
13 PDF Exploit Evasion Found	33939	123
14 MIT Kerberos kadmind RPC Library Unix Authentication Stack Overflow Vulnerability	30243	102
15 SSH2 Login Attempt	31914	98
16 SIP Register Request Attempt	33592	98
17 HTTP JavaScript Obfuscation Detected	31825	96
18 SSL Renegotiation Denial of Service Vulnerability	33862	94
19 PHP CGI Query String Parameter Handling Information Disclosure Vulnerability	34860	78
20 Apache Un-terminated Request With Content Length Denial Of Service Attack	32452	51
21 HTTP JavaScript Obfuscation Detected	31826	48

Forward Logs to External Services

Depending on the type and severity of the data in the log files, you may want to be alerted to critical events that require your attention, or you may have policies that require you to archive the data for longer than it can be stored on the firewall. In these cases you will want to forward your log data to an external service for archive, notification, and/or analysis.

To forward log data to an external service you must complete the following tasks:

- Configure the firewall to access the remote services that will be receiving the logs. See [Define Remote Logging Destinations](#).
- Configure each log type for forwarding. See [Enable Log Forwarding](#).
- For traffic and threat logs, enabling log forwarding includes setting up a log forwarding profile or a default log forwarding profile. For details, see [Log Forwarding Profiles](#).

Define Remote Logging Destinations

In order to reach an external service—such as a Syslog server or SNMP trap manager—the firewall must know the details of how to access and, if necessary, authenticate to the service. On the firewall, you define this information in a Server Profile. You must create a Server Profile for each external service you want the firewall to interact with. The type of logging destinations you need to set up and which logs you forward will depend on your needs. Some common log forwarding scenarios include the following:

- ❑ For immediate notification about critical system events or threats that require your attention, you can generate SNMP traps or send email alerts. See [Set Up Email Alerts](#) and/or [Set Up SNMP Trap Destinations](#).
- ❑ For long-term storage and archival of data and for centralized device monitoring, you can send the log data to a Syslog server. See [Define Syslog Servers](#). This enables integration with third-party security monitoring tools, such as Splunk! or ArcSight. You can also secure the channel between the firewall and the Syslog server. See [Configure the Firewall to Authenticate to the Syslog Server](#).
- ❑ For aggregation and reporting of log data from multiple Palo Alto Networks firewalls, you can forward logs to a Panorama Manager or Panorama Log Collector. See [Enable Log Forwarding](#).

You can define as many Server Profiles as you need. For example, you could use separate Server Profiles to send traffic logs to one Syslog server and system logs to a different one. Or, you could include multiple server entries in a single Server Profile to enable you to log to multiple Syslog servers for redundancy.



By default, all log data is forwarded over the MGT interface. If you plan to use an interface other than MGT, you will need to configure a Service Route for each service to which you plan to forward logs as described in Step 5 in the procedure to [Set Up Network Access for External Services](#).

Set Up Email Alerts


Set Up Email Alerts	
<p>Step 1 Create a Server Profile for your email server.</p>	<ol style="list-style-type: none"> 1. Select Device > Server Profiles > Email. 2. Click Add and then enter a Name for the profile. 3. (Optional) Select the virtual system to which this profile applies from the Location drop-down. 4. Click Add to add a new email server entry and enter the information required to connect to the Simple Mail Transport Protocol (SMTP) server and send email (you can add up to four email servers to the profile): <ul style="list-style-type: none"> • Server—Name to identify the mail server (1-31 characters). This field is just a label and does not have to be the host name of an existing SMTP server. • Display Name—The name to show in the From field of the email. • From—The email address where notification emails will be sent from. • To—The email address to which notification emails will be sent. • Additional Recipient—If you want the notifications sent to a second account, enter the additional address here. You can only add one additional recipient. To add multiple recipients, add the email address of a distribution list. • Gateway—The IP address or host name of the SMTP gateway to use to send the emails. 5. Click OK to save the server profile.
<p>Step 2 (Optional) Customize the format of the email messages the firewall sends.</p>	<p>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.</p>
<p>Step 3 Save the server profile and commit your changes.</p>	<ol style="list-style-type: none"> 1. Click OK to save the profile. 2. Click Commit to save the changes to the running configuration.

Set Up SNMP Trap Destinations

Simple Network Management Protocol (SNMP) is a standard facility for monitoring the devices on your network. You can configure the firewall to send SNMP traps to your SNMP management software to alert you to critical system events or threats that require your immediate attention.



You can also use SNMP to monitor the firewall. In this case, your SNMP manager must be configured to get statistics from the firewall rather than (or in addition to) having the firewall send traps to the manager. For more information, see [Configure the Firewall to Authenticate to the Syslog Server](#).

Set Up SNMP Trap Destinations	
<p>Step 1 (SNMP v3 only) Get the engine ID for the firewall.</p> <p> In many cases, the MIB browser or SNMP manager will automatically discover the engine ID upon successful connection to the SNMP agent on the firewall. You can usually find this information in the agent settings section of the interface. Refer to the documentation for your specific product for instructions on finding the agent information.</p>	<p>In order to find out the firewall’s engine ID, you must configure the firewall for SNMP v3 and send a GET message from your SNMP manager or MIB browser as follows:</p> <ol style="list-style-type: none"> 1. Enable the interface to allow inbound SNMP requests: <ul style="list-style-type: none"> • If you will be receiving SNMP GET messages on the MGT interface, select Device > Setup > Management and click to Edit in the Management Interface Settings section of the screen. In the Services section, select the SNMP check box and then click OK. • If you will be receiving SNMP GET messages on a different interface, you must associate a management profile with the interface and enable SNMP management. 2. Configure the firewall for SNMP v3 as described in Step 2 in Set Up SNMP Monitoring. If you do not configure the firewall for SNMP v3, your MIB browser will not allow you to GET the engine ID. 3. Connect your MIB browser or SNMP manager to the firewall and run a GET for OID 1.3.6.1.6.3.10.2.1.1.0. The value that is returned is the unique engine ID for the firewall.
<p>Step 2 Create a Server Profile that contains the information for connecting and authenticating to the SNMP manager(s).</p> <ol style="list-style-type: none"> 1. Select Device > Server Profiles > SNMP Trap. 2. Click Add and then enter a Name for the profile. 3. (Optional) Select the virtual system to which this profile applies from the Location drop-down. 4. Specify the version of SNMP you are using (V2c or V3). 5. Click Add to add a new SNMP Trap Receiver entry (you can add up to four trap receivers per server profile). The required values depend on whether you are using SNMP V2c or V3 as follows: <p style="margin-left: 20px;">SNMP V2c</p> <ul style="list-style-type: none"> • Server—Name to identify the SNMP manager (1-31 characters). This field is just a label and does not have to be the host name of an existing SNMP server. • Manager—The IP address of the SNMP manager to which you want to send traps. • Community—The community string required to authenticate to the SNMP manager. <p style="margin-left: 20px;">SNMP V3</p> <ul style="list-style-type: none"> • Server—Name to identify the SNMP manager (1-31 characters). This field is just a label and does not have to be the host name of an existing SNMP server. • Manager—The IP address of the SNMP manager to which you want to sent traps. • User—The username required to authenticate to the SNMP manager. • EngineID—The engine ID of the firewall, as identified in Step 1. This is a hexadecimal value from 5 to 64 bytes with a 0x prefix. Each firewall has a unique engine ID. • Auth Password—The password to be used for authNoPriv level messages to the SNMP manager. This password will be hashed using Secure Hash Algorithm (SHA-1), but will not be encrypted. • Priv Password—The password to be used for authPriv level messages to the SNMP manager. This password be hashed using SHA and will be encrypted using Advanced Encryption Standard (AES 128). 6. Click OK to save the server profile. 	

Set Up SNMP Trap Destinations (Continued)	
Step 3 (Optional) Set up a service route for SNMP traps.	By default, SNMP traps are sent over the MGT interface. If you want to use a different interface for SNMP traps, you must edit the service route to enable the firewall to reach your SNMP manager. See Set Up Network Access for External Services for instructions.
Step 4 Commit your changes.	Click Commit . The device may take up to 90 seconds to save your changes.
Step 5 Enable the SNMP manager to interpret the traps it receives from the firewall.	Load the PAN-OS MIB files into your SNMP management software and compile them. Refer to the documentation for your SNMP manager for specific instructions on how to do this.

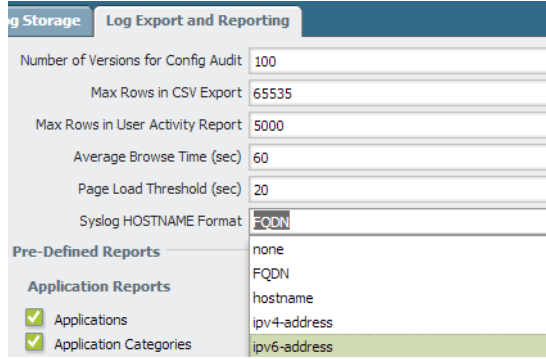


Define Syslog Servers

Syslog is a standard log transport mechanism that enables the aggregation of log data from different network devices—such as routers, firewalls, printers—from different vendors into a central repository for archive, analysis, and reporting.

The firewall generates six types of logs that can be forwarded to an external syslog server: traffic, threat, WildFire, host information profile (HIP) match, config, and system. If you want to forward all or some of these logs to an external service for long-term storage and analysis, you can use TCP or SSL for reliable and secure transport of logs, or UDP for non-secure transport.

Set Up Syslog Forwarding

- | | |
|--|---|
| <p>Step 1 Create a Server Profile that contains the information for connecting to the Syslog server(s).</p> | <ol style="list-style-type: none"> 1. Select Device > Server Profiles > Syslog. 2. Click Add and then enter a Name for the profile. 3. (Optional) Select the virtual system to which this profile applies from the Location drop-down. 4. Click Add to add a new Syslog server entry and enter the information required to connect to the Syslog server (you can add up to four Syslog servers to the same profile): <ul style="list-style-type: none"> • Name—Unique name for the server profile. • Server—IP address or fully qualified domain name (FQDN) of the Syslog server. • Transport—Select TCP, UDP, or SSL as the method of communication with the syslog server. • Port—The port number on which to send Syslog messages (default is UDP on port 514); you must use the same port number on the firewall and the Syslog server. • Format—Select the Syslog message format to use, BSD or IETF. Traditionally, BSD format is over UDP and IETF format is over TCP/SSL. For setting up secure syslog forwarding with client authentication, see Configure the Firewall to Authenticate to the Syslog Server. • Facility—Select one of the Syslog standard values, which is used to calculate the priority (PRI) field in your Syslog server implementation. You should select the value that maps to how you use the PRI field to manage your Syslog messages. 5. (Optional) To customize the format of the Syslog messages the firewall sends, select the Custom Log Format tab. For details on how to create custom formats for the various log types, refer to the Common Event Format Configuration Guide. 6. Click OK to save the server profile. |
|--|---|

Set Up Syslog Forwarding (Continued)	
<p>Step 2 (Optional) Configure the header format used in Syslog messages. Choosing the header format offers more flexibility in filtering and reporting on the log data for some SIEMs.</p>  <p> This is a global setting and applies to all syslog server profiles configured on the appliance.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Management and click the Edit icon in the Logging and Reporting Settings section. 2. Select Log Export and Reporting. 3. Select one of the following options from the Send Hostname in Syslog drop-down: <ul style="list-style-type: none"> • FQDN—(the default) Concatenates the hostname and domain name defined on the sending device. • hostname—Uses the hostname defined on the sending device. • ipv4-address—Uses the IPv4 address of the interface used to send logs on the device. By default, this is the MGT interface of the device. • ipv6-address—Uses the IPv6 address of the interface used to send logs on the device. By default, this is the MGT interface of the device. • none—Leaves the hostname field unconfigured on the device. There is no identifier for the device that sent the logs. 4. Click OK and Commit.
<p>Step 3 Commit your changes.</p>	<p>Click Commit. The device may take up to 90 seconds to save your changes.</p>
<p>Step 4 Enable log forwarding.</p>	<p>See Enable Log Forwarding.</p> <p>You must configure each log type for forwarding and specify the severity for which the event is logged.</p> <p> WildFire logs are a type of threat log, but they are not logged and forwarded along with threat logs. While WildFire logs use the same syslog format as threat logs, the threat subtype is preset to WildFire. Therefore, you must enable logging/ forwarding for WildFire logs distinctly from threat logs.</p>
<p>Step 5 Review the logs on the syslog server.</p>	<p>To parse the logs, see Syslog Field Descriptions.</p>

Configure the Firewall to Authenticate to the Syslog Server

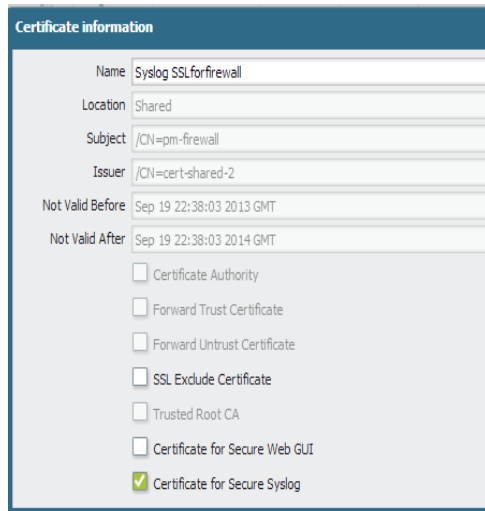
To enable client authentication for syslog over SSL, you can use a trusted CA or a self-signed CA for generating certificates that can be used for secure syslog communication. Check for the following when generating a certificate for secure syslog communication:

- The private key must be available on the sending device; the keys cannot be stored on a Hardware Security Module (HSM).
- The subject and the issuer for the certificate must not be identical.

- The certificate is neither a trusted CA nor a certificate signing request (CSR). Neither of these types of certificates can be enabled for secure syslog communication.

Configure the Firewall to Authenticate to the Syslog Server

Step 1 If the syslog server requires client authentication, generate the certificate for secure communication. For details on certificates, see [Certificate Management](#).



To verify that the sending device is authorized to communicate with the syslog server, you must enable the following:

- The server and the sending device must have certificates that are signed by the Enterprise CA; or you can generate a self-signed certificate on the firewall, export the CA root certificate from the firewall and import it in to the syslog server.
- Use the Enterprise CA or the self-signed certificate to generate a certificate with the IP address of the sending device (as the Common Name) and enabled for use in secure syslog communication. The syslog server uses this certificate to verify that the firewall is authorized to communicate with the syslog server.

Use the following steps to generate the certificate on the firewall or Panorama:

1. Select **Device > Certificate Management > Certificates > Device Certificates**.
2. Click **Generate** to create a new certificate that will be signed by a trusted CA or the self-signed CA.
3. Enter a **Name** for the certificate.
4. In **Common Name**, enter the IP address of the device sending logs to the syslog server.
5. Select **Shared** if you want the certificate to be a shared certificate on Panorama or to be shared by all virtual systems in a multiple virtual system firewall.
6. In **Signed by**, select the trusted CA or the self-signed CA that is trusted by both the syslog server and the sending device.
7. Click **Generate**. The certificate and the keypair will be generated.
8. Click the link with name of the certificate and enable the **Certificate for Secure Syslog** option for secure access to the syslog server.
9. **Commit** the changes.
10. Verify the certificate details and that it is marked for **Usage** as **Certificate for Secure Syslog**.

Name	Location	Subject	Issuer	CA	Key	Expires	Status	Usage
cert-shared-2	Shared	cert-shared-2	cert-shared-2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Feb 26 23:03:44 2014 GMT	valid	Certificate for Secure Web GUI Trusted Root CA Certificate
Syslog SSL for firewall	Shared	pm-firewall	cert-shared-2	<input type="checkbox"/>	<input checked="" type="checkbox"/>	Sep 19 22:38:03 2014 GMT	valid	Certificate for Secure Syslog

Enable Log Forwarding

After you create the Server Profiles that define where to send your logs (see [Define Remote Logging Destinations](#)), you must enable log forwarding. For each log type, you can specify whether to forward it to Syslog, email, SNMP trap receiver, and/or Panorama.



Before you can forward log files to a Panorama Manager or a Panorama Log Collector, the firewall must be configured as a [managed device](#). You can then enable log forwarding to Panorama for each type of log. For logs forwarded to Panorama, support for centralized log forwarding to an external syslog server is available.

You can use Secure Copy (SCP) commands from the CLI to export the entire log database (logdb) to an SCP server and import it to another firewall: refer to the [PAN-OS Command Line Interface \(CLI\) Reference Guide](#). Because the log database is too large for an export or import to be practical on the following platforms, they do not support these options: PA-7050 firewall (all PAN-OS releases), Panorama virtual appliance running Panorama 6.0 or later releases, and Panorama M-Series appliances (all Panorama releases).

The way you enable forwarding depends on the log type:

- **Traffic Logs**—You enable forwarding of Traffic logs by creating a Log Forwarding Profile (**Objects > Log Forwarding**) and adding it to the security policies you want to trigger the log forwarding. Only traffic that matches a specific rule within the security policy will be logged and forwarded. For details on setting up a log forwarding profile, see [Log Forwarding Profiles](#).
- **Threat Logs**—You enable forwarding of Threat logs by creating a Log Forwarding Profile (**Objects > Log Forwarding**) that specifies which severity levels you want to forward and then adding it to the security policies for which you want to trigger the log forwarding. A Threat log entry will only be created (and therefore forwarded) if the associated traffic matches a Security Profile (Antivirus, Anti-spyware, Vulnerability, URL Filtering, File Blocking, Data Filtering, or DoS Protection). For details on setting up a log forwarding profile, see [Log Forwarding Profiles](#). The following table summarizes the threat severity levels:

Severity	Description
Critical	Serious threats, such as those that affect default installations of widely deployed software, result in root compromise of servers, and the exploit code is widely available to attackers. The attacker usually does not need any special authentication credentials or knowledge about the individual victims and the target does not need to be manipulated into performing any special functions.
High	Threats that have the ability to become critical but have mitigating factors; for example, they may be difficult to exploit, do not result in elevated privileges, or do not have a large victim pool.
Medium	Minor threats in which impact is minimized, such as DoS attacks that do not compromise the target or exploits that require an attacker to reside on the same LAN as the victim, affect only non-standard configurations or obscure applications, or provide very limited access. In addition, WildFire Submissions log entries with a malware verdict are logged as Medium.
Low	Warning-level threats that have very little impact on an organization's infrastructure. They usually require local or physical system access and may often result in victim privacy or DoS issues and information leakage. Data Filtering profile matches are logged as Low.

Severity	Description
Informational	Suspicious events that do not pose an immediate threat, but that are reported to call attention to deeper problems that could possibly exist. URL Filtering log entries and WildFire Submissions log entries with a benign verdict are logged as Informational.

- **Config Logs**—You enable forwarding of Config logs by specifying a Server Profile in the log settings configuration. (**Device > Log Settings > Config Logs**).
- **System Logs**—You enable forwarding of System logs by specifying a Server Profile in the log settings configuration. (**Device > Log Settings > System Logs**). You must select a Server Profile for each severity level you want to forward. For a partial list of system log messages and their corresponding severity levels, refer to the [System Log Reference](#). The following table summarizes the system log severity levels:

Severity	Description
Critical	Hardware failures, including HA failover and link failures.
High	Serious issues, including dropped connections with external devices, such as LDAP and RADIUS servers.
Medium	Mid-level notifications, such as antivirus package upgrades.
Low	Minor severity notifications, such as user password changes.
Informational	Log in/log off, administrator name or password change, any configuration change, and all other events not covered by the other severity levels.

Log Forwarding Profiles


Log forwarding profiles allow you to forward traffic and threat logs to Panorama or an external system. A log forwarding profile can be added to a security zone to forward zone protection logs or to a security policy to forward logs for traffic that matches that policy. You can also configure a default log forwarding profile—the settings in the default profile will be used as the default log forwarding settings for new security zones and new security policies. This allows you to consistently include your organization’s preferred log forwarding settings in new policies and zones automatically, without administrators having to manually add them each time.

The following sections show how to create a log forwarding profile and how to enable a profile to be used as the default log forwarding settings for new security policies or security zones:

- ▲ [Create a Log Forwarding Profile](#)
- ▲ [Set Up or Override a Default Log Forwarding Profile](#)

Create a Log Forwarding Profile

Create a log forwarding profile that can be added to security policies and security zones, in order to forward traffic and threat logs to Panorama or an external system. Forwarded logs can be sent as SNMP traps, syslog messages, or email notifications.

Enable a Log Forwarding Profile	
<p>Step 1 Add a log forwarding profile.</p>	<ol style="list-style-type: none"> 1. Select Objects > Log Forwarding Profile and Add a new security profile group. 2. Give the profile group a descriptive Name to help identify it when adding the profile to security policies or security zones. 3. If the firewall is in Multiple Virtual System Mode, enable the profile to be Shared by all virtual systems. 4. Add settings for the Traffic logs, Threat logs, and WildFire logs: <ul style="list-style-type: none"> • Select the Panorama check box for the severity of the Traffic, Threat, or WildFire logs that you want to be forwarded to Panorama. • Specify logs that you want to forward to additional destinations: SNMP Trap destinations, Email servers, or Syslog servers. 5. Click OK to save the log forwarding profile.
<p>Step 1 Add the log forwarding profile to a security policy.</p> <p>Traffic that matches the security policy and your log forwarding profile will be forwarded to the destinations defined in your profile.</p> <p>For more details on security policies, see Security Policy.</p>	<ol style="list-style-type: none"> 1. Select Policies > Security and Add or modify a security policy. 2. Select Actions and the log forwarding profile you created in the Log Forwarding Profile drop-down. <div data-bbox="735 1108 792 1171" style="display: inline-block; vertical-align: middle;">  </div> Threat log entries are generated according to the security profiles you have configured, in addition to the settings defined in the log forwarding profile. For more details on security profiles, see Security Profiles. 3. Click OK to save the security policy.
<p>Step 1 Add the log forwarding profile to a security zone.</p> <p>For more details on setting up security zones, see Configure Interfaces and Zones.</p>	<ol style="list-style-type: none"> 1. Select Network > Zones and Add or modify a security zone. 2. Select the log forwarding profile from the Log Setting drop-down. 3. Click OK to save the security zone.
<p>Step 1 Save your changes.</p>	<p>Commit.</p>

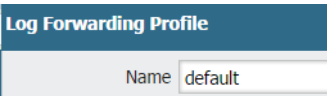
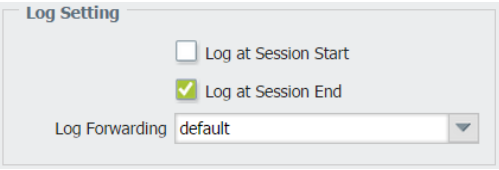
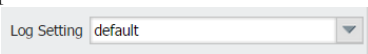
Set Up or Override a Default Log Forwarding Profile

Add a new log forwarding profile or modify an existing one to be used as the default log forwarding settings for new security policy rules or new security zones. When an administrator creates a new security policy or a new security zone, the *default* log forwarding profile will be automatically selected as the policy or zone’s log forwarding settings (the administrator can choose to manually select different log forwarding settings if desired). Use the following options to set up a default log forwarding profile or to override your default settings.



If no default security profile exists, the profile settings for a new security policy are set to **None** by default.

Set Up or Override a Default Log Forwarding Profile

<ul style="list-style-type: none"> Set up a default log forwarding profile. 	<ol style="list-style-type: none"> Select Objects > Log Forwarding and Add a new log forwarding profile or modify an existing profile. Name the security profile group <i>default</i>:  Click OK and Commit. Confirm that the <i>default</i> log forwarding profile is included in new security policies by default: <ol style="list-style-type: none"> Select Policies > Security and Add a new security policy. Select the Actions tab and ensure the Log Forwarding field shows the <i>default</i> profile selected:  Confirm the <i>default</i> log forwarding profile is included in new security zones by default: <ol style="list-style-type: none"> Select Network > Zones and Add a new security zone. Ensure the Log Setting field shows the <i>default</i> log forwarding profile selected: 
<ul style="list-style-type: none"> Override a default log forwarding profile. 	<p>If you have an existing default log forwarding profile, and you do not want the log forwarding settings defined in that profile to be applied to a new security policy or a new security zone, continue to modify the Log Setting field in the policy or zone according to your preference.</p>

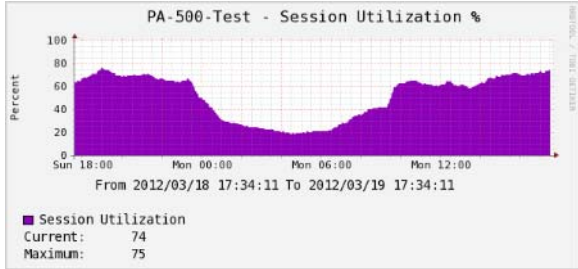
Monitor the Firewall Using SNMP

All Palo Alto Networks firewalls support standard networking SNMP management information base (MIB) modules as well as proprietary Enterprise MIB modules. You can configure an SNMP manager to get statistics from the firewall. For example, you could configure your SNMP manager to monitor the interfaces, active sessions, concurrent sessions, session utilization percentage, temperature, and/or system uptime on the firewall.



Palo Alto Networks firewalls support SNMP GET requests only; SNMP SET requests are not supported.

Set Up SNMP Monitoring	
<p>Step 1 Enable the interface to allow inbound SNMP requests.</p>	<ul style="list-style-type: none"> • If you will be receiving SNMP GET messages on the MGT interface, select Device > Setup > Management and click to Edit the Management Interface Settings section of the screen. In the Services section, select the SNMP check box and then click OK. • If you will be receiving SNMP GET messages on a different interface, you must associate a management profile with the interface and enable SNMP management.
<p>Step 2 From the web interface on the firewall, configure the settings to allow the SNMP agent on the firewall to respond to incoming GET requests from the SNMP manager.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Operations > SNMP Setup. 2. Specify the Physical Location of the firewall and the name or email address of an administrative Contact. 3. Select the SNMP Version and then enter the configuration details as follows (depending on which SNMP version you are using) and then click OK: <ul style="list-style-type: none"> • V2c—Enter the SNMP Community String that will allow the SNMP manager access to the SNMP agent on the firewall. The default value is public, however because this is a well-known community string, it is a best practice to use a value that is not easily guessed. • V3—You must create at least one View and one User in order to use SNMPv3. The view specifies which management information the manager has access to. If you want to allow access to all management information, just enter the top-level OID of .1.3.6.1 and specify the Option as include (you can also create views that exclude certain objects). Use 0xf0 as the Mask. Then when you create a user, select the View you just created and specify the Auth Password and Priv Password. <p>The authentication settings (the community string for V2c or the username and passwords for V3) configured on the firewall must match the value configured on the SNMP manager.</p> 4. Click OK to save the settings. 5. Click Commit to save the SNMP settings.
<p>Step 3 Enable the SNMP manager to interpret firewall statistics.</p>	<p>Load the PAN-OS MIB files into your SNMP management software and, if necessary, compile them. Refer to the documentation for your SNMP manager for specific instructions on how to do this.</p>

Set Up SNMP Monitoring (Continued)	
<p>Step 4 Identify the statistics you want to monitor.</p>	<p>Using a MIB browser, walk the PAN-OS MIB files to identify the object identifiers (OIDs) that correspond to the statistics you want to monitor. For example, suppose you want to monitor Session Utilization Percentage on the firewall. Using a MIB browser you will see that this statistic corresponds to OID 1.3.6.1.4.1.25461.2.1.2.3.1.0 in the PAN-COMMON-MIB.</p>
<p>Step 5 Configure the SNMP management software to monitor the OIDs you are interested in.</p>	<p>Refer to the documentation for your SNMP manager for specific instructions on how to do this.</p>
<p>Step 6 After you complete the configuration on both the firewall and the SNMP manager, you can begin monitoring the firewall from your SNMP management software.</p>	<p>The following is an example of how an SNMP manager displays real-time session utilization percentage statistics for a monitored PA-500 firewall:</p> 

Monitor the Firewall Using NetFlow

NetFlow is an industry-standard protocol that enables the firewall to record statistics on the IP traffic that traverses its interfaces. The firewall exports the statistics as NetFlow fields to a NetFlow collector. The NetFlow collector is a server you use to analyze network traffic for security, administration, accounting and troubleshooting purposes. All firewalls support NetFlow Version 9 except the PA-4000 Series and PA-7050 firewalls. The firewalls support only unidirectional NetFlow, not bidirectional. 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 interfaces, you can export records for the aggregate group but not for individual interfaces within the group. To identify firewall interfaces in a NetFlow collector, see [Identify Firewall Interfaces in External Monitoring Systems](#).

The firewall supports standard and enterprise (PAN-OS specific) [NetFlow Templates](#), which NetFlow collectors use to decipher the exported fields. The firewall selects a template based on the type of data it exports: IPv4 or IPv6 traffic, with or without NAT, and with standard or enterprise-specific fields. The firewall periodically refreshes templates to 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. You configure the refresh frequency according to the requirements of the particular NetFlow collector you use.

Monitor the Firewall Using NetFlow	
<p>Step 1 Create a NetFlow server profile.</p> <p>The profile defines which NetFlow collectors will receive the exported records and specifies export parameters.</p>	<ol style="list-style-type: none"> 1. Select Device > Server Profiles > NetFlow and click Add. 2. Enter a Name for 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. For the Active Timeout, specify the frequency in minutes at which the firewall exports records (default is 5). 5. Select the PAN-OS Field Types check box if you want the firewall to export App-ID and User-ID fields. 6. For each NetFlow collector (up to two per profile) that will receive fields, click Add and enter an identifying server Name, hostname or IP address (NetFlow Server), and access Port (default 2055). 7. Click OK to save the profile.
<p>Step 2 Assign the NetFlow server profile to the interfaces that carry the traffic you want to analyze.</p> <p>The interfaces can be any type except HA, log card, or decrypt mirror. In this example, you assign the profile to an existing Ethernet interface.</p>	<ol style="list-style-type: none"> 1. Select Network > Interfaces > Ethernet and click an Interface name to edit it. 2. In the NetFlow Profile drop-down, select the NetFlow server profile and click OK. 3. Click Commit and OK.

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
IPv6 Standard	258
IPv6 Enterprise	259
IPv4 with NAT Standard	260
IPv4 with NAT Enterprise	261

The following table lists the NetFlow fields that the firewall can send, along with the templates that define them:

Value	Field	Description	Templates
1	IN_BYTES	Incoming counter with length $N * 8$ bits for the number of bytes associated with an IP flow. By default, N is 4.	All templates
2	IN_PKTS	Incoming counter with length $N * 8$ bits for the number of packets associated with an IP 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 incoming interface.	All templates
6	TCP_FLAGS	Total of all the TCP flags in this flow.	All templates
7	L4_SRC_PORT	TCP/UDP source port number (for example, FTP, Telnet, or equivalent).	All templates
8	IPV4_SRC_ADDR	IPv4 source address.	IPv4 standard IPv4 enterprise IPv4 with NAT standard IPv4 with NAT enterprise

Value	Field	Description	Templates
10	INPUT_SNMP	Input interface index. The value length is 2 bytes by default, but higher values are possible. For details on how Palo Alto Networks firewalls generate interface indexes, see Identify Firewall Interfaces in External Monitoring Systems .	All templates
11	L4_DST_PORT	TCP/UDP destination port number (for example, FTP, Telnet, or equivalent).	All templates
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 Identify Firewall Interfaces in External Monitoring Systems .	All templates
21	LAST_SWITCHED	System uptime in milliseconds when the last packet of this flow was switched.	All templates
22	FIRST_SWITCHED	System uptime in milliseconds when the first packet of this flow was switched.	All templates
27	IPV6_SRC_ADDR	IPv6 source address.	IPv6 standard IPv6 enterprise
28	IPV6_DST_ADDR	IPv6 destination address.	IPv6 standard IPv6 enterprise
32	ICMP_TYPE	Internet Control Message Protocol (ICMP) packet type. This is reported as: ICMP Type * 256 + ICMP code	All templates
61	DIRECTION	Flow direction: <ul style="list-style-type: none"> • 0 = ingress • 1 = egress 	All templates
148	flowId	An identifier of a flow that is unique within an <i>observation domain</i> . You can use this <i>information element</i> to distinguish between different flows if <i>flow keys</i> 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

Value	Field	Description	Templates
233	firewallEvent	Indicates a firewall event: <ul style="list-style-type: none"> • 0 = Ignore (invalid) • 1 = Flow created • 2 = Flow deleted • 3 = Flow denied • 4 = Flow alert • 5 = Flow update (the session state changed from active to deny) 	All templates
225	postNATSourceIPv4Address	The definition of this information element is identical to that of sourceIPv4Address, except that it reports a modified value that a NAT middlebox function caused after the packet passed the <i>observation point</i> .	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 a NAT middlebox function caused after the packet passed the observation point.	IPv4 with NAT standard IPv4 with NAT enterprise
227	postNAPTSourceTransportPort	The definition of this information element is identical to that of sourceTransportPort, except that it reports a modified value that a Network Address Port Translation (NAPT) middlebox function caused after the packet passed the observation point.	IPv4 with NAT standard IPv4 with NAT enterprise
228	postNAPTDestinationTransportPort	The definition of this information element is identical to that of destinationTransportPort, except that it reports a modified value that a Network Address Port Translation (NAPT) middlebox function caused after the packet passed the observation point.	IPv4 with NAT standard IPv4 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
56701	App-ID	The name of an application that App-ID identified. The name can be up to 32 bytes.	IPv4 enterprise IPv4 with NAT enterprise IPv6 enterprise
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

Identify Firewall Interfaces in External Monitoring Systems



When you use a NetFlow collector (see [Monitor the Firewall Using NetFlow](#)) or SNMP manager (see [Monitor the Firewall Using SNMP](#)) to monitor traffic flows, an interface index (SNMP ifindex object) identifies the firewall interface that carried a particular flow. The formula that the Palo Alto Networks firewall uses to calculate interface indexes varies by platform and whether the interface is physical or logical.



You cannot use SNMP to monitor logical interfaces, only physical interfaces. You can use NetFlow to monitor logical or physical interfaces.

Most NetFlow collectors use SNMP to determine the name of a physical interface based on the SNMP interface index.

Physical interface indexes have a range of 1-9999, which the firewall calculates as follows:

Firewall Platform	Calculation	Example Interface Index
Non-chassis based: VM-Series, PA-200, PA-500, PA-2000 Series, PA-3000 Series, PA-4000 Series, PA-5000 Series  The PA-4000 Series platform supports SNMP but not NetFlow.	MGT port + physical port offset <ul style="list-style-type: none"> MGT port—This is a constant that depends on the platform: <ul style="list-style-type: none"> 2 for hardware-based firewalls (for example, the PA-5000 Series firewall) 1 for the VM-Series firewall Physical port offset—This is the physical port number. 	PA-5000 Series firewall, Eth1/4 = $2 \text{ (MGT port)} + 4 \text{ (physical port)} = 6$
Chassis based: PA-7050 firewalls  This platform supports SNMP but not NetFlow.	(Max. ports * slot) + physical port offset + MGT port <ul style="list-style-type: none"> Maximum ports—This is a constant of 64. Slot—This is the chassis slot number of the network interface card. Physical port offset—This is the physical port number. MGT port—This is a constant of 5 for PA-7050 firewalls. 	PA-7050 firewall, Eth3/9 = $[64 \text{ (max. ports)} * 3 \text{ (slot)}] + 9 \text{ (physical port)} + 5 \text{ (MGT port)} = 206$

Logical interface indexes for all platforms are nine-digit numbers that the firewall calculates as follows:

Interface Type	Range	Digit 9	Digits 7-8	Digits 5-6	Digits 1-4	Example Interface Index
Layer 3 subinterface	101010001-199999999	Type: 1	Interface slot: 1-9 (01-09)	Interface port: 1-9 (01-09)	Subinterface: suffix 1-9999 (0001-9999)	Eth1/5.22 = 100000000 (type) + 100000 (slot) + 50000 (port) + 22 (suffix) = 101050022
Layer 2 subinterface	101010001-199999999	Type: 1	Interface slot: 1-9 (01-09)	Interface port: 1-9 (01-09)	Subinterface: suffix 1-9999 (0001-9999)	Eth2/3.6 = 100000000 (type) + 200000 (slot) + 30000 (port) + 6 (suffix) = 102030006

Interface Type	Range	Digit 9	Digits 7-8	Digits 5-6	Digits 1-4	Example Interface Index
Vwire subinterface	101010001-199999999	Type: 1	Interface slot: 1-9 (01-09)	Interface port: 1-9 (01-09)	Subinterface: suffix 1-9999 (0001-9999)	Eth4/2.312 = 100000000 (type) + 400000 (slot) + 20000 (port) + 312 (suffix) = 104020312
VLAN	200000001-200009999	Type: 2	00	00	VLAN suffix: 1-9999 (0001-9999)	VLAN.55 = 200000000 (type) + 55 (suffix) = 200000055
Loopback	300000001-300009999	Type: 3	00	00	Loopback suffix: 1-9999 (0001-9999)	Loopback.55 = 300000000 (type) + 55 (suffix) = 300000055
Tunnel	400000001-400009999	Type: 4	00	00	Tunnel suffix: 1-9999 (0001-9999)	Tunnel.55 = 400000000 (type) + 55 (suffix) = 400000055
Aggregate group	500010001-500089999	Type: 5	00	AE suffix: 1-8 (01-08)	Subinterface: suffix 1-9999 (0001-9999)	AE5.99 = 500000000 (type) + 50000 (AE Suffix) + 99 (suffix) = 500050099

Manage Reporting

The reporting capabilities on the firewall allow you to keep a pulse on your network, validate your policies, and focus your efforts on maintaining network security for keeping your users safe and productive.

- ▲ [About Reports](#)
- ▲ [View Reports](#)
- ▲ [Disable Predefined Reports](#)
- ▲ [Generate Custom Reports](#)
- ▲ [Generate Botnet Reports](#)
- ▲ [Manage PDF Summary Reports](#)
- ▲ [Generate User/Group Activity Reports](#)
- ▲ [Manage Report Groups](#)
- ▲ [Schedule Reports for Email Delivery](#)

About Reports

The firewall includes predefined reports that you can use as-is, or you can build custom reports that meet your needs for specific data and actionable tasks, or you can combine predefined and custom reports to compile information you need. The firewall provides the following types of reports:

- **Predefined Reports**—Allow you to view a quick summary of the traffic on your network. A suite of predefined reports are available in four categories—Applications, Traffic, Threat, and URL Filtering. See [View Reports](#).
- **User or Group Activity Reports**—Allow you to schedule or create an on-demand report on the application use and URL activity for a specific user or for a user group. The report includes the URL categories and an estimated browse time calculation for individual users. See [Generate User/Group Activity Reports](#).
- **Custom Reports**—Create and schedule custom reports that show exactly the information you want to see by filtering on conditions and columns to include. You can also include query builders for more specific drill down on report data. See [Generate Custom Reports](#).
- **PDF Summary Reports**—Aggregate up to 18 predefined or custom reports/graphs from Threat, Application, Trend, Traffic, and URL Filtering categories into one PDF document. See [Manage PDF Summary Reports](#).
- **Botnet Reports**—Allow you to use behavior-based mechanisms to identify potential botnet-infected hosts in the network. See [Interpret Botnet Report Output](#).
- **Report Groups**—Combine custom and predefined reports into report groups and compile a single PDF that is emailed to one or more recipients. See [Manage Report Groups](#).

Reports can be generated on demand, on a recurring schedule, and can be scheduled for email delivery.

View Reports

The firewall provides an assortment of over 40 predefined reports that are generated every day; these reports can be viewed directly on the firewall. In addition to these reports, you can view *scheduled* custom reports and summary reports.

About 200 MB of storage is allocated for saving reports on the firewall. This storage space is not user-configurable, and older reports are purged to store recent reports. Therefore, for long-term retention of reports, you can either export the reports or schedule the reports for email delivery. To disable selected reports and conserve system resources on the firewall, see [Disable Predefined Reports](#).



User/group activity reports must be generated on demand or scheduled for email delivery. Unlike the other reports, these reports cannot be saved on the firewall.

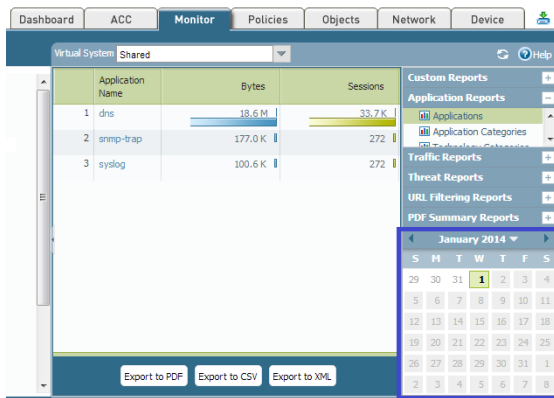
View Reports

Step 1 Select **Monitor > Reports**.

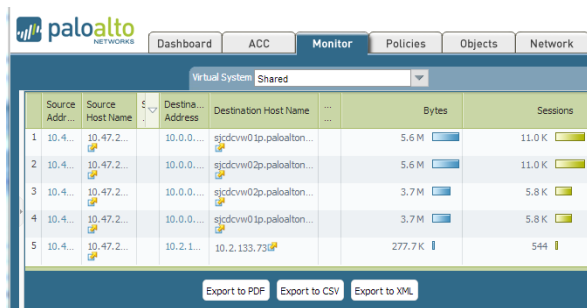
The reports are chunked into sections on the right-hand side of the window: **Custom Reports**, **Application Reports**, **Traffic Reports**, **Threat Reports**, **URL Filtering Reports**, and **PDF Summary Reports**.

Step 2 Select a report to view. When you select a report, the previous day's report is displayed on screen.

To view reports for any of the previous days, select an available date from the calendar at the bottom of the page and select a report within the same section. If you change sections, the time selection is reset.



Step 3 To view a report offline, you can export the report to PDF, CSV or to XML formats. Click **Export to PDF**, **Export to CSV**, or **Export to XML** at the bottom of the page. Then print or save the file.



Disable Predefined Reports

The firewall includes about 40 predefined reports that are automatically generated each day. If you do not use some or all of these predefined reports, you can disable selected reports and conserve system resources on the firewall.

Before disabling one or more predefined reports, make sure that the report is not included in a Group Report or a PDF Report. If the disabled predefined report is included in a group or PDF report, the Group/PDF report will be rendered without any data.

Disable Predefined Reports

1. Select **Device > Setup > Management** on the firewall.
2. Click the Edit icon in the Logging and Reporting Settings section and select the **Log Export and Reporting** tab.
3. To disable reports:
 - Clear the check box corresponding to each report that you want to disable.
 - Select **Deselect All** to disable all predefined reports.

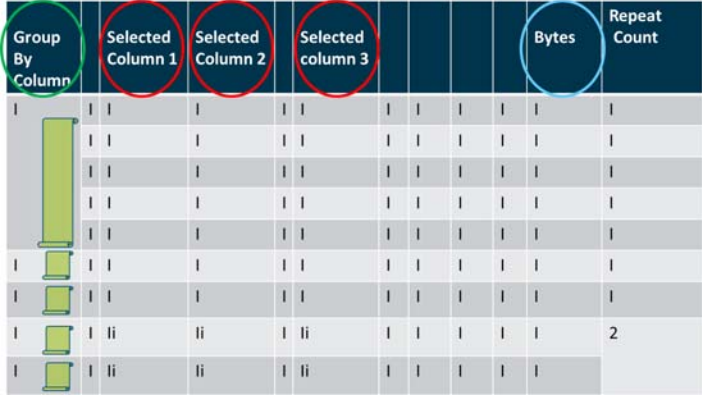


4. Click **OK**, and **Commit** the changes.

Generate Custom Reports

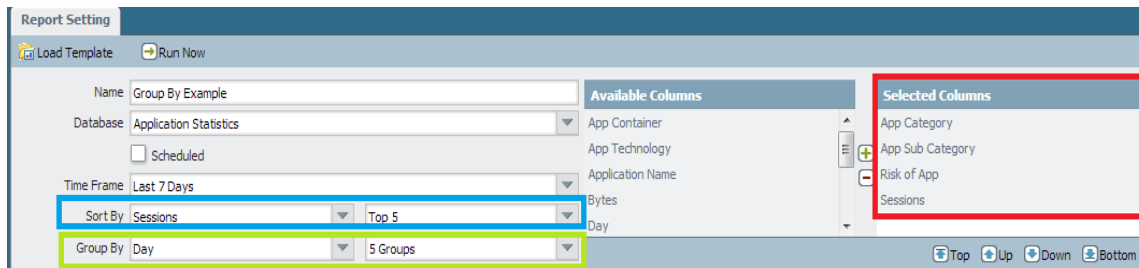
In order to create purposeful custom reports, you must consider the attributes or key pieces of information that you want to retrieve and analyze. This consideration guides you in making the following selections in a custom report:

Selection	Description
Data Source	<p>The data file that is used to generate the report. The firewall offers two types of data sources—Summary databases and Detailed logs.</p> <ul style="list-style-type: none"> • Summary databases are available for traffic, threat, and application statistics. The firewall aggregates the detailed logs on traffic, application, and threat at 15-minute intervals. The data is condensed—duplicate sessions are grouped together and incremented with a repeat counter, and some attributes (or columns) are not included in the summary—to allow faster response time when generating reports. • Detailed logs are itemized and are a complete listing of all the attributes (or columns) that pertain to the log entry. Reports based on detailed logs take much longer to run and are not recommended unless absolutely necessary.
Attributes	<p>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).</p>
Sort By/ Group By	<p>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.</p> <p>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.</p> <p>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.</p>

Selection	Description
	<p>The following example illustrates how the Selected Columns and Sort By/Group By criteria work together when generating reports:</p>  <p>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.</p> <p>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.</p> <p>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.</p>

Selection	Description
-----------	-------------

For example, if a report has the following selections



The output will display as follows:

Day Received	App Category	App Sub Category	Risk	Sessions
1 Wed, Jan 8, 2014	networking	infrastructure	0.1	33.3K
	general-internet	internet-utility	0.4	3.5K
	networking	encrypted-tunnel	0.4	1.5K
	media	audio-streaming	0.3	265
	business-systems	management	0.2	262
2 Thu, Jan 2, 2014	networking	infrastructure	0.1	35.8K
	networking	encrypted-tunnel	0.4	1.3K
	general-internet	internet-utility	0.4	430
	media	audio-streaming	0.3	347
	business-systems	management	0.2	330
3 Tue, Jan 7, 2014	networking	infrastructure	0.1	19.6K
	networking	encrypted-tunnel	0.4	1.1K
	general-internet	internet-utility	0.4	411
	business-systems	management	0.2	153
	networking	infrastructure	0.3	151
4 Fri, Jan 3, 2014	networking	infrastructure	0.1	20.4K
	networking	encrypted-tunnel	0.4	210
	business-systems	management	0.2	175
	networking	infrastructure	0.3	166
	general-internet	internet-utility	0.4	98
5 Wed, Jan 1, 2014	networking	infrastructure	0.1	3.1K
	business-systems	management	0.2	23
	networking	infrastructure	0.3	23

The report is anchored by **Day** and sorted by **Sessions**. It lists the 5 days (**5 Groups**) with maximum traffic in the **Last 7 Days** time frame. The data is enumerated by the **Top 5** sessions for each day for the selected columns—**App Category, App Subcategory** and **Risk**.

Time Period	The date range for which you want to analyze data. You can define a custom range or select a time period ranging from last 15 minutes to the last 30 days. The reports can be run on demand or scheduled to run at a daily or weekly cadence.
Query Builder	The query builder allows you to define specific queries to further refine the selected attributes. It allows you see just what you want in your report using and and or operators and a match criteria, and then include or exclude data that matches or negates the query in the report. Queries enable you to generate a more focused collation of information in a report.

Generate Custom Reports


1. Select **Monitor > Manage Custom Reports**.

2. Click **Add** and then enter a **Name** for the report.



To base a report on an predefined template, click **Load Template** and choose the template. You can then edit the template and save it as a custom report.

3. Select the database to use for the report.



Each time you create a custom report, a **Log View** report is automatically created. This report show the logs that were used to build the custom report. The log view report uses the same name as the custom report, but appends the phrase (Log View) to the report name.

When creating a report group, you can include the log view report with the custom report. For more information, see [Manage Report Groups](#).


4. Select the **Scheduled** check box to run the report each night. The report is then available for viewing in the **Reports** column on the side.

5. Define the filtering criteria. Select the **Time Frame**, the **Sort By** order, **Group By** preference, and select the columns that must display in the report.

6. (Optional) Select the **Query Builder** attributes if you want to further refine the selection criteria. To build a report query, specify the following and click **Add**. Repeat as needed to construct the full query.

- **Connector**—Choose the connector (and/or) to precede the expression you are adding.
- **Negate**—Select the check box to interpret the query as a negation. If, for example, you choose to match entries in the last 24 hours and/or are originating from the untrust zone, the negate option causes a match on entries that are not in the past 24 hours and/or are not from the untrust zone.
- **Attribute**—Choose a data element. The available options depend on the choice of database.
- **Operator**—Choose the criterion to determine whether the attribute applies (such as =). The available options depend on the choice of database.
- **Value**—Specify the attribute value to match.

For example, the following figure (based on the Traffic Log database) shows a query that matches if the traffic log entry was received in the past 24 hours and is from the “untrust” zone.



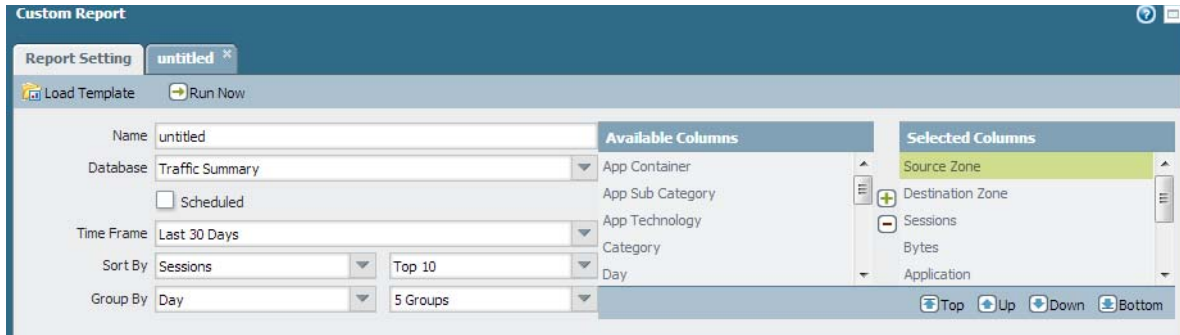
7. To test the report settings, select **Run Now**. Modify the settings as required to change the information that is displayed in the report.

8. Click **OK** to save the custom report.

Generate Custom Reports (Continued)

Examples of Custom Reports

If you want to set up a simple report in which you use the traffic summary database from the last 30 days, and sort the data by the top 10 sessions and these sessions are grouped into 5 groups by day of the week. You would set up the custom report to look like this:



And the PDF output for the report would look as follows:

Sample Report

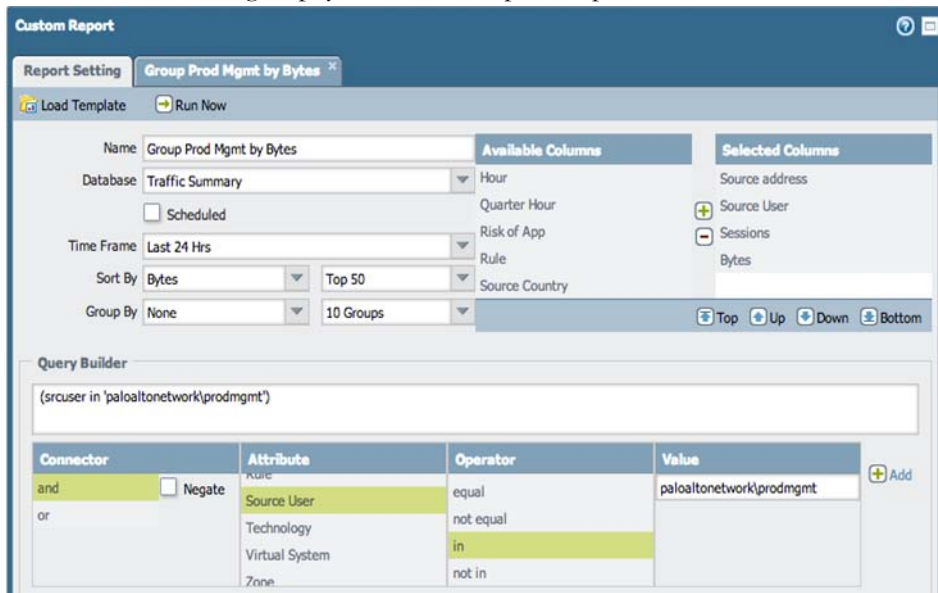
Group By
Column

Sort By
Column

pm-firewall : 2013/11/20 16:07:33 - 2013/12/20 16:07:32											
Day Received	Source Zone	Destination Zone	Application	Source address	Source Host Name	Destination address	Destination Host Name	Risk	App Category	Sessions	Bytes
Mon, Dec 16, 2013	trust	untrust	dns	10.47.20.20	10.47.20.20	10.0.0.247	qjstcve02p.paloaltonetworks.local	4	networking	11,264	6.41 M
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.0.0.246	qjstcve01p.paloaltonetworks.local	4	networking	11,264	6.40 M
	trust	untrust	dns	10.47.20.3	10.47.20.3	10.0.0.247	qjstcve02p.paloaltonetworks.local	4	networking	5,834	3.44 M
	trust	untrust	dns	10.47.20.3	10.47.20.3	10.0.0.246	qjstcve01p.paloaltonetworks.local	4	networking	5,834	3.43 M
	trust	untrust	dns	10.47.20.5	10.47.20.5	10.43.2.10	qjbt1dcve01p.paloaltonetworks.local	4	networking	1,764	414.15 k
	trust	untrust	sql	10.47.20.5	10.47.20.5	10.101.2.90	hoxevone.paloaltonetworks.local	4	networking	1,024	49.03 M
	trust	untrust	syslog	10.47.20.20	10.47.20.20	10.2.133.73	10.2.133.73	2	business-systems	280	103.22 k
	trust	untrust	snmp-trap	10.47.20.20	10.47.20.20	10.2.133.73	10.2.133.73	3	networking	260	101.84 k
	trust	untrust	dns	10.47.20.5	10.47.20.5	10.44.2.10	qjbt2dcve01p.paloaltonetworks.local	4	networking	196	43.11 k
	trust	untrust	smtp	10.47.20.5	10.47.20.5	10.101.101.234	10.101.101.234	4	general-internet	170	80.07 k
Sun, Dec 15, 2013	trust	untrust	dns	10.47.20.20	10.47.20.20	10.0.0.246	qjstcve01p.paloaltonetworks.local	4	networking	11,044	6.59 M
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.0.0.247	qjstcve02p.paloaltonetworks.local	4	networking	11,044	6.58 M
	trust	untrust	dns	10.47.20.3	10.47.20.3	10.0.0.247	qjstcve02p.paloaltonetworks.local	4	networking	5,834	3.87 M
	trust	untrust	dns	10.47.20.3	10.47.20.3	10.0.0.246	qjstcve01p.paloaltonetworks.local	4	networking	5,824	3.86 M
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.43.2.10	qjbt1dcve01p.paloaltonetworks.local	4	networking	1,964	471.26 k
	trust	untrust	sql	10.47.20.20	10.47.20.20	10.101.2.90	hoxevone.paloaltonetworks.local	4	networking	418	12.02 M
	trust	untrust	syslog	10.47.20.20	10.47.20.20	10.2.133.73	10.2.133.73	2	business-systems	278	104.10 k
	trust	untrust	snmp-trap	10.47.20.20	10.47.20.20	10.2.133.73	10.2.133.73	3	networking	278	183.38 k
	trust	untrust	sql	10.47.20.20	10.47.20.20	23.72.38.120	a23-72-38-120.deploy.static.akamaitechnologies.com	4	networking	101	937.11 k
Thu, Dec 19, 2013	trust	untrust	web-browsing	10.47.20.20	10.47.20.20	213.27.164.28	213.27.164.28	4	general-internet	90	515.49 k
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.0.0.246	qjstcve01p.paloaltonetworks.local	4	networking	11,354	5.95 M
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.0.0.247	qjstcve02p.paloaltonetworks.local	4	networking	11,344	6.48 M
	trust	untrust	dns	10.47.20.3	10.47.20.3	10.0.0.246	qjstcve01p.paloaltonetworks.local	4	networking	5,834	3.51 M
	trust	untrust	dns	10.47.20.3	10.47.20.3	10.0.0.247	qjstcve02p.paloaltonetworks.local	4	networking	5,824	3.50 M
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.43.2.10	qjbt1dcve01p.paloaltonetworks.local	4	networking	1,444	348.48 k
	trust	untrust	sql	10.47.20.20	10.47.20.20	10.101.2.90	hoxevone.paloaltonetworks.local	4	networking	482	24.22 M
	trust	untrust	snmp-trap	10.47.20.20	10.47.20.20	10.2.133.73	10.2.133.73	3	networking	280	190.26 k
	trust	untrust	syslog	10.47.20.20	10.47.20.20	10.2.133.73	10.2.133.73	2	business-systems	280	107.89 k
	trust	untrust	sql	10.47.20.20	10.47.20.20	10.47.0.8	pm-fwesd1.paloaltonetworks.local	4	networking	222	3.79 M
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.44.2.10	qjbt2dcve01p.paloaltonetworks.local	4	networking	186	41.49 k
Wed, Dec 18, 2013	trust	untrust	dns	10.47.20.20	10.47.20.20	10.0.0.247	qjstcve02p.paloaltonetworks.local	4	networking	11,334	6.49 M
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.0.0.246	qjstcve01p.paloaltonetworks.local	4	networking	11,324	6.50 M
	trust	untrust	dns	10.47.20.3	10.47.20.3	10.0.0.247	qjstcve02p.paloaltonetworks.local	4	networking	5,834	3.51 M
	trust	untrust	dns	10.47.20.3	10.47.20.3	10.0.0.246	qjstcve01p.paloaltonetworks.local	4	networking	5,824	3.50 M
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.43.2.10	qjbt1dcve01p.paloaltonetworks.local	4	networking	1,164	276.79 k
	trust	untrust	sql	10.47.20.20	10.47.20.20	10.101.2.90	hoxevone.paloaltonetworks.local	4	networking	496	64.20 M
	trust	untrust	snmp-trap	10.47.20.20	10.47.20.20	10.2.133.73	10.2.133.73	3	networking	269	190.63 k
	trust	untrust	syslog	10.47.20.20	10.47.20.20	10.2.133.73	10.2.133.73	2	business-systems	289	108.07 k
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.44.2.10	qjbt2dcve01p.paloaltonetworks.local	4	networking	228	64.16 k
Thu, Dec 12, 2013	trust	untrust	web-browsing	10.47.20.20	10.47.20.20	213.27.164.28	213.27.164.28	4	general-internet	62	4.27 M
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.0.0.246	qjstcve01p.paloaltonetworks.local	4	networking	11,404	5.93 M
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.0.0.247	qjstcve02p.paloaltonetworks.local	4	networking	11,394	6.51 M
	trust	untrust	dns	10.47.20.3	10.47.20.3	10.0.0.246	qjstcve01p.paloaltonetworks.local	4	networking	3,894	2.19 M
	trust	untrust	dns	10.47.20.3	10.47.20.3	10.0.0.247	qjstcve02p.paloaltonetworks.local	4	networking	3,884	2.19 M
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.43.2.10	qjbt1dcve01p.paloaltonetworks.local	4	networking	2,354	603.08 k
	trust	untrust	sql	10.47.20.20	10.47.20.20	10.101.2.90	hoxevone.paloaltonetworks.local	4	networking	649	64.13 M
	trust	untrust	dns	10.47.20.20	10.47.20.20	10.44.2.10	qjbt2dcve01p.paloaltonetworks.local	4	networking	417	90.68 k
	trust	untrust	snmp-trap	10.47.20.20	10.47.20.20	10.2.133.73	10.2.133.73	3	networking	292	196.16 k
	trust	untrust	syslog	10.47.20.20	10.47.20.20	10.2.133.73	10.2.133.73	2	business-systems	292	110.95 k
	trust	untrust	sql	10.47.20.20	10.47.20.20	10.101.2.49	10.101.2.49	4	networking	130	6.88 M

Generate Custom Reports (Continued)

Now, if you want to use the query builder to generate a custom report that represents the top consumers of network resources within a user group, you would set up the report to look like this:



The report would display the top users in the product management user group sorted by bytes, as follows:

	Source Address	Source Host Name	Source User	Sessions	Bytes
1	10.1.1.16.35	10.1.1.16.35	paloaltonetwork\p_acc...	136	445.3 K
2	10.1.1.35.48	10.1.1.35.48	paloaltonetwork\p...	123	410.5 K
3	10.1.1.16.49	10.1.1.16.49	paloaltonetwork\p...	103	360.4 K
4	10.1.1.13.23	10.1.1.13.23	paloaltonetwork\p...	103	347.2 K
5	10.1.1.14.51	10.1.1.14.51	paloaltonetwork\p...	103	334.9 K
6	10.1.1.11.55	10.1.1.11.55	paloaltonetwork\p...	96	326.3 K
7	10.1.1.11.161	10.1.1.11.161	paloaltonetwork\j...	84	306.7 K
8	10.1.1.16.42	panv...h.paloaltonetwork...	paloaltonetwork\c...	49	232.2 K
9	10.1.1.14.145	panv...h.paloaltonetwork...	paloaltonetwork\l...	65	201.2 K
10	10.1.1.14.42	10.1.1.14.42	paloaltonetwork\p...	54	191.4 K
11	10.1.1.14.145	pa...paloaltonetworks.lo...	paloaltonetwork\p...	28	171.2 K
12	10.1.1.17.26	10.1.1.17.26	paloaltonetwork\p...	32	110.0 K
13	10.1.1.17.74	10.1.1.17.74	paloaltonetwork\p...	13	98.6 K

Generate Botnet Reports

The botnet report enables you to use heuristic and behavior-based mechanisms to identify potential malware- or botnet-infected hosts in your network. To evaluate botnet activity and infected hosts, the firewall correlates user and network activity data in Threat, URL, and Data Filtering logs with the list of malware URLs in PAN-DB, known dynamic DNS domain providers, and domains registered within the last 30 days. You can configure the report to identify hosts that visited those sites, as well as hosts that communicated with Internet Relay Chat (IRC) servers or that used unknown applications. Malware often use dynamic DNS to avoid IP blacklisting, while IRC servers often use bots for automated functions.



The firewall requires Threat Prevention and URL Filtering licenses to use the botnet report.

- ▲ [Configure a Botnet Report](#)
- ▲ [Interpret Botnet Report Output](#)

Configure a Botnet Report

You can schedule a botnet report or run it on demand. The firewall generates scheduled botnet reports every 24 hours because behavior-based detection requires correlating traffic across multiple logs over that timeframe.

Configure a Botnet Report	
<p>Step 1 Define the types of traffic that indicate possible botnet activity.</p>	<ol style="list-style-type: none"> 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.

Configure a Botnet Report (Continued)

<p>Step 2 Schedule the report or run it on demand.</p>	<ol style="list-style-type: none"> 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, you can add not (addr.src in 10.0.1.35) as a query to exclude that host from the report output. For details, see Interpret Botnet Report Output. 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

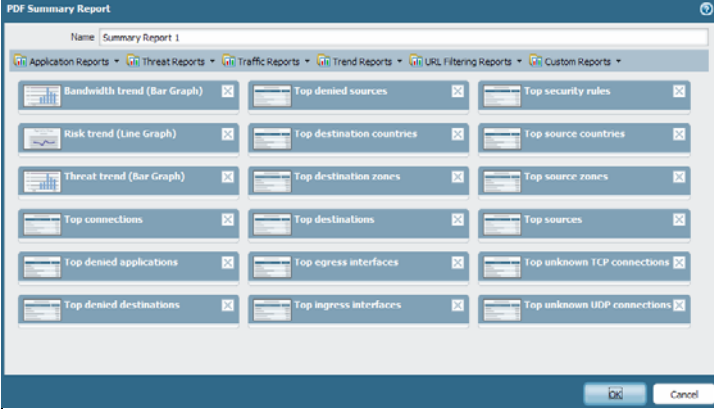
The botnet report displays a line for each host that is associated with traffic you defined as suspicious when configuring the report. For each host, the report displays a confidence score of 1 to 5 to indicate the likelihood of botnet infection, where 5 indicates the highest likelihood. The scores correspond 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](#).

Manage PDF Summary Reports

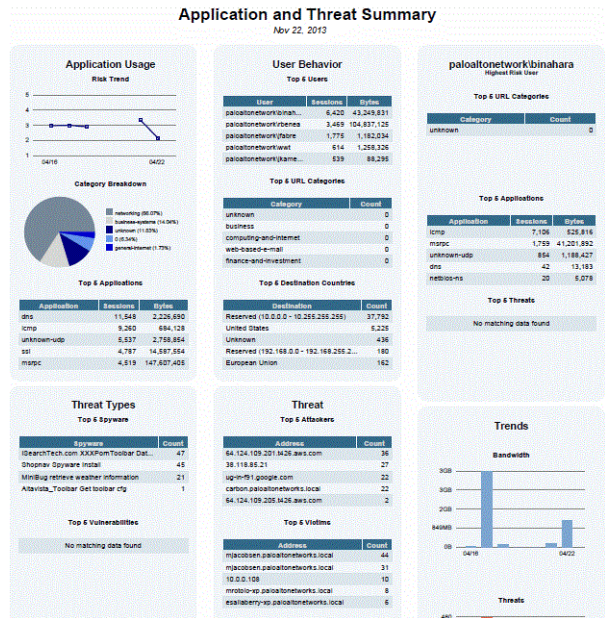
PDF summary reports contain information compiled from existing reports, based on data for the top 5 in each category (instead of top 50). They also contain trend charts that are not available in other reports.

Generate PDF Summary Reports	
<p>Step 1 Set up a PDF Summary Report.</p>	<ol style="list-style-type: none"> 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.  <ol style="list-style-type: none"> 4. Click OK to save the report. 5. Commit the changes. <ul style="list-style-type: none"> • 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 icons to another area of the report.

Generate PDF Summary Reports

Step 2 View the report.

To download and view the PDF Summary Report, see [View Reports](#).



Generate User/Group Activity Reports


User/Group activity reports summarize the web activity of individual users or user groups. Both reports include the same information with a couple exceptions—**Browsing Summary by URL Category** and **Browse time calculations** are included in User Activity Reports, but are not included in Group Activity Reports.

The User Activity Report includes two variables—the **Average Browse Time (sec)** and the **Page Load Threshold (sec)** in the Logging and Reporting sections of the **Device > Setup > Management tab**—when calculating the browse time. The **Average Browse Time** is the average time that the admin estimates that a user will browse a web page. Any request made after the average browse time has elapsed will be considered a new browsing activity. The calculation uses **Container Pages** (logged in the URL filtering logs) as the basis and ignores any new web pages that are loaded between the time of the first request (start time) and the average browse time. Example: If the average browse time setting is 2 minutes and a user opens a web page and views that page for 5 minutes, the browse time for that page will still be 2 minutes. This is done because there is no way to determine how long a user views a given page. The Average Browse Time calculation ignores sites categorized as web advertisements and content delivery networks. The **Page Load Threshold** is the estimated time for page elements to load on the page. Any request that occurs between the first page load and the page load threshold is assumed to be elements of the page. Any requests that occur outside of the page load threshold are assumed to be the user clicking a link within the page.

User-ID must be configured on the firewall in order to access the list of user/user groups.

Generate User/Group Activity Reports

1. Select **Monitor > PDF Reports > User Activity Report**.
 2. Click **Add** and then enter a **Name** for the report.
 3. Create the report:
 - For a User Activity Report: Select **User** and enter the **Username** or **IP address** (IPv4 or IPv6) of the user who will be the subject of the report.
 - For Group Activity Report: Select **Group** and select the **Group Name** for which to retrieve user group information in the report.
 4. Select the time frame for the report from the drop-down.



The number of logs that are analyzed in a user activity report is determined by the number of rows defined on the **Max Rows in User Activity Report** on the Logging and Reporting Settings section on **Device > Setup > Management**.
 5. Select **Include Detailed Browsing** to include detailed URL logs in the report.

The detailed browsing information can include a large volume of logs (thousands of logs) for the selected user or user group and can make the report very large.
 6. To run the report on demand, click **Run Now**.
 7. To save the report, click **OK**. User/Group activity reports cannot be saved on the firewall; to schedule the report for email delivery, see [Schedule Reports for Email Delivery](#).
-

Manage Report Groups

Report groups allow you to create sets of reports that the system can compile and send as a single aggregate PDF report with an optional title page and all the constituent reports included.

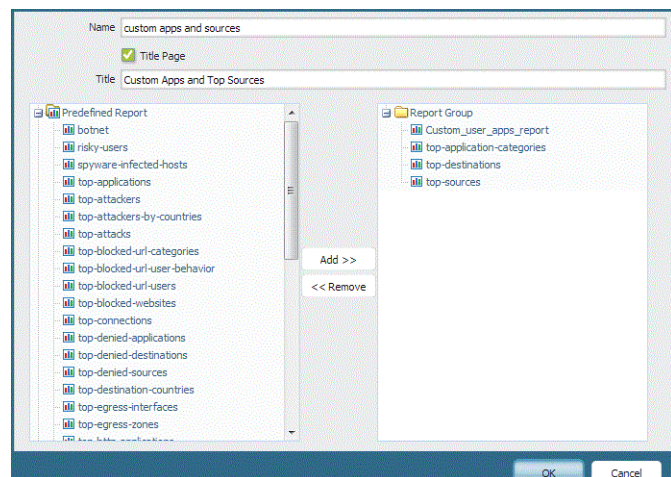
Set up Report Groups

Step 1 Set up report groups.



You must set up a **Report Group** to email report(s).

1. [Create a Server Profile for your email server.](#)
2. Define the **Report Group**. A report group can compile predefined reports, PDF Summary reports, custom reports, and Log View report into a single PDF.
 - a. Select **Monitor > Report Group**.
 - b. Click **Add** and then enter a **Name** for the report group.
 - c. (Optional) Select **Title Page** and add a **Title** for the PDF output.
 - d. Select reports from the left column and click **Add** to move each report to the report group on the right.



The **Log View** report is a report type that is automatically created each time you create a custom report and uses the same name as the custom report. This report will show the logs that were used to build the contents of the custom report.

To include the log view data, when creating a report group, add your custom report under the **Custom Reports** list and then add the log view report by selecting the matching report name from the **Log View** list. The report will include the custom report data and the log data that was used to create the custom report.

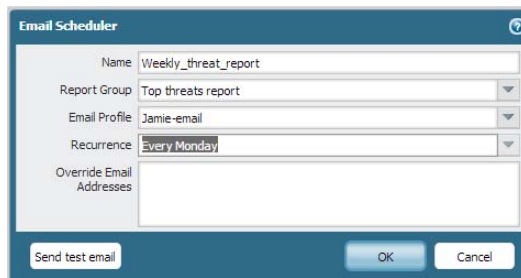
- e. Click **OK** to save the settings.
- f. To use the report group, see [Schedule Reports for Email Delivery](#).

Schedule Reports for Email Delivery

Reports can be scheduled for daily delivery or delivered weekly on a specified day. Scheduled reports are executed starting at 2:00 AM, and email delivery starts after all scheduled reports have been generated.

Schedule Reports for Email Delivery

1. Select **Monitor > PDF Reports > Email Scheduler**.
2. Select the **Report Group** for email delivery. To set up a report group; see [Manage Report Groups](#).
3. Select the frequency at which to generate and send the report in **Recurrence**.
4. Select the email server profile to use for delivering the reports. To set up an email server profile, see [Create a Server Profile for your email server](#).
5. The **Override Recipient email(s)** allows you to send this report exclusively to the recipients specified in this field. When you add recipients to the **Override Recipient email(s)**, the report is not sent to the recipients configured in the email server profile. Use this option for those occasions when the report is for the attention of someone other than the administrators or recipients defined in the email server profile.



The screenshot shows the 'Email Scheduler' dialog box with the following fields and values:

- Name: Weekly_threat_report
- Report Group: Top threats report
- Email Profile: Jamie-email
- Recurrence: Every Monday
- Override Email Addresses: (empty text area)

Buttons at the bottom: Send test email, OK, Cancel.

Syslog Field Descriptions

This is a list of the standard fields for each of the five log types that are forwarded to an external server. For ease of parsing, the comma is the delimiter; each field is a comma-separated value (CSV) string. The FUTURE_USE tag applies to fields that the devices do not currently implement.



WildFire logs are a subtype of threat logs and use the same Syslog format.

- ▲ [Traffic Logs](#)
- ▲ [Threat Logs](#)
- ▲ [HIP Match Logs](#)
- ▲ [Config Logs](#)
- ▲ [System Logs](#)
- ▲ [Syslog Severity](#)
- ▲ [Custom Log/Event Format](#)
- ▲ [Escape Sequences](#)

Traffic Logs

Format: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source IP, Destination IP, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Ingress Interface, Egress Interface, Log Forwarding Profile, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Bytes, Bytes Sent, Bytes Received, Packets, Start Time, Elapsed Time, Category, FUTURE_USE, Sequence Number, Action Flags, Source Location, Destination Location, FUTURE_USE, Packets Sent, Packets Received, Session End Reason*

Field Name	Description
Receive Time (receive_time)	Time the log was received at the management plane
Serial Number (serial)	Serial number of the device that generated the log
Type (type)	Specifies type of log; values are traffic, threat, config, system and hip-match
Subtype (subtype)	Subtype of traffic log are as follows: <ul style="list-style-type: none"> • start—session started • end—session ended • drop—session dropped before the application is identified and there is no rule that allows the session. • deny—session dropped after the application is identified and there is a rule to block or no rule that allows the session.
Generated Time (time_generated)	Time the log was generated on the dataplane

Field Name	Description
Source IP (src)	Original session source IP address
Destination IP (dst)	Original session destination IP address
NAT Source IP (natsrc)	If Source NAT performed, the post-NAT Source IP address
NAT Destination IP (natdst)	If Destination NAT performed, the post-NAT Destination IP address
Rule Name (rule)	Name of the rule that the session matched
Source User (srcuser)	Username of the user who initiated the session
Destination User (dstuser)	Username of the user to which the session was destined
Application (app)	Application associated with the session
Virtual System (vsys)	Virtual System associated with the session
Source Zone (from)	Zone the session was sourced from
Destination Zone (to)	Zone the session was destined to
Ingress Interface (inbound_if)	Interface that the session was sourced from
Egress Interface (outbound_if)	Interface that the session was destined to
Log Forwarding Profile (logset)	Log Forwarding Profile that was applied to the session
Session ID (sessionid)	An internal numerical identifier applied to each session
Repeat Count (repeatcnt)	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only
Source Port (sport)	Source port utilized by the session
Destination Port (dport)	Destination port utilized by the session
NAT Source Port (natsport)	Post-NAT source port
NAT Destination Port (natdport)	Post-NAT destination port

Field Name	Description
Flags (flags)	32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value: <ul style="list-style-type: none"> 0x80000000—session has a packet capture (PCAP) 0x02000000—IPv6 session 0x01000000—SSL session was decrypted (SSL Proxy) 0x00800000—session was denied via URL filtering 0x00400000—session has a NAT translation performed (NAT) 0x00200000—user information for the session was captured via the captive portal (Captive Portal) 0x00080000—X-Forwarded-For value from a proxy is in the source user field 0x00040000—log corresponds to a transaction within a http proxy session (Proxy Transaction) 0x00008000—session is a container page access (Container Page) 0x00002000—session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above. 0x00000800—symmetric return was used to forward traffic for this session
Protocol (proto)	IP protocol associated with the session
Action (action)	Action taken for the session; values are allow or deny: <ul style="list-style-type: none"> Allow—session was allowed by policy Deny—session was denied by policy
Bytes (bytes)	Number of total bytes (transmit and receive) for the session
Bytes Sent (bytes_sent)	Number of bytes in the client-to-server direction of the session Available on all models except the PA-4000 Series
Bytes Received (bytes_received)	Number of bytes in the server-to-client direction of the session Available on all models except the PA-4000 Series
Packets (packets)	Number of total packets (transmit and receive) for the session
Start Time (start)	Time of session start
Elapsed Time (elapsed)	Elapsed time of the session
Category (category)	URL category associated with the session (if applicable)
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7050 firewalls.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama
Source Location (srcloc)	Source country or Internal region for private addresses; maximum length is 32 bytes

Field Name	Description
Destination Location (dstloc)	Destination country or Internal region for private addresses. Maximum length is 32 bytes
Packets Sent (pkts_sent)	Number of client-to-server packets for the session Available on all models except the PA-4000 Series
Packets Received (pkts_received)	Number of server-to-client packets for the session Available on all models except the PA-4000 Series
Session End Reason (session_end_reason) New in v6.1!	<p>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):</p> <ul style="list-style-type: none"> • threat—The firewall detected a threat associated with a reset, drop, or block (IP address) action. • policy-deny—The session matched a security policy with a deny or drop action. • 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: <ul style="list-style-type: none"> • 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 <i>unknown</i> 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 <i>unknown</i>. • n/a—This value applies when the traffic log type is not <i>end</i>.

Threat Logs

Format: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source IP, Destination IP, NAT Source IP, NAT Destination IP, Rule Name, Source User, Destination User, Application, Virtual System, Source Zone, Destination Zone, Ingress Interface, Egress Interface, Log Forwarding Profile, FUTURE_USE, Session ID, Repeat Count, Source Port, Destination Port, NAT Source Port, NAT Destination Port, Flags, Protocol, Action, Miscellaneous, Threat ID, Category, Severity, Direction, Sequence Number, Action Flags, Source Location, Destination Location, FUTURE_USE, Content Type, PCAP_id, Filedigest, Cloud, FUTURE_USE, User Agent*, File Type*, X-Forwarded-For*, Referer*, Sender*, Subject*, Recipient*, Report ID*

Field Name	Description
Receive Time (receive_time)	Time the log was received at the management plane
Serial Number (serial)	Serial number of the device that generated the log
Type (type)	Specifies type of log; values are traffic, threat, config, system and hip-match
Subtype (subtype)	Subtype of threat log. Values include the following: <ul style="list-style-type: none"> • data—Data pattern matching a Data Filtering profile. • file—File type matching a File Blocking profile. • flood—Flood detected via a Zone Protection profile. • packet—Packet-based attack protection triggered by a Zone Protection profile. • scan—Scan detected via a Zone Protection profile. • spyware—Spyware detected via an Anti-Spyware profile. • url—URL filtering log. • virus—Virus detected via an Antivirus profile. • vulnerability—Vulnerability exploit detected via a Vulnerability Protection profile. • wildfire—A WildFire verdict generated when the firewall submits a file to WildFire per a WildFire Analysis profile and a verdict (malicious, grayware, or benign, depending on what you are logging) is logged in the WildFire Submissions log. • wildfire-virus—Virus detected via an Antivirus profile.
Generated Time (time_generated)	Time the log was generated on the dataplane
Source IP (src)	Original session source IP address
Destination IP (dst)	Original session destination IP address
NAT Source IP (natsrc)	If source NAT performed, the post-NAT source IP address
NAT Destination IP (natdst)	If destination NAT performed, the post-NAT destination IP address
Rule Name (rule)	Name of the rule that the session matched
Source User (srcuser)	Username of the user who initiated the session
Destination User (dstuser)	Username of the user to which the session was destined

Field Name	Description
Application (app)	Application associated with the session
Virtual System (vsys)	Virtual System associated with the session
Source Zone (from)	Zone the session was sourced from
Destination Zone (to)	Zone the session was destined to
Ingress Interface (inbound_if)	Interface that the session was sourced from
Egress Interface (outbound_if)	Interface that the session was destined to
Log Forwarding Profile (logset)	Log Forwarding Profile that was applied to the session
Session ID (sessionid)	An internal numerical identifier applied to each session
Repeat Count (repeatcnt)	Number of sessions with same Source IP, Destination IP, Application, and Subtype seen within 5 seconds; used for ICMP only
Source Port (sport)	Source port utilized by the session
Destination Port (dport)	Destination port utilized by the session
NAT Source Port (natsport)	Post-NAT source port
NAT Destination Port (natdport)	Post-NAT destination port
Flags (flags)	32-bit field that provides details on session; this field can be decoded by AND-ing the values with the logged value: <ul style="list-style-type: none"> 0x80000000—session has a packet capture (PCAP) 0x02000000—IPv6 session 0x01000000—SSL session was decrypted (SSL Proxy) 0x00800000—session was denied via URL filtering 0x00400000—session has a NAT translation performed (NAT) 0x00200000—user information for the session was captured via the captive portal (Captive Portal) 0x00080000—X-Forwarded-For value from a proxy is in the source user field 0x00040000—log corresponds to a transaction within a http proxy session (Proxy Transaction) 0x00008000—session is a container page access (Container Page) 0x00002000—session has a temporary match on a rule for implicit application dependency handling. Available in PAN-OS 5.0.0 and above 0x00000800—symmetric return was used to forward traffic for this session
Protocol (proto)	IP protocol associated with the session

Field Name	Description
Action (action)	<p>Action taken for the session; values are alert, allow, deny, drop, drop-all-packets, reset-client, reset-server, reset-both, block-url.</p> <ul style="list-style-type: none"> • Alert—threat or URL detected but not blocked • Allow—flood detection alert • Deny—flood detection mechanism activated and deny traffic based on configuration • Drop—threat detected and associated session was dropped • Drop-all-packets—threat detected and session remains, but drops all packets • Reset-client—threat detected and a TCP RST is sent to the client • Reset-server—threat detected and a TCP RST is sent to the server • Reset-both—threat detected and a TCP RST is sent to both the client and the server • Block-url—URL request was blocked because it matched a URL category that was set to be blocked
Miscellaneous (misc)	<p>Field with variable length with a maximum of 1023 characters</p> <p>The actual URI when the subtype is URL</p> <p>File name or file type when the subtype is file</p> <p>File name when the subtype is virus</p> <p>File name when the subtype is WildFire</p>
Threat ID (threatid)	<p>Palo Alto Networks identifier for the threat. It is a description string followed by a 64-bit numerical identifier in parentheses for some Subtypes:</p> <ul style="list-style-type: none"> • 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 • 100000 – 2999999—virus detection • 3000000 – 39999999—WildFire signature feed • 4000000-49999999—DNS Botnet signatures
Category (category)	<p>For URL Subtype, it is the URL Category; For WildFire subtype, it is the verdict on the file and is either 'malicious' or 'benign'; For other subtypes, the value is 'any'.</p>
Severity (severity)	<p>Severity associated with the threat; values are informational, low, medium, high, critical</p>

Field Name	Description
Direction (direction)	Indicates the direction of the attack, <i>client-to-server</i> or <i>server-to-client</i> <ul style="list-style-type: none"> 0—direction of the threat is client to server 1—direction of the threat is server to client
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially. Each log type has a unique number space. This field is not supported on PA-7050 firewalls.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Source Location (srcloc)	Source country or Internal region for private addresses. Maximum length is 32 bytes.
Destination Location (dstloc)	Destination country or Internal region for private addresses. Maximum length is 32 bytes.
Content Type (contenttype)	Applicable only when Subtype is URL. Content type of the HTTP response data. Maximum length 32 bytes.
PCAP ID (pcap_id)	Pcap-ID is a 64 bit unsigned integral denoting an ID to correlate threat pcap files with extended pcaps taken as a part of that flow. All threat logs will contain either a pcap_id of 0 (no associated pcap), or an ID referencing the extended pcap file.
File Digest (filedigest)	Only for WildFire subtype; all other types do not use this field 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.
New in v6.1! 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.
New in v6.1! 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.
New in v6.1! 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.
New in v6.1! 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.
New in v6.1! Sender (sender)	Only for WildFire subtype; all other types do not use this field. Specifies the name of the sender of an email that WildFire determined to be malicious when analyzing an email link forwarded by the firewall.

Field Name	Description
Subject (subject) New in v6.1!	Only for WildFire subtype; all other types do not use this field. Specifies the subject of an email that WildFire determined to be malicious when analyzing an email link forwarded by the firewall.
Recipient (recipient) New in v6.1!	Only for WildFire subtype; all other types do not use this field. Specifies the name of the receiver of an email that WildFire determined to be malicious when analyzing an email link forwarded by the firewall.
Report ID (reportid) New in v6.1!	Only for WildFire subtype; all other types do not use this field. Identifies the analysis request on the WildFire cloud or the WildFire appliance.

HIP Match Logs

Format: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Source User, Virtual System, Machine name, OS, Source Address, HIP, Repeat Count, HIP Type, FUTURE_USE, FUTURE_USE, Sequence Number, Action Flags

Field Name	Description
Receive Time (receive_time)	Time the log was received at the management plane
Serial Number (serial)	Serial number of the device that generated the log
Type (type)	Type of log; values are traffic, threat, config, system and hip-match
Subtype (subtype)	Subtype of HIP match log; unused
Generated Time (time_generated)	Time the log was generated on the dataplane
Source User (srcuser)	Username of the user who initiated the session
Virtual System (vsys)	Virtual System associated with the HIP match log
Machine Name (machinename)	Name of the user's machine
OS	The operating system installed on the user's machine or device (or on the client system)
Source Address (src)	IP address of the source user
HIP (matchname)	Name of the HIP object or profile
Repeat Count (repeatcnt)	Number of times the HIP profile matched
HIP Type (matchtype)	Whether the hip field represents a HIP object or a HIP profile
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7050 firewalls.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama

Config Logs

Format: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Host, Virtual System, Command, Admin, Client, Result, Configuration Path, Sequence Number, Action Flags, Before Change Detail*, After Change Detail*

Field Name	Description
Receive Time (receive_time)	Time the log was received at the management plane
Serial Number (serial)	Serial number of the device that generated the log
Type (type)	Type of log; values are traffic, threat, config, system and hip-match
Subtype (subtype)	Subtype of configuration log; unused
Generated Time (time_generated)	Time the log was generated on the dataplane
Host (host)	Host name or IP address of the client machine
Virtual System (vsys)	Virtual System associated with the configuration log
Command (cmd)	Command performed by the Admin; values are add, clone, commit, delete, edit, move, rename, set.
Admin (admin)	Username of the Administrator performing the configuration
Client (client)	Client used by the Administrator; values are Web and CLI
Result (result)	Result of the configuration action; values are Submitted, Succeeded, Failed, and Unauthorized
Configuration Path (path)	The path of the configuration command issued; up to 512 bytes in length
Sequence Number (seqno)	A 64bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7050 firewalls.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama.
Before Change Detail (before_change_detail) New in v6.1!	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) New in v6.1!	This field is in custom logs only; it is not in the default format. It contains the full xpath after the configuration change.

System Logs

Format: FUTURE_USE, Receive Time, Serial Number, Type, Subtype, FUTURE_USE, Generated Time, Virtual System, Event ID, Object, FUTURE_USE, FUTURE_USE, Module, Severity, Description, Sequence Number, Action Flags

Field Name	Description
Receive Time (receive_time)	Time the log was received at the management plane
Serial Number (serial)	Serial number of the device that generated the log
Type (type)	Type of log; values are traffic, threat, config, system and hip-match
Subtype (subtype)	Subtype of the system log; refers to the system daemon generating the log; values are crypto, dhcp, dnsproxy, dos, general, global-protect, ha, hw, nat, ntpd, pbf, port, pppoe, ras, routing, satd, sslmgr, sslvpn, userid, url-filtering, vpn
Generated Time (time_generated)	Time the log was generated on the dataplane
Virtual System (vsys)	Virtual System associated with the configuration log
Event ID (eventid)	String showing the name of the event
Object (object)	Name of the object associated with the system event
Module (module)	This field is valid only when the value of the Subtype field is general. It provides additional information about the sub-system generating the log; values are general, management, auth, ha, upgrade, chassis
Severity (severity)	Severity associated with the event; values are informational, low, medium, high, critical
Description (opaque)	Detailed description of the event, up to a maximum of 512 bytes
Sequence Number (seqno)	A 64-bit log entry identifier incremented sequentially; each log type has a unique number space. This field is not supported on PA-7050 firewalls.
Action Flags (actionflags)	A bit field indicating if the log was forwarded to Panorama

Syslog Severity

The syslog severity is set based on the log type and contents.

Log Type/Severity	Syslog Severity
Traffic	Info
Config	Info
Threat/System—Informational	Info
Threat/System—Low	Notice
Threat/System—Medium	Warning

Log Type/Severity	Syslog Severity
Threat/System—High	Error
Threat/System—Critical	Critical

Custom Log/Event Format

To facilitate the integration with external log parsing systems, the firewall allows you to customize the log format; it also allows you to add custom *Key: Value* attribute pairs. Custom message formats can be configured under **Device > Server Profiles > Syslog > Syslog Server Profile > Custom Log Format**.

To achieve ArcSight Common Event Format (CEF) compliant log formatting, refer to the [CEF Configuration Guide](#).

Escape Sequences

Any field that contains a comma or a double-quote is enclosed in double quotes. Furthermore, if a double-quote appears inside a field it is escaped by preceding it with another double-quote. To maintain backward compatibility, the Misc field in threat log is always enclosed in double-quotes.



User-ID

User Identification (User-ID) is a Palo Alto Networks next-generation firewall feature that allows you to create policies and perform reporting based on users and groups rather than individual IP addresses. The following sections describe the Palo Alto Networks User-ID feature and provide instructions on setting up user- and group-based access:

- ▲ [User-ID Overview](#)
- ▲ [User-ID Concepts](#)
- ▲ [Enable User-ID](#)
- ▲ [Map Users to Groups](#)
- ▲ [Map IP Addresses to Users](#)
- ▲ [Configure a Firewall to Share User Mapping Data with Other Firewalls](#)
- ▲ [Enable User- and Group-Based Policy](#)
- ▲ [Verify the User-ID Configuration](#)

User-ID Overview

User-ID seamlessly integrates Palo Alto Networks firewalls with a range of enterprise directory and terminal services offerings, enabling you to tie application activity and security policies to users and groups—not just IP addresses. In addition, with User-ID enabled, the Application Command Center (ACC), App-Scope, reports, and logs all include usernames in addition to user IP addresses.

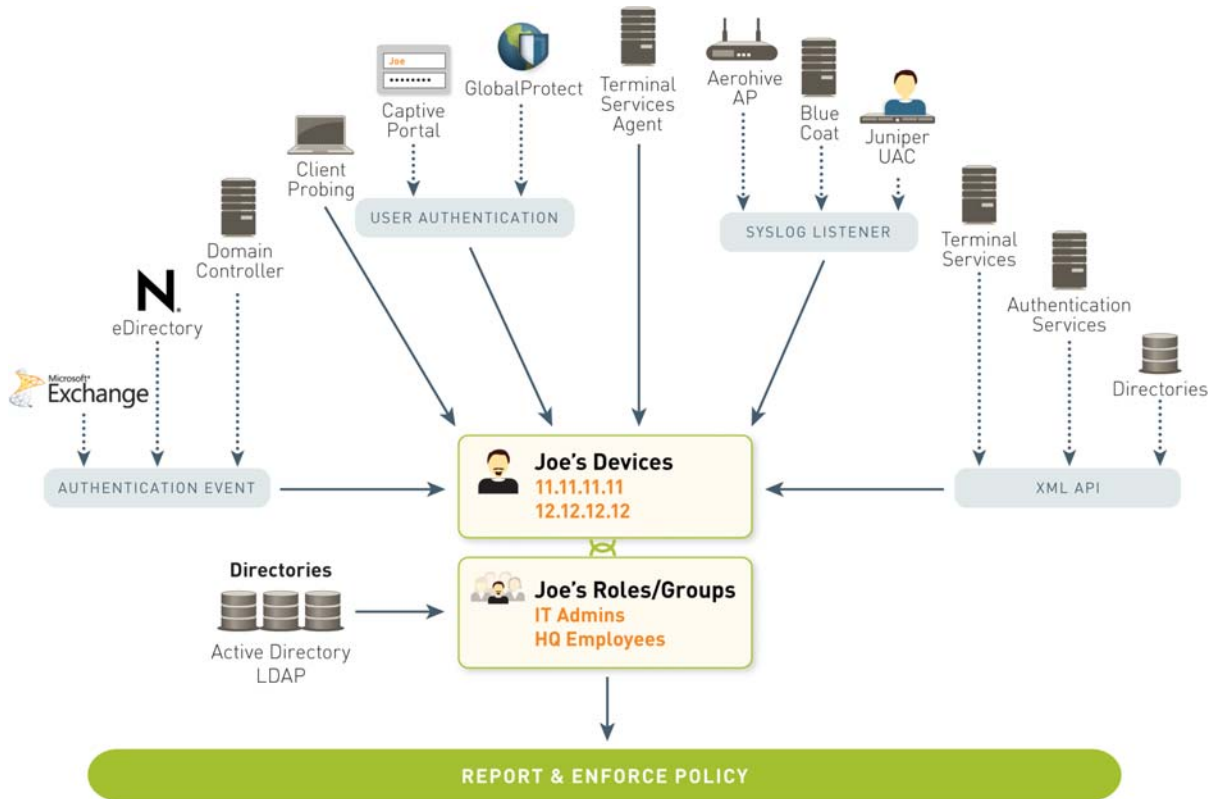
The Palo Alto Networks next-generation firewall supports monitoring of the following enterprise services:

- Microsoft Active Directory
- Lightweight Directory Access Protocol (LDAP)
- Novell eDirectory
- Citrix Metaframe Presentation Server or XenApp
- Microsoft Terminal Services

To be able to create policy based on user and group, the firewall must have a list of all available users and their corresponding group mappings that you can select from when defining your policies. It gets this [Group Mapping](#) information by connecting directly to your LDAP directory server.

To be able to enforce the user- and group-based policies, the firewall must be able to map the IP addresses in the packets it receives to usernames. User-ID provides many mechanisms to get these IP address-to-username mappings. For example, it uses agents to monitor server logs for login events and/or probe clients, and/or listen for syslog messages from authenticating services. To identify mappings for IP addresses that were not mapped using one of the agent mechanisms, you can configure the firewall to redirect HTTP requests to a captive portal login. You can tailor the mechanisms you use for [User Mapping](#) to suit your environment, and you can even use different mechanisms at different sites.

Figure: User-ID



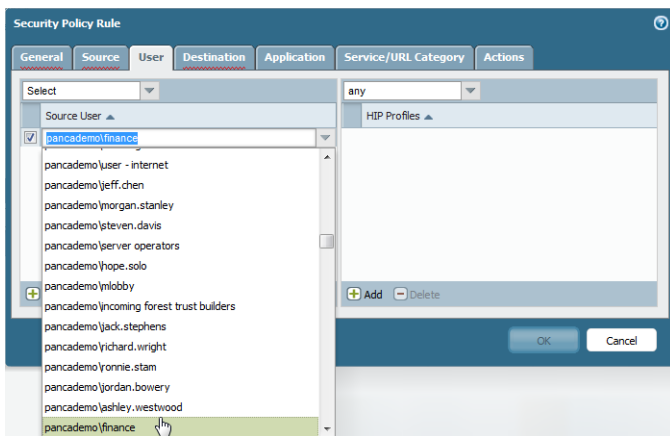
Proceed to [User-ID Concepts](#) for information on how User-ID works and [Enable User-ID](#) for instructions on setting up User-ID on the firewall.

User-ID Concepts

- ▲ Group Mapping
- ▲ User Mapping

Group Mapping

In order to define security policies based on user or group, the firewall must retrieve the list of groups and the corresponding list of members from your directory server. To enable this functionality to [Map Users to Groups](#), you must create an LDAP server profile that instructs the firewall how to connect and authenticate to the server and how to search the directory for the user and group information. After you connect to the LDAP server and configure the group mapping functionality for user identification, you will be able to select users or groups when defining your security policies. The firewall supports a variety of LDAP directory servers, including Microsoft Active Directory (AD), Novell eDirectory, and Sun ONE Directory Server.



You can then define policies based on group membership rather than on individual users for simplified administration. For example, the following security policy allows access to specific internal applications based on group membership:

	Name	Tags	Zone	Source			Destination		Application	Service	Action
				Address	User	HIP Profile	Zone	Address			
1	CRM access	none	13-trust	any	Finance	Missing Patch ...	13-trust	any	sap	application-default	✓
2	Eng access	none	13-trust	any	Engineering	Missing Patch ...	13-trust	any	bugzilla perforce	application-default	✓

User Mapping

Having the names of the users and groups is only one piece of the puzzle. The firewall also needs to know which IP addresses map to which users so that security policies can be enforced appropriately. [Figure: User-ID](#) illustrates the different methods that are used to identify users and groups on your network and shows how user mapping and group mapping work together to enable user- and group-based security enforcement and visibility.

The following topics describe the different methods of user mapping:

- ▲ Server Monitoring
- ▲ Client Probing
- ▲ Port Mapping
- ▲ Syslog
- ▲ Captive Portal
- ▲ GlobalProtect
- ▲ User-ID XML API

Server Monitoring

With server monitoring a User-ID agent—either a Windows-based agent running on a domain server in your network, or the integrated PAN-OS User-ID agent running on the firewall—monitors the security event logs for specified Microsoft Exchange Servers, domain controllers, or Novell eDirectory servers for login events. For example, in an AD environment, you can configure the User-ID agent to monitor the security logs for Kerberos ticket grants or renewals, Exchange server access (if configured), and file and print service connections. Keep in mind that in order for these events to be recorded in the security log, the AD domain must be configured to log successful account login events. In addition, because users can log in to any of the servers in the domain, you must set up server monitoring for all servers in order to capture all user login events.

Because server monitoring requires very little overhead and because the majority of users can generally be mapped using this method, it is recommended as the base user mapping method for most User-ID deployments. See [Configure User Mapping Using the Windows User-ID Agent](#) or [Configure User Mapping Using the PAN-OS Integrated User-ID Agent](#) for details.

Client Probing

In a Microsoft Windows environment, you can configure the User-ID agent to probe client systems using Windows Management Instrumentation (WMI). The Windows-based User-ID agent can also perform NetBIOS probing (not supported on the PAN-OS integrated User-ID agent). Probing is particularly useful in environments with a high IP address turnover because changes will be reflected on the firewall more quickly, enabling more accurate enforcement of user-based policies. However, if the correlation between IP addresses and users is fairly static, you probably do not need to enable client probing. Because probing can generate a large amount of network traffic (based on the total number of mapped IP addresses), the agent that will be initiating the probes should be located as close as possible to the end clients.

If probing is enabled, the agent will probe each learned IP address periodically (every 20 minutes by default, but this is configurable) to verify that the same user is still logged in. In addition, when the firewall encounters an IP address for which it has no user mapping, it will send the address to the agent for an immediate probe.

See [Configure User Mapping Using the Windows User-ID Agent](#) or [Configure User Mapping Using the PAN-OS Integrated User-ID Agent](#) for details.

Port Mapping

In environments with multi-user systems—such as Microsoft Terminal Server or Citrix environments—many users share the same IP address. In this case, the user-to-IP address mapping process requires knowledge of the source port of each client. To perform this type of mapping, you must install the Palo Alto Networks Terminal Services Agent on the Windows/Citrix terminal server itself to intermediate the assignment of source ports to the various user processes. For terminal servers that do not support the Terminal Services Agent, such as Linux terminal servers, you can use the XML API to send user mapping information from login and logout events to User-ID. See [Configure User Mapping for Terminal Server Users](#) for configuration details.

Syslog

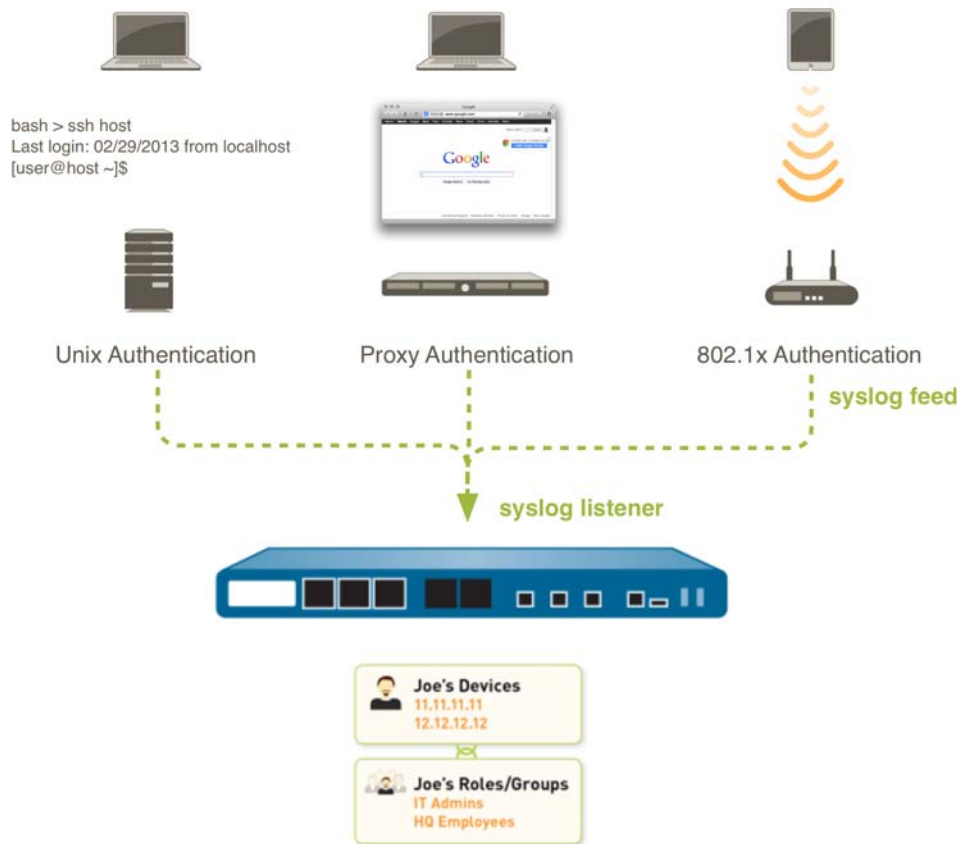
Your environment might have existing network services that authenticate users. These services include wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, and other Network Access Control (NAC) mechanisms. You can configure these services to send syslog messages and configure the User-ID agent to parse the messages for login events. The agent then maps IP addresses to usernames based on the login events.

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. 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 Receive User Mappings from a Syslog Sender](#) for configuration details.

Figure: User-ID Integration with Syslog

Captive Portal

If the firewall or the User-ID agent are unable to map an IP address to a user—for example if the user is not logged in or is using an operating system such as Linux that is not supported by your domain servers—you can configure Captive Portal. When configured, any web traffic (HTTP or HTTPS) matching your Captive Portal policy requires user authentication, either transparently via an NT LAN Manager (NTLM) challenge to the browser, or actively by redirecting the user to a web authentication form for authentication against a RADIUS, LDAP, Kerberos, or local authentication database or using client certificate authentication. See [Map IP Addresses to User Names Using Captive Portal](#) for details.

GlobalProtect

For mobile or roaming users, the GlobalProtect client provides the user mapping information to the firewall directly. In this case, every GlobalProtect user has an agent or app running on the client that requires the user to enter login credentials for VPN access to the firewall. This login information is then added to the User-ID user mapping table on the firewall for visibility and user-based security policy enforcement. Because GlobalProtect users must authenticate to gain access to the network, the IP address-to-username mapping is

explicitly known. This is the best solution in sensitive environments where you must be certain of who a user is in order to allow access to an application or service. For more information on setting up GlobalProtect, refer to the [GlobalProtect Administrator's Guide](#).

User-ID XML API

For other types of user access that cannot be mapped using any of the standard user mapping methods or Captive Portal—for example, to add mappings of users connecting from a third-party VPN solution or users connecting to a 802.1x enabled wireless network—you can use the User-ID XML API to capture login events and send them to the User-ID agent or directly to the firewall. See [Send User Mappings to User-ID Using the XML API](#) for details.

Enable User-ID

You must complete the following tasks to set up the firewall to user users and groups in policy enforcement, logging, and reporting:

- [Map Users to Groups](#)
- [Map IP Addresses to Users](#)
- [Enable User- and Group-Based Policy](#)
- [Verify the User-ID Configuration](#)

Map Users to Groups

Use the following procedure to connect to your LDAP directory to enable the firewall to retrieve user-to-group mapping information:



Best practices for group mapping in an Active Directory environment:

- If you have a single domain, you only need one LDAP server profile that connects the firewall to the domain controller with the best connectivity. You can add additional domain controllers for fault tolerance.
- If you have multiple domains and/or multiple forests, you must create a server profile to connect to a domain server in each domain/forest. Take steps to ensure unique usernames in separate forests.
- If you have Universal Groups, create a server profile to connect to the Global Catalogue server.

Map Users to Groups

Step 1 Create an LDAP Server Profile that specifies how to connect to the directory servers you want the firewall to use to obtain group mapping information.

1. Select **Device > Server Profiles > LDAP**.
2. Click **Add** and then enter a **Name** for the profile.
3. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down.

4. Click **Add** to add a new LDAP server entry and then enter a **Server** name to identify the server (1-31 characters) and the IP **Address** and **Port** number the firewall should use to connect to the LDAP server (default=389 for LDAP; 636 for LDAP over SSL). You can add up to four LDAP servers to the profile, however, all the servers you add to a profile must be of the same type. For redundancy you should add at least two servers.
5. Enter the LDAP **Domain** name to prepend to all objects learned from the server. The value you enter here depends on your deployment:
 - If you are using Active Directory, you must enter the NetBIOS domain name; NOT a FQDN (for example, enter acme, not acme.com). Note that if you need to collect data from multiple domains, you will need to create separate server profiles.
 - If you are using a global catalog server, leave this field blank.
6. Select the **Type** of LDAP server you are connecting to. The correct LDAP attributes in the group mapping settings will automatically be populated based on your selection. However, if you have customized your LDAP schema, you may need to modify the default settings.
7. In the **Base** field, select the DN that corresponds to the point in the LDAP tree where you want the firewall to begin its search for user and group information.
8. Enter the authentication credentials for binding to the LDAP tree in the **Bind DN**, **Bind Password**, and **Confirm Bind Password** fields. The Bind DN can be in either User Principal Name (UPN) format (for example, administrator@acme.local) or it can be a fully qualified LDAP name (for example, cn=administrator,cn=users,dc=acme,dc=local).
9. If you want the firewall to communicate with the LDAP server(s) over a secure connection, select the SSL check box. If you enable SSL, make sure that you have also specified the appropriate port number.
10. Click **OK** to save the profile.

Map Users to Groups (Continued)

Step 2 Add the LDAP server profile to the User-ID Group Mapping configuration.

1. Select **Device > User Identification > Group Mapping Settings** and click **Add**.
2. Enter a unique **Name** to identify the group mapping configuration.
3. Select the **Server Profile** you created in [Step 1](#).
4. (Optional) To filter the groups that User-ID tracks for group mapping, in the Group Objects section enter a **Search Filter** (LDAP query), then specify the **Object Class** (group definition), **Group Name**, and **Group Member**.
5. (Optional) To filter the users that User-ID tracks for group mapping, in the User Objects section enter a **Search Filter** (LDAP query), then specify the **Object Class** (user definition) and **User Name**.
6. (Optional) To match User-ID information with email header information identified in the links and attachments of emails forwarded to WildFire, enter the list of email domains in your organization in the Mail Domains section, **Domain List** field. Use commas to separate multiple domains (up to 256 characters). After you click **OK** (Step 9), PAN-OS automatically populates the **Mail Attributes** field based on your LDAP server type (Sun/RFC, Active Directory, or Novell). When a match occurs, the username in the WildFire log email header section will contain a link that opens the **ACC** tab, filtered by the user or group of users.
7. Make sure the **Enabled** check box is selected.
8. (Optional) To limit which groups appear in security policy, select the **Group Include List** tab and browse the LDAP tree to locate the desired groups. For each group you want to include, select it in the **Available Groups** list and click the add **+** icon to move it to the **Included Groups** list.
9. Click **OK** to save the settings.

Step 3 **Commit** the configuration.

Map IP Addresses to Users

The tasks you need to perform to map IP addresses to usernames depends on the type and location of the client systems on your network. Complete as many of the following tasks as necessary to enable mapping of your client systems:

- ❑ To map users as they log in to your Exchange servers, domain controllers, or eDirectory servers, or Windows clients that can be directly probed, you must configure a User-ID agent to monitor the server logs and/or probe client systems. You can either install the standalone Windows User-ID agent on one or more member servers in the domain that contains the servers and clients to be monitored (see [Configure User Mapping Using the Windows User-ID Agent](#)) or you can configure the on-firewall User-ID agent that is integrated with PAN-OS ([Configure User Mapping Using the PAN-OS Integrated User-ID Agent](#)). For guidance as to which agent configuration is appropriate for your network and the number and placements of agents that are required, refer to [Architecting User Identification Deployments](#).
- ❑ If you have clients running multi-user systems such as Microsoft Terminal Server or Citrix Metaframe Presentation Server or XenApp, see [Configure the Palo Alto Networks Terminal Server Agent for User Mapping](#) for instructions on how to install and configure the agent on a Windows server. If you have a multi-user system that is not running on Windows, you can use the User-ID XML API to send IP address to username mappings directly to the firewall. See [Retrieve User Mappings from a Terminal Server Using the User-ID XML API](#).
- ❑ To obtain user mappings from existing network services that authenticate users, such as wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, or other Network Access Control (NAC) mechanisms, configure the User-ID agent (either the Windows agent or the agentless user mapping feature on the firewall) to listen for authentication syslog messages from those services. See [Configure User-ID to Receive User Mappings from a Syslog Sender](#).
- ❑ If you have users with client systems that are not logged into your domain servers—for example, users running Linux clients that do not log in to the domain—see [Map IP Addresses to User Names Using Captive Portal](#).
- ❑ For other clients that you are unable to map using the previous methods, you can use the User-ID XML API to add user mappings directly to the firewall. See [Send User Mappings to User-ID Using the XML API](#).
- ❑ Because policy is local to each firewall, each firewall must have a current list of IP address-to-username mappings in order to accurately enforce security policy by group or user. However, you can configure one firewall to collect all the user mappings and distribute them to the other firewalls. For details, see [Configure a Firewall to Share User Mapping Data with Other Firewalls](#).

Configure User Mapping Using the Windows User-ID Agent

In most cases, the majority of your network users will have logins to your monitored domain services. For these users, the Palo Alto Networks User-ID agent monitors the servers for login events and performs the IP address-to-user mapping. The way you configure the User-ID agent depends on the size of your environment and the location of your domain servers. As a best practice, you should locate your User-ID agents near your monitored servers (that is, the monitored servers and the Windows User-ID agent should not be across a WAN link from each other). This is because most of the traffic for user mapping occurs between the agent and the monitored server, with only a small amount of traffic—the delta of IP address mappings since the last update—from the agent to the firewall.

The following topics describe how to install and configure the User-ID Agent and how to configure the firewall to retrieve user mapping information from the agent:

- ▲ [Install the User-ID Agent](#)
- ▲ [Configure the User-ID Agent for User Mapping](#)


Install the User-ID Agent


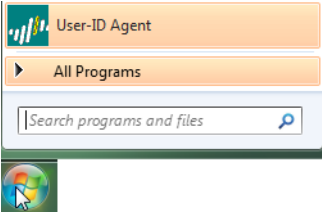
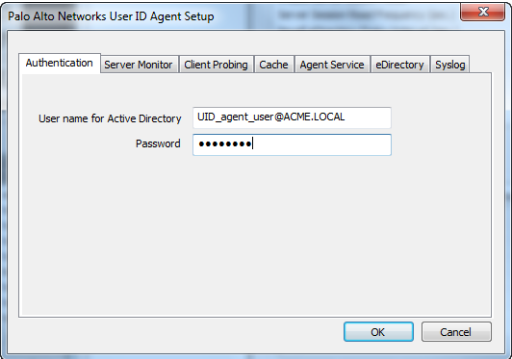
The following procedure shows how to install the User-ID agent on a member server in the domain and set up the service account with the required permissions. If you are upgrading, the installer will automatically remove the older version, however, it is a good idea to back up the config.xml file before running the installer.



For information about the system requirements for installing the Windows-based User-ID agent and for information on the supported server OS versions are supported, refer to “Operating System (OS) Compatibility User-ID Agent” in the User-ID Agent Release Notes, which are available on the Palo Alto Networks [Software Updates](#) page.

Install the Windows User-ID Agent

<p>Step 1 Decide where to install the User-ID agent(s).</p> <p>The User-ID agent queries the Domain Controller and Exchange server logs using Microsoft Remote Procedure Calls (MSRPCs), which require a complete transfer of the entire log at each query. Therefore, you should always install one or more User-ID agents at each site that has servers to be monitored.</p> <p> For more detailed information on where to install User-ID agents, refer to Architecting User Identification (User-ID) Deployments.</p>	<ul style="list-style-type: none"> • You must install the User-ID agent on a system running one of the supported OS versions: see “Operating System (OS) Compatibility User-ID Agent” in the User-ID Agent Release Notes. • Make sure the system you plan to install the User-ID agent on is a member of the domain that the servers it will be monitoring belong to. • As a best practice, install the User-ID agent close to the servers it will be monitoring (there is more traffic between the User-ID agent and the monitored servers than there is between the User-ID agent and the firewall, so locating the agent close to the monitored servers optimizes bandwidth usage). • To ensure the most comprehensive mapping of users, you must monitor all servers that contain user login information. You may need to install multiple User-ID agents to efficiently monitor all of your resources.
---	---

Install the Windows User-ID Agent (Continued)	
<p>Step 2 Download the User-ID agent installer.</p> <p>As a best practice, install the User-ID agent version that is the same as the PAN-OS version running on the firewalls.</p>	<ol style="list-style-type: none"> 1. Log in to Palo Alto Networks Support site. 2. Select Software Updates from the Manage Devices section. 3. Scroll to the User Identification Agent section of the screen and Download the version of the User-ID agent you want to install. 4. Save the UaInstall-x.x.x-xx.msi file on the system(s) where you plan to install the agent.
<p>Step 3 Run the installer as an administrator.</p> 	<ol style="list-style-type: none"> 1. To launch a command prompt as an administrator, click Start and right-click Command Prompt and then select Run as administrator. 2. From the command line, run the .msi file you downloaded. For example, if you saved the .msi file to the Desktop you would enter the following: <pre>C:\Users\administrator.acme>cd Desktop C:\Users\administrator.acme\Desktop>UaInstall-6.0.0-1.msi</pre> 3. Follow the setup prompts to install the agent using the default settings. By default, the agent gets installed to the C:\Program Files (x86)\Palo Alto Networks\User-ID Agent folder, but you can Browse to a different location. 4. When the installation completes, Close the setup window.
<p>Step 4 Launch the User-ID Agent application.</p> 	<ol style="list-style-type: none"> 1. Click Start and select User-ID Agent.
<p>Step 5 (Optional) Change the service account that the User-ID agent uses to log in.</p> 	<p>By default, the agent uses the administrator account used to install the .msi file. However, you may want to switch this to a restricted account as follows:</p> <ol style="list-style-type: none"> 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.

Install the Windows User-ID Agent (Continued)

- Step 6** (Optional) Assign account permissions to the installation folder.
- You only need to perform this step if the service account you configured for the User-ID agent is not a member of the administrators group for the domain or a member of both the Server Operators and the Event Log Readers groups.
1. Give the service account permissions to the installation folder:
 - a. From the Windows Explorer, navigate to `C:\Program Files\Palo Alto Networks` and right-click the folder and select **Properties**.
 - b. On the **Security** tab, **Add** the User-ID agent service account and assign it permissions to **Modify, Read & execute, List folder contents**, and **Read** and then click **OK** to save the account settings.
 2. Give the service account permissions to the User-ID Agent registry sub-tree:
 - a. Run `regedit32` and navigate to the Palo Alto Networks sub-tree in one of the following locations:
 - **32-bit systems**—`HKEY_LOCAL_MACHINE\Software\Palo Alto Networks`
 - **64-bit systems**—`HKEY_LOCAL_MACHINE\Software\WOW6432Node\Palo Alto Networks`
 - b. Right-click the Palo Alto Networks node and select **Permissions**.
 - c. Assign the User-ID service account **Full Control** and then click **OK** to save the setting.
 3. On the domain controller, add the service account to the builtin groups to enable privileges to read the security log events (Event Log Reader group) and open sessions (Server Operator group):
 - a. Run the MMC and Launch the Active Directory Users and Computers snap-in.
 - b. Navigate to the Builtin folder for the domain and then right-click each group you need to edit (Event Log Reader and Server Operator) and select **Add to Group** to open the properties dialog.
 - c. Click **Add** and enter the name of the service account that you configured the User-ID service to use and then click **Check Names** to validate that you have the proper object name.
 - d. Click **OK** twice to save the settings.

Configure the User-ID Agent for User Mapping

The Palo Alto Networks User-ID agent is a Windows service that connects to servers on your network—for example, Active Directory servers, Microsoft Exchange servers, and Novell eDirectory servers—and monitors the logs for login events. The agent uses this information to map IP addresses to usernames. Palo Alto Networks firewalls connect to the User-ID agent to retrieve this user mapping information, enabling visibility into user activity by username rather than IP address and enables user- and group-based security enforcement.




For information about the server OS versions supported by the User-ID agent, refer to “Operating System (OS) Compatibility User-ID Agent” in the *User-ID Agent Release Notes*, which are available on the Palo Alto Networks [Software Updates](#) page.

Map IP Addresses to Users Using the User-ID Agent

Step 1 Launch the User-ID Agent application.

1. Select **User-ID Agent** from the Windows Start menu.



Step 2 Define the servers the User-ID agent should monitor to collect IP address to user mapping information.

The User-ID agent can monitor a total of up to 100 servers, of which up to 50 can be syslog senders.

Keep in mind that to collect all of the required mappings, you must connect to all servers that your users log in to in order to monitor the security log files on all servers that contain login events.

1. Select **User Identification > Discovery**.
2. In the **Servers** section of the screen, click **Add**.
3. Enter a **Name** and **Server Address** for the server to be monitored. The network address can be a FQDN or an IP address.
4. Select the **Server Type (Microsoft Active Directory, Microsoft Exchange, Novell eDirectory, or Syslog Sender)** and then click **OK** to save the server entry. Repeat this step for each server to be monitored.
5. (Optional) To enable the firewall to automatically discover domain controllers on your network using DNS lookups, click **Auto Discover**.




The auto-discovery locates domain controllers in the local domain only; you must manually add Exchange servers, eDirectory servers, and syslog senders.

- e. (Optional) To tune the frequency at which the firewall polls configured servers for mapping information, select **User Identification > Setup** and **Edit** the Setup section. On the **Server Monitor** tab, modify the value in the **Server Log Monitor Frequency (seconds)** field. As a best practice, you should increase the value in this field to 5 seconds in environments with older Domain Controllers or high-latency links. Click **OK** to save the changes.

Map IP Addresses to Users Using the User-ID Agent (Continued)

<p>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.</p>	<ol style="list-style-type: none"> 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: <ul style="list-style-type: none"> • Search Base—The starting point or root context for agent queries, for example: <code>dc=domain1, dc=example, dc=com</code>. • Bind Distinguished Name—The account to use to bind to the directory, for example: <code>cn=admin, ou=IT, dc=domain1, dc=example, dc=com</code>. • 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 <code>objectClass=Person</code>). • 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.
<p>Step 4 (Optional) Enable client probing.</p> <p>Client probing is useful in environments where IP addresses are not tightly bound to users because it ensures that previously mapped addresses are still valid. However, as the total number of learned IP addresses grows, so does the amount of traffic generated. As a best practice, only enable probing on network segments where IP address turnover is high.</p> <p>For more details on the placement of User-ID agents using client probing, refer to Architecting User Identification (User-ID) Deployments.</p>	<ol style="list-style-type: none"> 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. <div data-bbox="824 1220 883 1283" data-label="Image"> </div> <p>For NetBIOS probing to work effectively, each probed client PC must allow port 139 in the Windows firewall and must also have file and printer sharing services enabled. WMI probing is always preferred over NetBIOS whenever possible.</p>
<p>Step 5 Save the configuration.</p>	<p>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.</p>

Map IP Addresses to Users Using the User-ID Agent (Continued)	
<p>Step 6 (Optional) Define the set of users for which you do not need to provide IP address-to-user name mappings, such as kiosk accounts.</p> <p> You can also use the <code>ignore-user</code> list to identify users whom you want to force to authenticate using Captive Portal.</p>	<p>Create an <code>ignore_user_list.txt</code> file and save it to the User-ID Agent folder on the domain server where the agent is installed.</p> <p>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:</p> <pre>SPAdmin SPInstall TFSReport</pre>
<p>Step 7 Configure the firewalls to connect to the User-ID agent.</p>	<p>Complete the following steps on each firewall you want to connect to the User-ID agent to receive user mappings:</p> <ol style="list-style-type: none"> 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 on which the agent will be listening for user mapping requests. This value must match the value configured on the User-ID agent. By default, the port is set to 5007 on the firewall and on newer versions of the User-ID agent. However, some older User-ID agent versions use port 2010 as the default. 5. Make sure that the configuration is Enabled and then click OK. 6. Commit the changes. 7. Verify that the Connected status displays as connected (a green light).
<p>Step 8 Verify that the User-ID agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent.</p>	<ol style="list-style-type: none"> 1. Launch the User-ID agent and select User Identification. 2. Verify that the agent status shows Agent is running. If the Agent is not running, click Start. 3. To verify that the User-ID agent can connect to monitored servers, make sure the Status for each Server is Connected. 4. To verify that the firewalls can connect to the User-ID agent, make sure the Status for each of the Connected Devices is Connected. 5. To verify that the User-ID agent is mapping IP addresses to usernames, select Monitoring and make sure that the mapping table is populated. You can also Search for specific users, or Delete user mappings from the list.

Configure User Mapping Using the PAN-OS Integrated User-ID Agent

The following procedure shows how to configure the PAN-OS integrated agent on the firewall for user mapping. The integrated User-ID agent performs the same tasks as the Windows-based agent with the exception of NetBIOS client probing (WMI probing is supported).

Map IP Addresses to Users Using the Integrated User-ID Agent

<p>Step 1 Create an Active Directory (AD) account for the firewall agent that has the privilege levels required to log in to each service or host you plan to monitor to collect user mapping data.</p>	<ul style="list-style-type: none"> • Windows 2008 or later domain servers—Add the account to the Event Log Readers group. If you are using the on-device User-ID agent, the account must also be a member of the Distributed COM Users Group. • Windows 2003 domain servers—Assign Manage Auditing and Security Logs permissions through group policy. • WMI probing—Make sure the account has rights to read the CIMV2 namespace; by default, Domain Administrator and Server Operator accounts have this permission. • NTLM authentication—Because the firewall must join the domain if you are using Captive Portal NTLM authentication with an on-device User-ID agent, the Windows account you create for NTLM access must have administrative privileges. Note that due to AD restrictions on virtual systems running on the same host, if you have configured multiple virtual systems, only vsys1 will be able to join the domain.
--	--


Map IP Addresses to Users Using the Integrated User-ID Agent (Continued)


Step 2 Define the servers the firewall should monitor to collect IP address-to-user mapping information.

Within the total maximum of 100 monitored servers per firewall, you can define no more than 50 syslog senders for any single virtual system.


Keep in mind that in order to collect all of the required mappings, you must connect to all servers that your users log in to so that the firewall can monitor the security log files on all servers that contain login events.


1. Select **Device > User Identification > User Mapping**.
2. In the Server Monitoring section of the screen, click **Add**.
3. Enter a **Name** and **Network Address** for the server. The network address can be a FQDN or an IP address.
4. Select the **Type** of server.
5. Make sure the **Enabled** check box is selected and then click **OK**.
6. (Optional) To enable the firewall to automatically discover domain controllers on your network using DNS lookups, click **Discover**.




 The auto-discovery feature is for domain controllers only; you must manually add any Exchange servers or eDirectory servers you want to monitor.

7. (Optional) Specify the frequency at which the firewall polls Windows servers for mapping information. This is the interval between the end of the last query and the start of the next query.

 If the query load is high, the observed delay between queries might significantly exceed the specified frequency.


- a. In the **Palo Alto Networks User ID Agent Setup** section, click the Edit  icon.
- b. Select the **Server Monitor** tab and specify the **Server Log Monitor Frequency** in seconds (default is 2, range is 1-3600).

 As a best practice, increase the value in this field to 5 seconds in environments with older domain controllers or high-latency links.




- c. Click **OK** to save the changes.

Step 3 Set the domain credentials for the account the firewall will use to access Windows resources. This is required for monitoring Exchange servers and domain controllers as well as for WMI probing.

1. Edit the Palo Alto Networks User ID Agent Setup section of the screen.
2. On the **WMI Authentication** tab, enter the **User Name** and **Password** for the account that will be used to probe the clients and monitor servers. Enter the user name using the domain\username syntax.



Map IP Addresses to Users Using the Integrated User-ID Agent (Continued)

<p>Step 4 (Optional) Enable WMI probing.</p> <p> The on-device agent does not support NetBIOS probing; it is supported on the Windows-based User-ID agent only.</p>	<ol style="list-style-type: none"> 1. On the Client Probing tab, select the Enable Probing check box. 2. (Optional) If necessary, modify the Probe Interval (in minutes) to ensure it is long enough for the User-ID agent to probe all the learned IP addresses (default is 20, range is 1-1440). This is the interval between the end of the last probe request and the start of the next request. <p> If the request load is high, the observed delay between requests might significantly exceed the specified interval.</p> 3. Make sure the Windows firewall will allow client probing by adding a remote administration exception to the Windows firewall for each probed client.
<p>Step 5 Save the configuration.</p>	<ol style="list-style-type: none"> 1. Click OK to save the User-ID agent setup settings. 2. Click Commit to save the configuration.
<p>Step 6 (Optional) Define the set of users for which you do not need to provide IP address-to-user name mappings, such as kiosk accounts.</p> <p> You can also use the <code>ignore-user</code> list to identify users whom you want to force to authenticate using Captive Portal.</p>	<ol style="list-style-type: none"> 1. Open a CLI session to the firewall. 2. To add the list of user accounts for which you do not want the firewall to perform mapping, run the following command: set user-id-collector ignore-user <value> where <value> is a list of the user accounts to ignore; there is no limit to the number of accounts you can add to the list. Separate entries with a space and do not include the domain name with the username. For example: <pre>set user-id-collector ignore-user SPAdmin SPInstall TFSReport</pre> 3. Commit your changes.
<p>Step 7 Verify the configuration.</p>	<ol style="list-style-type: none"> 1. From the CLI, enter the following command: <pre>show user server-monitor state all</pre> 2. On the Device > User Identification > User Mapping tab in the web interface, verify that the Status of each server you configured for server monitoring is Connected.

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

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.

- ▲ [Configure the Integrated User-ID Agent as a Syslog Listener](#)
- ▲ [Configure the Windows User-ID Agent as a Syslog Listener](#)

Configure the Integrated User-ID Agent as a Syslog Listener


To configure the PAN-OS Integrated User-ID agent to create new user mappings based on syslog monitoring, start by defining Syslog Parse profiles. The User-ID agent uses the profiles to find login 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 format:

```
[Tue Jul 5 13:15:04 2016 CDT] Administrator authentication success User: johndoe1
Source: 192.168.3.212
```

After configuring the Syslog Parse profiles, you specify syslog senders for the User-ID agent to monitor.



The PAN-OS integrated User-ID agent accepts syslogs over SSL and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, always use SSL to listen for syslog messages. However, if you must use UDP, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending UDP traffic to the firewall.

Collect User Mappings from Syslog Senders	
<p>Step 1 Determine whether there is a predefined Syslog Parse profile for your particular syslog senders.</p> <p>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.</p> <p> 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.</p>	<ol style="list-style-type: none"> 1. Install the latest Applications or Applications and Threats update: <ol style="list-style-type: none"> a. Select Device > Dynamic Updates and Check Now. b. Download and Install any new update. 2. Determine which predefined Syslog Parse profiles are available: <ol style="list-style-type: none"> a. Select Device > User Identification > User Mapping click Add in the Server Monitoring section. b. 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.

Collect User Mappings from Syslog Senders (Continued)

Step 2 Define custom Syslog Parse profiles to extract IP address-to-username mapping information from syslog messages.

1. Review the syslog messages that the syslog sender generates to identify the syntax for successful login 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. Specify the **Type** of parsing to extract user mapping information:
 - **Regex Identifier**—Regular expressions.
 - **Field Identifier**—Text strings.

The following steps describe how to configure these parsing types.

Collect User Mappings from Syslog Senders (Continued)

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. For the example message, the regex `(authentication\success){1}` extracts the first {1} instance of the string `authentication success`. The backslash (`\`) before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character.
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:

4. Click **OK** twice to save the profile.

Collect User Mappings from Syslog Senders (Continued)

Step 4 (Field Identifier parsing only) Define string matching patterns.

1. Enter an **Event String** to identify successful login events. For the example message, the string `authentication success` identifies login 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:

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


Collect User Mappings from Syslog Senders (Continued)

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 section.
 2. Enter a **Name** to identify the sender.
 3. Make sure the sender profile is **Enabled** (default).
 4. Select **Syslog Sender** as the server **Type**.
 5. Enter the **Network Address** of the syslog sender (IP address or FQDN).
 6. Select a custom or predefined Syslog Parse profile as a **Filter**.
 7. Select **UDP** or **SSL** (default) as the **Connection Type**.



Use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, always use SSL to listen for syslog messages when using agentless User Mapping on a firewall. However, if you must use UDP, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending UDP traffic to the firewall.

A syslog server using SSL to connect will show a Status of Connected only when there is an active SSL connection. Syslog servers using UDP will not show a Status value.
 8. **(Optional)** If the 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.
-

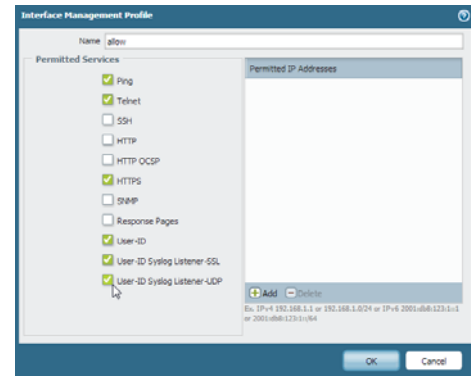
Collect User Mappings from Syslog Senders (Continued)

Step 6 Enable syslog listener services in the management profile associated with the interface used for user mapping.

1. Select **Network > Network Profiles > Interface Mgmt** and then select an interface profile to edit or click **Add** to create a new profile.
2. Select **User-ID Syslog Listener-SSL** and/or **User-ID Syslog Listener-UDP**, depending on the protocols you defined when you set up your Syslog Senders in the Server Monitor list.



The listening ports (514 for UDP and 6514 for SSL) are not configurable; they are enabled through the management service only.



3. Click **OK** to save the interface management profile.

Even after enabling the User-ID Syslog Listener service on the interface, the interface will only accept syslog connections from servers that have a corresponding entry in the User-ID monitored servers configuration. Connections or messages from servers that are not on the list will be discarded.
4. Assign the Interface Management profile to the interface that the firewall uses to collect user mappings:
 - a. Select **Network > Interfaces** and edit the interface.
 - b. Select **Advanced > Other info**, select the Interface Management Profile you just added, and click **OK**.
5. **Commit** your changes.

Collect User Mappings from Syslog Senders (Continued)

Step 7 Verify the configuration by [logging in to the firewall CLI](#) and running the following commands:

To see the status of a particular syslog sender:

```
admin@PA-5050> show user server-monitor state Syslog2
UDP Syslog Listener Service is enabled
SSL Syslog Listener Service is enabled

Proxy: Syslog2(vsys: vsys1)   Host: Syslog2(10.5.204.41)
number of log messages       : 1000
number of auth. success messages : 1000
number of active connections  : 0
total connections made       : 4
```

To see how many log messages came in from syslog senders and how many entries were successfully mapped:

```
admin@PA-5050> show user server-monitor statistics

Directory Servers:
Name           TYPE      Host           Vsys    Status
-----
AD             AD        10.2.204.43    vsys1   Connected

Syslog Servers:
Name           Connection Host           Vsys    Status
-----
Syslog1       UDP         10.5.204.40    vsys1   N/A
Syslog2       SSL         10.5.204.41    vsys1   Not connected
```

To see how many user mappings were discovered through syslog senders:

```
admin@PA-5050> show user ip-user-mapping all type SYSLOG

IP           Vsys  From      User                               IdleTimeout(s) M
axTimeout(s)
-----
192.168.3.8  vsys1  SYSLOG    acme\jreddick                     2476           2
476
192.168.5.39 vsys1  SYSLOG    acme\jdonaldson                   2480           2
480
192.168.2.147 vsys1  SYSLOG    acme\ccrisp                       2476           2
476
192.168.2.175 vsys1  SYSLOG    acme\jjaso                         2476           2
476
192.168.4.196 vsys1  SYSLOG    acme\jblevins                     2480           2
480
192.168.4.103 vsys1  SYSLOG    acme\bmoos                        2480           2
480
192.168.2.193 vsys1  SYSLOG    acme\esogard                      2476           2
476
192.168.2.119 vsys1  SYSLOG    acme\acallaspo                   2476           2
476
192.168.3.176 vsys1  SYSLOG    acme\jlowrie                      2478           2
478

Total: 9 users
```

Configure the Windows User-ID Agent as a Syslog Listener

To configure the Windows-based User-ID agent to create new user mappings based on syslog monitoring, start by defining Syslog Parse profiles. The User-ID agent uses the profiles to find login 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 format:

```
[Tue Jul 5 13:15:04 2016 CDT] Administrator authentication success User:johndoe1
Source:192.168.3.212
```

After configuring the Syslog Parse profiles, you specify syslog senders for the User-ID agent to monitor.



The Windows User-ID agent accepts syslogs over TCP and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, use TCP instead of UDP. In either case, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending syslogs to the User-ID agent.

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

Step 1 Define custom Syslog Parse profiles to filter syslog messages for successful login events.

1. Review the syslogs generated by the authenticating service to identify the syntax of the login events. This enables you to create the matching patterns that will allow the firewall to identify and extract the authentication events from the syslogs.



While reviewing the syslogs, also determine whether the domain name is included in the log entries. If the authentication logs do not contain domain information, consider defining a default domain name when adding the syslog sender to the monitored servers list in [Step 4](#).

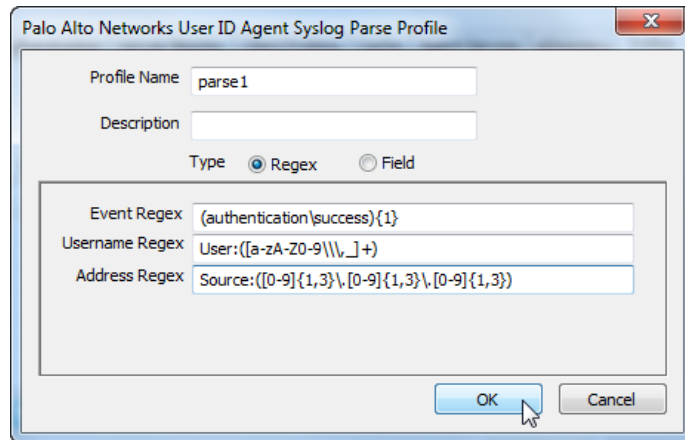
2. Open the Windows **Start** menu and select **User-ID Agent**.
3. Select **User Identification > Setup** and **Edit** the Setup.
4. Select **Syslog, Enable Syslog Service**, and **Add** a Syslog Parse profile.
5. Enter a **Profile Name** and **Description**.
6. Select the **Type** of parsing to find login events in syslog messages:
 - **Regex**—Regular expressions.
 - **Field**—Text strings.

The following steps describe how to configure these parsing types.

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

Step 2 (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** to identify successful login events.
 For the example message, the regex (authentication\succes*){1} extracts the first {1} instance of the string authentication success. The backslash before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character.
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.

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

Step 3 (Field Identifier parsing only) Define string matching patterns.

1. Enter an **Event String** to identify successful login events. For the example message, the string `authentication success` identifies login 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:

The screenshot shows a configuration window titled "Palo Alto Networks User ID Agent Syslog Parse Profile". The window contains the following fields and options:

- Profile Name: `parse2`
- Description: (empty)
- Type: Field (selected), Regex
- Event String: `authentication success`
- Username Prefix: `User:`
- Username Delimiter: `\s`
- Address Prefix: `Source:`
- Address Delimiter: `\n`
- Buttons: OK, Cancel

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

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

- Step 4** 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 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.
 6. Select the Syslog Parse profile you configured as a **Filter**.
 7. Click **OK** to save the settings.
 8. **Commit** your changes to the User-ID agent configuration.

Step 5 Verify the configuration by [logging in to the firewall CLI](#) and running the following commands:

To see the status of a particular syslog sender:

```
admin@PA-5050> show user server-monitor state Syslog2
UDP Syslog Listener Service is enabled
SSL Syslog Listener Service is enabled

Proxy: Syslog2(vsys: vsys1)      Host: Syslog2(10.5.204.41)
      number of log messages      : 1000
      number of auth. success messages : 1000
      number of active connections  : 0
      total connections made       : 4
```

To see how many log messages came in from syslog senders and how many entries were successfully mapped:

```
admin@PA-5050> show user server-monitor statistics

Directory Servers:
Name                TYPE      Host           Vsys    Status
-----
AD                  AD        10.2.204.43   vsys1   Connected

Syslog Servers:
Name                Connection Host           Vsys    Status
-----
Syslog1             UDP          10.5.204.40   vsys1   N/A
Syslog2             SSL          10.5.204.41   vsys1   Not connected
```

To see how many user mappings were discovered through syslog senders:

```
admin@PA-5050> show user ip-user-mapping all type SYSLOG

IP          Vsys  From  User                               IdleTimeout(s) M
axTimeout(s)
-----
192.168.3.8 vsys1 SYSLOG acme\jreddick                     2476           2
476
192.168.5.39 vsys1 SYSLOG acme\jdonaldson                   2480           2
480
192.168.2.147 vsys1 SYSLOG acme\ccrisp                       2476           2
476
192.168.2.175 vsys1 SYSLOG acme\jjaso                        2476           2
476
192.168.4.196 vsys1 SYSLOG acme\jblevins                     2480           2
480
192.168.4.103 vsys1 SYSLOG acme\bmoss                        2480           2
480
192.168.2.193 vsys1 SYSLOG acme\esogard                      2476           2
476
192.168.2.119 vsys1 SYSLOG acme\acallaspo                   2476           2
476
192.168.3.176 vsys1 SYSLOG acme\jlowrie                     2478           2
478

Total: 9 users
```

Map IP Addresses to User Names Using Captive Portal

If the firewall receives a request from a zone that has User-ID enabled and the source IP address does not have any user data associated with it yet, it checks its Captive Portal policy for a match to determine whether to perform authentication. This is useful in environments where you have clients that are not logged in to your domain servers, such as Linux clients. This user mapping method is triggered only for web traffic (HTTP or HTTPS) that matches a security rule/policy, but that has not been mapped using a different method.

- ▲ [Captive Portal Authentication Methods](#)
- ▲ [Captive Portal Modes](#)
- ▲ [Configure Captive Portal](#)

Captive Portal Authentication Methods

Captive Portal uses the following methods to obtain user data from the client when a request matches a Captive Portal policy:

Authentication Method	Description
NT LAN Manager (NTLM)	The firewall uses an encrypted challenge-response mechanism to obtain the user's credentials from the browser. When configured properly, the browser will provide the credentials to the firewall transparently without prompting the user, but will display a prompt for credentials if necessary. If the browser cannot perform NTLM or if NTLM authentication fails, the firewall falls back to web form or client certificate authentication, depending on your Captive Portal configuration. By default, IE supports NTLM. Firefox and Chrome can be configured to use it. You cannot use NTLM to authenticate non-Windows clients.
Web Form	Requests are redirected to a web form for authentication. You can configure Captive Portal to use a local user database, RADIUS, LDAP, or Kerberos to authenticate users. Although users will always be prompted for credentials, this authentication method works with all browsers and operating systems.
Client Certificate Authentication	Prompts the browser to present a valid client certificate for authenticating the user. To use this method, you must provision client certificates on each user system and install the trusted CA certificate used to issue those certificates on the firewall. This is the only authentication method that enables transparent authentication for Mac OS and Linux clients.

Captive Portal Modes

The Captive Portal mode defines how web requests are captured for authentication:

Mode	Description
Transparent	The firewall intercepts the browser traffic per the Captive Portal rule and impersonates the original destination URL, issuing an HTTP 401 to invoke authentication. However, because the firewall does not have the real certificate for the destination URL, the browser will display a certificate error to users attempting to access a secure site. Therefore you should only use this mode when absolutely necessary, such as in Layer 2 or virtual wire deployments.
Redirect	The firewall intercepts unknown HTTP or HTTPS sessions and redirects them to a Layer 3 interface on the firewall using an HTTP 302 redirect in order to perform authentication. This is the preferred mode because it provides a better end-user experience (no certificate errors). However, it does require additional Layer 3 configuration. Another benefit of the Redirect mode is that it provides for the use of session cookies, which enable the user to continue browsing to authenticated sites without requiring re-mapping each time the time outs expire. This is especially useful for users who roam from one IP address to another (for example, from the corporate LAN to the wireless network) because they will not need to re-authenticate upon IP address change as long as the session stays open. In addition, if you plan to use NTLM authentication, you must use Redirect mode because the browser will only provide credentials to trusted sites.

Configure Captive Portal

The following procedure shows how to configure Captive Portal using the PAN-OS integrated User-ID agent to redirect requests that match a Captive Portal policy to a Layer 3 interface on the firewall.



If you plan to use Captive Portal without using the other User-ID functions (user mapping and group mapping), you do not need to configure an agent.

Configure Captive Portal Using the PAN-OS Integrated User-ID Agent

Step 1	Make sure the firewall has a route to the servers it will monitor to gather user data (for example, your Domain Controllers and your Exchange servers).	The firewall must be able to communicate with the servers over a designated interface (the management [MGT] interface by default), so make sure that the network where your directory servers reside is accessible from that interface. If this configuration doesn't work in your environment, you must configure Captive Portal using the Window-based User-ID agent.
Step 2	Make sure DNS is configured to resolve your Domain Controller addresses.	To verify proper resolution, ping the server FQDN. For example: admin@PA-200> ping host dc1.acme.com

Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued)

Step 3 (Redirect mode only) Create a Layer 3 interface to which to redirect Captive Portal requests.

1. Create a management profile to enable the interface to display Captive Portal response pages:
 - a. Select **Network > Interface Mgmt** and click **Add**.
 - b. Enter a **Name** for the profile, select **Response Pages**, and then click **OK**.
2. Create the Layer 3 interface. Be sure to attach the management profile you just created (on the **Advanced > Other Info** tab of the Ethernet Interface dialog).
3. Create a DNS “A” record that maps the IP address you configured on the Layer 3 interface to an intranet host name (that is, a hostname that does not have a dot in the name, such as ntlmhost).

Step 4 (Redirect mode only) To transparently redirect users without displaying certificate errors, install a certificate that matches the IP address of the interface to which you are redirecting requests. You can either generate a self-signed certificate or import a certificate that is signed by an external CA.



When setting up Captive Portal for the first time, imported certificates may not work. If you plan to use an imported certificate, complete the initial configuration without specifying a **Server Certificate**. After you get Captive Portal working, you can go back and switch to the imported certificate.



To use a self-signed certificate, you must first create a root CA certificate and then use that CA to sign the certificate you will use for Captive Portal as follows:

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 Captive Portal, click **Generate**. Enter a **Certificate Name** and enter the DNS name of the intranet host for the interface as the **Common Name**. In the **Signed By** field, select the CA you created in the previous step. Add an IP address attribute and specify the IP address of the Layer 3 interface to which you will be redirecting requests. **Generate** the certificate.

Name	Subject	Issuer	CA	Key	Expires	Status
<input type="checkbox"/> Acme-CertAutho...	ca.acme.com	ca.acme.com	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Nov 19 19:08:32 2013 GMT	valid
<input type="checkbox"/> CP-Cert	cp.acme.com	ca.acme.com	<input type="checkbox"/>	<input checked="" type="checkbox"/>	Nov 19 19:10:11 2013 GMT	valid

3. To configure clients to trust the certificate, select the CA certificate on the **Device Certificates** tab and click **Export**. You must then import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory Group Policy Object (GPO).

Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued)


<p>Step 5 Set up an authentication mechanism to use when the web form is invoked. Note that even if you plan to use NTLM, you must also set up a secondary authentication mechanism that can be used if NTLM authentication fails or if the user agent does not support it.</p> <p> Best Practices:</p> <ul style="list-style-type: none"> • If using RADIUS to authenticate users from the web form, be sure to enter a RADIUS domain. This will be used as the default domain if users don't supply one upon login. • If using AD to authenticate users from the web form, make sure to enter sAMAccountName as the LogonAttribute. 	<ol style="list-style-type: none"> 1. Configure the firewall to connect to the authentication service you plan to use so that it can access the authentication credentials. <ul style="list-style-type: none"> • If you plan to authenticate using LDAP, Kerberos, or RADIUS, you must create a server profile that instructs the firewall how to connect to the service and access the authentication credentials for your users. Select Device > Server Profiles and add a new profile for the specific service you will be accessing. • If you plan to use local database authentication, you must first create the local database. Select Device > Local User Database and add the users and groups to be authenticated. 2. Create an authentication profile that references the server profile or local user database you just created. Select Device > Authentication Profile and add a new profile for use with Captive Portal. For details on creating a specific type of authentication profile, refer to the online help.
<p>Step 6 (Optional) Set up client certificate authentication. Note that you do not need to set up both an authentication profile and a client certificate profile to enable Captive Portal. If you configure both, the user will be required to authenticate using both methods.</p> <p> For details on other certificate profile fields, such as whether to use CRL or OCSP, refer to the online help.</p>	<ol style="list-style-type: none"> 1. Use a root CA certificate to generate a client certificate for each user who will authenticate to Captive Portal. The CA in this case is usually your enterprise CA, not the firewall. 2. Export the CA certificate in Base64 format to a system that the firewall can access. 3. Import the root CA certificate from the CA that generated the client certificates onto the firewall: <ol style="list-style-type: none"> a. Select Device > Certificate Management > Certificates > Device Certificates and click Import. b. Enter a Certificate Name that identifies the certificate as your client CA certificate. c. Browse to the Certificate File you downloaded from the CA. d. Select Base64 Encoded Certificate (PEM) as the File Format and then click OK. e. Select the certificate you just imported on the Device Certificates tab to open it. f. Select Trusted Root CA and then click OK. 4. Create the client certificate profile that you will use when you configure Captive Portal. <ol style="list-style-type: none"> a. Select Device > Certificates > Certificate Management > Certificate Profile, click Add, and enter a profile Name. b. In the Username Field drop-down, select the certificate field that contains the user's identity information. c. In the CA Certificates field, click Add, select the Trusted Root CA certificate you just imported and then click OK.

Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued)

Step 7 Enable [NT LAN Manager \(NTLM\)](#) authentication.



When using the on-device User-ID agent, the firewall must be able to successfully resolve the DNS name of your Domain Controller in order for the firewall to join the domain. The credentials you supply here will be used to join the firewall to the domain upon successful DNS resolution.

1. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User ID Agent Setup section of the screen.
2. On the **NTLM** tab, select the **Enable NTLM authentication processing** check box.
3. Enter the **NTLM Domain** against which the User-ID agent on the firewall should check NTLM credentials.
4. Enter the **Admin User Name** and **Password** for the Active Directory account you created in [Step 1 in Map IP Addresses to Users Using the Integrated User-ID Agent](#) for NTLM authentication.
 Do not include the domain in the **Admin User Name** field. Otherwise, the firewall will fail to join the domain. Palo Alto Networks recommends that you use a User-ID agent account that is separate from your firewall administrator account.
5. You don't need to configure any other settings for the User-ID agent: click **OK**.

Configure Captive Portal Using the PAN-OS Integrated User-ID Agent (Continued)

Step 8 Configure the Captive Portal settings.

1. Select **Device > User Identification > Captive Portal Settings** and Edit the **Captive Portal** section of the screen.
2. Make sure the **Enabled** check box is selected.
3. Set the **Mode**. This example shows how to set up **Redirect** mode.
4. (Redirect mode only) Select the **Server Certificate** the firewall should use to redirect requests over SSL. This is the certificate you created in [Step 4](#).
5. (Redirect mode only) Specify the **Redirect Host**, which is the intranet hostname that resolves to the IP address of the Layer 3 interface to which you are redirecting requests, as specified in [Step 3](#).
6. Select the authentication method to use if NTLM fails (or if you are not using NTLM):
 - If you are using LDAP, Kerberos, RADIUS, or local database authentication, select the **Authentication Profile** you created in [Step 5](#).
 - If you are using client certificate authentication, select the **Certificate Profile** you created in [Step 6](#).

The screenshot displays the 'Captive Portal' configuration page. At the top, the 'Enabled' checkbox is checked. Below it, 'Idle Timer (min)' is set to 15 and 'Expiration (min)' is set to 60. The 'Server Certificate' dropdown is set to 'None' and the 'Authentication Profile' dropdown is set to 'Corp'. The 'Mode' section has 'Transparent' unselected and 'Redirect' selected. The 'Session Cookie' section has 'Enable' checked, 'Timeout (min)' set to 1440, and 'Roaming' checked. The 'Redirect Host' field contains 'corp-fw'. The 'Certificate Authentication' section has 'Certificate Profile' set to 'None'. The 'NTLM Authentication' section has 'Attempts' set to 1 and 'Timeout (sec)' set to 2.

7. Click **OK** to save your settings.
8. Click **Commit** to save the Captive Portal configuration.

Configure User Mapping for Terminal Server Users

Individual terminal server users appear to have the same IP address and therefore an IP address-to-username mapping is not sufficient to identify a specific user. To enable identification of specific users on Windows-based terminal servers, the Palo Alto Networks Terminal Services agent (TS agent) allocates a port range to each user. It then notifies every connected firewall about the allocated port range, which allows the firewall to create an IP address-port-user mapping table and enable user- and group-based security policy enforcement. For non-Windows terminal servers, you can configure the User-ID XML API to extract user mapping information.

The following sections describe how to configure user mapping for terminal server users:

- ▲ [Configure the Palo Alto Networks Terminal Server Agent for User Mapping](#)
- ▲ [Retrieve User Mappings from a Terminal Server Using the User-ID XML API](#)

Configure the Palo Alto Networks Terminal Server Agent for User Mapping

Use the following procedure to install the TS agent on the terminal server. You must install the TS agent on all terminal servers that your users log in to in order to successfully map all your users.



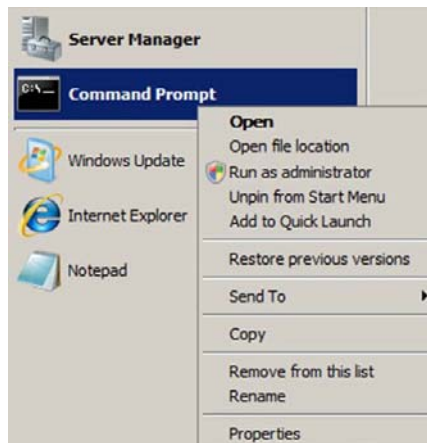
For information about the supported terminal servers supported by the TS Agent, refer to “Operating System (OS) Compatibility TS Agent” in the Terminal Services Agent Release Notes, which are available on the Palo Alto Networks [Software Updates](#) page.

Install the Windows Terminal Server Agent

<p>Step 1 Download the TS Agent installer.</p>	<ol style="list-style-type: none"> 1. Log in to the Palo Alto Networks Support site. 2. Select Software Updates from the Manage Devices section. 3. Scroll to the Terminal Services Agent section and Download the version of the agent you want to install. 4. Save the <code>TaInstall164.x64-x.x.x-xx.msi</code> or <code>TaInstall-x.x.x-xx.msi</code> file (be sure to select the appropriate version based on whether the Windows system is running a 32-bit OS or a 64-bit OS) on the system(s) where you plan to install the agent.
---	--

Install the Windows Terminal Server Agent (Continued)

Step 2 Run the installer as an administrator.



1. To launch a command prompt as an administrator, click Start and right-click **Command Prompt** and then select **Run as administrator**.
2. From the command line, run the .msi file you downloaded. For example, if you saved the .msi file to the Desktop you would enter the following:


```
C:\Users\administrator.acme>cd Desktop
C:\Users\administrator.acme\Desktop>TaInstall-6.0.0-1.msi
```
3. Follow the setup prompts to install the agent using the default settings. By default, the agent gets installed to the C:\Program Files (x86)\Palo Alto Networks\Terminal Server Agent folder, but you can **Browse** to a different location.
4. When the installation completes, **Close** the setup window.



If you are upgrading to a TS Agent version that has a newer driver than the existing installation, the installation wizard prompts you to reboot the system after upgrading in order to use the new driver.

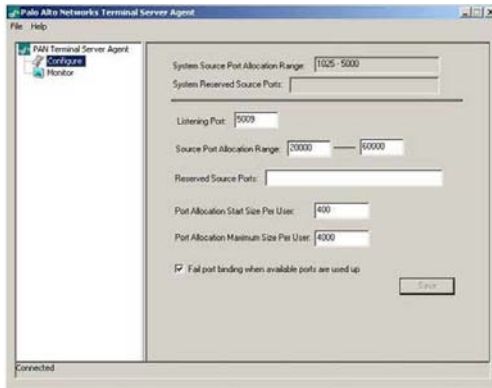
Step 3 Launch the Terminal Server Agent application.

Click Start and select **Terminal Server Agent**.



Install the Windows Terminal Server Agent (Continued)

Step 4 Define the range of ports for the TS Agent to allocate to end users.



The **System Source Port Allocation Range** and **System Reserved Source Ports** fields specify the range of ports that will be allocated to non-user sessions. Make sure the values specified in these fields do not overlap with the ports you designate for user traffic. These values can only be changed by editing the corresponding Windows registry settings.

1. Select **Configure**.
2. Set the **Source Port Allocation Range** (default 20000-39999). This is the full range of port numbers that the TS Agent will allocate for user mapping. The port range you specify cannot overlap with the **System Source Port Allocation Range**.
3. (Optional) If there are ports/port ranges within the source port allocation that you do not want the TS Agent to allocate to user sessions, specify them as **Reserved Source Ports**. To include multiple ranges, use commas with no spaces, for example: 2000-3000, 3500, 4000-5000.
4. Specify the number of ports to allocate to each individual user upon login to the terminal server in the **Port Allocation Start Size Per User** field (default 200).
5. Specify the **Port Allocation Maximum Size Per User**, which is the maximum number of ports the Terminal Server agent can allocate to an individual user.
6. Specify whether to continue processing traffic from the user if the user runs out of allocated ports. By default, the **Fail port binding when available ports are used up** is selected, which indicates that the application will fail to send traffic when all ports are used. To enable users to continue using applications when they run out of ports, clear this check box. Keep in mind that this traffic may not be identified with User-ID.

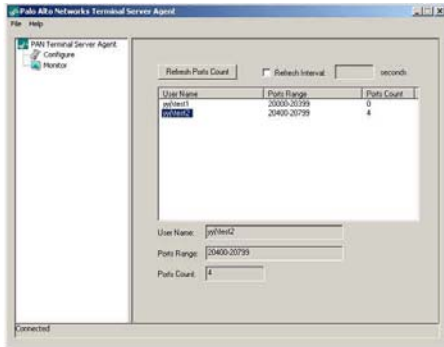
Step 5 Configure the firewalls to connect to the Terminal Server agent.

Complete the following steps on each firewall you want to connect to the Terminal Server agent to receive user mappings:

1. Select **Device > User Identification > Terminal Server Agents** and click **Add**.
2. Enter a **Name** for the Terminal Server agent.
3. Enter the IP address of the Windows **Host** on which the Terminal Server agent is installed.
4. Enter the **Port** number on which the agent will listen for user mapping requests. This value must match the value configured on the Terminal Server agent. By default, the port is set to 5009 on the firewall and on the agent. If you change it here, you must also change the **Listening Port** field on the Terminal Server agent **Configure** screen.
5. Make sure that the configuration is **Enabled** and then click **OK**.
6. **Commit** the changes.
7. Verify that the **Connected** status displays as connected (a green light).

Install the Windows Terminal Server Agent (Continued)

Step 6 Verify that the Terminal Server agent is successfully mapping IP addresses to usernames and that the firewalls can connect to the agent.



1. Launch the Terminal Server agent and verify that the firewalls can connect by making sure the **Connection Status** for each of Device in the Connection List is **Connected**.
2. To verify that the Terminal Server agent is successfully mapping port ranges to usernames, select **Monitoring** and make sure that the mapping table is populated.

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

The User-ID XML API is a RESTful API that uses standard HTTP requests to send and receive data. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports RESTful services.

To enable a non-Windows terminal server to send user mapping information directly to the firewall, create scripts that extract the user login and logout events and use them for input to the User-ID XML API request format. Then define the mechanisms for submitting the XML API request(s) to the firewall using cURL or wget and providing the firewall's API key for secure communication. Creating user mappings from multi-user systems such as terminal servers requires use of the following API messages:

- **<multiusersystem>**—Sets up the configuration for an XML API Multi-user System on the firewall. This message allows for definition of the terminal server IP address (this will be the source address for all users on that terminal server). In addition, the **<multiusersystem>** setup message specifies the range of source port numbers to allocate for user mapping and the number of ports to allocate to each individual user upon login (called the *block size*). If you want to use the default source port allocation range (1025-65534) and block size (200), you do not need to send a **<multiusersystem>** setup event to the firewall. Instead, the firewall will automatically generate the XML API Multi-user System configuration with the default settings upon receipt of the first user login event message.
- **<blockstart>**—Used with the **<login>** and **<logout>** messages to indicate the starting source port number allocated to the user. The firewall then uses the block size to determine the actual range of port numbers to map to the IP address and username in the login message. For example, if the **<blockstart>** value is 13200 and the block size configured for the multi-user system is 300, the actual source port range allocated to the user is 13200 through 13499. Each connection initiated by the user should use a unique source port number within the allocated range, enabling the firewall to identify the user based on its IP address-port-user mappings for enforcement of user- and group-based security policy rules. When a user exhausts all the ports allocated, the terminal server must send a new **<login>** message allocating a new port range for the user so that the firewall can update the IP address-port-user mapping. In addition, a single


username can have multiple blocks of ports mapped simultaneously. When the firewall receives a <logout> message that includes a <blockstart> parameter, it removes the corresponding IP address-port-user mapping from its mapping table. When the firewall receives a <logout> message with a username and IP address, but no <blockstart>, it removes the user from its table. And, if the firewall receives a <logout> message with an IP address only, it removes the multi-user system and all mappings associated with it.





The XML files that the terminal server sends to the firewall can contain multiple message types and the messages do not need to be in any particular order within the file. However, upon receiving an XML file that contains multiple message types, the firewall will process them in the following order: multiusersystem requests first, followed by logins, then logouts.

The following workflow provides an example of how to use the User-ID XML API to send user mappings from a non-Windows terminal server to the firewall.


Use the User-ID XML API to Map Non-Windows Terminal Services Users

<p>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).</p> <p> Any special characters in the password must be URL/percent-encoded.</p>	<p>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:</p> <pre>https://<Firewall-IPAddress>/api/?type=keygen&user=<username>&password=<password></pre> <p>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:</p> <pre>https://10.1.2.5/api/?type=keygen&user=admin&password=admin</pre> <p>The firewall responds with a message containing the key, for example:</p> <pre><response status="success"> <result> <key>k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UOn1ZRg=</key> </result> </response></pre>
--	--

Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued)

<p>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.</p> <p>If the terminal services agent does not send a setup message, the firewall will automatically create a terminal server agent configuration using the following default settings upon receipt of the first login message:</p> <ul style="list-style-type: none"> • Default port range: 1025 to 65534 • Per user block size: 200 • Maximum number of multi-user systems: 1,000 	<p>The following shows a sample setup message:</p> <pre><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></pre> <p>where <code>entry ip</code> specifies the IP address assigned to terminal server users, <code>startport</code> and <code>endport</code> specify the port range to use when assigning ports to individual users, and <code>blocksize</code> specifies the number of ports to assign to each user. The maximum <code>blocksize</code> is 4000 and each multi-user system can allocate a maximum of 1,000 blocks.</p> <p>If you define a custom <code>blocksize</code> 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.</p>
<p>Step 3 Create a script that will extract the login events and create the XML input file to send to the firewall.</p> <p>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 <code>blockstart</code> values would be 1000, 1200, 1400, 1600, or 1800. <code>Blockstart</code> values of 1001, 1300, or 1850 would be unacceptable because some of the port numbers in the range would be left unused.</p> <p> The login event payload that the terminal server sends to the firewall can contain multiple login events.</p>	<p>The following shows the input file format for a user-ID XML login event:</p> <pre><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></pre> <p>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 <code>jparker</code> for policy enforcement.</p> <p> Each multi-user system can allocate a maximum of 1,000 port blocks.</p>

Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued)

<p>Step 4 Create a script that will extract the logout events and create the XML input file to send to the firewall.</p> <p>Upon receipt of a <code>logout</code> event message with a <code>blockstart</code> parameter, the firewall removes the corresponding IP address-port-user mapping. If the <code>logout</code> message contains a username and IP address, but no <code>blockstart</code> parameter, the firewall removes all mappings for the user. If the <code>logout</code> message contains an IP address only, the firewall removes the multi-user system and all associated mappings.</p>	<p>The following shows the input file format for a User-ID XML logout event:</p> <pre><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></pre> <p> You can also clear the multiuser system entry from the firewall using the following CLI command: <code>clear xml-api multiusersystem</code></p>
<p>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.</p>	<p>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 <code>jjaso</code> is mapped to a source network address translation (SNAT) value of <code>10.1.1.23:20000-20099</code>, the script you create should include the following:</p> <pre>[root@ts1 ~]# iptables -t nat -A POSTROUTING -m owner --uid-owner jjaso -p tcp -j SNAT --to-source 10.1.1.23:20000-20099</pre> <p>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:</p> <pre>[root@ts1 ~]# iptables -t nat -D POSTROUTING 1</pre>
<p>Step 6 Define how to package the XML input files containing the setup, login, and logout events into <code>wget</code> or <code>cURL</code> messages for transmission to the firewall.</p>	<p>To apply the files to the firewall using wget:</p> <pre>> wget --post file <filename> "https://<Firewall-IPaddress>/api/?type=user-id&key=<key>&file-name=<input_filename.xml>&client=wget&vsys=<VSYS_name>"</pre> <p>For example, the syntax for sending an input file named <code>login.xml</code> to the firewall at <code>10.2.5.11</code> using key <code>k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9U0n1ZRg</code> using <code>wget</code> would look as follows:</p> <pre>> wget --post file login.xml "https://10.2.5.11/api/?type=user-id&key=k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9U0n1ZRg&file-name=login.xml&client=wget&vsys=vsys1"</pre> <p>To apply the file to the firewall using cURL:</p> <pre>> curl --form file=@<filename> https://<Firewall-IPaddress>/api/?type=user-id&key=<key>&vsys=<VSYS_name></pre> <p>For example, the syntax for sending an input file named <code>login.xml</code> to the firewall at <code>10.2.5.11</code> using key <code>k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9U0n1ZRg</code> using <code>cURL</code> would look as follows:</p> <pre>> curl --form file@login.xml "https://10.2.5.11/api/?type=user-id&key=k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9U0n1ZRg&vsys=vsys1"</pre>

Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued)

Step 7 Verify that the firewall is successfully receiving login events from the terminal servers.

Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:

To verify if the terminal server is connecting to the firewall over XML:

```
admin@PA-5050> show user xml-api multiusersystem
```

Host	Vsys	Users	Blocks
10.5.204.43	vsys1	5	2

To verify that the firewall is receiving mappings from a terminal server over XML:

```
admin@PA-5050> show user ip-port-user-mapping all
```

```
Global max host index 1, host hash count 1
```

```
XML API Multi-user System 10.5.204.43
```

```
Vsys 1, Flag 3
```

```
Port range: 20000 - 39999
```

```
Port size: start 200; max 2000
```

```
Block count 100, port count 20000
```

```
20000-20199: acme\administrator
```

```
Total host: 1
```

Send User Mappings to User-ID Using the XML API


Although the User-ID functionality provides many out-of-the box methods for obtaining user mapping information, you may have some applications or devices that capture user information that cannot be natively integrated with User-ID. In this case you can use the User-ID XML API to create custom scripts that allow you to leverage existing user data and send it to the User-ID agent or directly to the firewall.

The User-ID XML API is a RESTful API that uses standard HTTP requests to send and receive data. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports RESTful services. To leverage user data from an existing system—such as a custom application developed internally or another device that is not supported by one of the existing user mapping mechanisms—you can create custom scripts to extract the data and send it to the firewall or the User-ID agent using the XML API.


To enable an external system to send user mapping information to the User-ID agent or directly to the firewall, you can create scripts that extract the user login and logout events and use them for input to the User-ID XML API request format. Then define the mechanisms for submitting the XML API request(s) to the firewall using cURL or wget using the firewall's API key for secure communication. For more details, refer to the [PAN-OS XML API Usage Guide](#).

Configure a Firewall to Share User Mapping Data with Other Firewalls


Because policy is local to each firewall, each firewall needs current user mapping and group mapping information to accurately enforce security policy by user and group. However, you can configure one firewall to collect all the mapping information and act as a User-ID agent to share that information with other firewalls. The redistribution firewall can share only the information it collects using local methods (for example, the PAN-OS integrated User-ID agent or Captive Portal), not the information collected from the Windows-based User-ID and Terminal Services agents. You configure the receiving firewalls to retrieve the mapping information from the redistribution firewall; they don't need to communicate directly with domain servers.



The following procedure describes how to set up redistribution of User-ID information.

Configure a Firewall to Redistribute User Mappings	
<p>Step 1 Configure the redistribution firewall.</p> <p> User-ID configurations apply to a single virtual system only. To redistribute User-ID mappings from multiple virtual systems, you must configure the user mapping settings on each virtual system separately, using a unique pre-shared key in each configuration.</p>	<ol style="list-style-type: none"> 1. Select Device > User Identification > User Mapping and edit the Palo Alto Networks User-ID Agent Setup section. 2. Select Redistribution. 3. Enter a Collector Name. 4. Enter and confirm the Pre-Shared Key that will enable other firewalls to connect to this firewall to retrieve user mapping information. 5. Click OK to save the redistribution configuration.
<p>Step 2 Create an interface management profile that enables the User-ID service and attach it to the interface that the other firewalls will connect to in order to retrieve user mappings.</p>	<ol style="list-style-type: none"> 1. Select Network > Network Profiles > Interface Mgmt and click Add. 2. Enter a Name for the profile and then select the Permitted Services. At a minimum, select User-ID Service and HTTPS. 3. Click OK to save the profile. 4. Select Network > Interfaces > Ethernet and select the interface you plan to use for redistribution. 5. On the Advanced > Other Info tab, select the Management Profile you just created. 6. Click OK and Commit.

Configure a Firewall to Redistribute User Mappings (Continued)

<p>Step 3 Configure the other firewalls to retrieve user mappings from the redistribution firewall.</p> <p> If the redistribution firewall has multiple virtual systems configured for redistribution, make sure you are using the pre-shared key that corresponds to the virtual system from which you want this firewall to retrieve User-ID mappings.</p>	<p>Perform the following steps on each firewall that you want to be able to retrieve user mappings:</p> <ol style="list-style-type: none"> 1. Select Device > User Identification > User-ID Agents. 2. Click Add and enter a User-ID agent Name for the redistribution firewall. 3. Enter the hostname or IP address of the firewall interface that you configured for redistribution in the Host field. 4. Enter 5007 as the Port number on which the redistribution firewall will listen for User-ID requests. 5. Enter the Collector Name that you specified in the redistribution firewall configuration (Step 1-3). 6. Enter and confirm the Collector Pre-Shared Key. The key value you enter here must match the value configured on the redistribution firewall (Step 1-4). 7. (Optional) If you are using the redistribution firewall to retrieve group mappings in addition to user mappings, select the Use as LDAP Proxy check box. 8. (Optional) If you are using the redistribution firewall for Captive Portal authentication, select the Use for NTLM Authentication check box. 9. Make sure that the configuration is Enabled and then click OK. 10. Commit the changes.
<p>Step 4 Verify the configuration.</p>	<p>On the User-ID Agents tab, verify that the redistribution firewall entry you just added shows a green icon in the Connected column. If a red icon appears, check traffic logs (Monitor > Logs > Traffic) to identify the issue. You can also check to see if any user mapping data has been received by running the following operational commands from the CLI:</p> <pre>show user ip-user-mapping (to view user mapping information on the dataplane)</pre> <pre>show user ip-user-mapping-mp (to view mappings on the management plane).</pre>

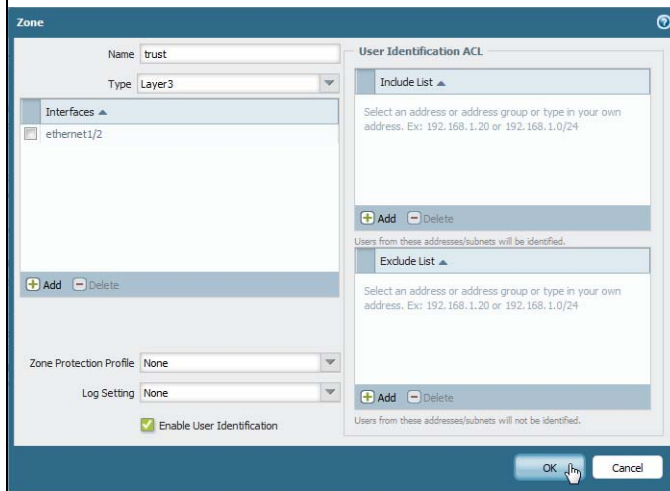
Enable User- and Group-Based Policy

In order to enable security policy based on user and/or group, you must enable User-ID for each zone that contains users you want to identify. You can then define policies that allow or deny traffic based on user name or group membership. Additionally, you can create Captive Portal policies to enable identification for IP addresses that do not yet have any user data associated with them.

Enable User- and Group-Based Policy

Step 1 Enable User-ID on the source zones that contain the users that will send requests that require user-based access controls.

1. Select **Network > Zones**.
2. Click the **Name** of the zone in which you want to enable User-ID to open the Zone dialog.
3. Select the **Enable User Identification** check box and then click **OK**.



Enable User- and Group-Based Policy (Continued)

Step 2 Create security policies based on user and/or group.



As a best practice, create policies based on group rather than user whenever possible. This prevents you from having to continually update your policies (which requires a commit) whenever your user base changes.

1. After configuring User-ID, you will be able to choose a user name or group name when defining the source or destination of a security rule:
 - a. Select **Policies > Security** and click **Add** to create a new policy or click on an existing policy rule name to open the Security Policy Rule dialog.
 - b. Specify which users and/or groups to match in the policy in one of the following ways:
 - If you want to specify specific users/groups as matching criteria, select the **User** tab and click the **Add** button in the Source User section to display a list of users and groups discovered by the firewall group mapping function. Select the users and/or groups to add to the policy.
 - If you want the policy to match any user who has or has not successfully authenticated and you don't need to know the specific user or group name, select **known-user** or **unknown** from the drop-down above the **Source User** list.
2. Configure the rest of the policy as appropriate and then click **OK** to save it. For details on other fields in the security policy, see [Set Up Basic Security Policies](#).

Enable User- and Group-Based Policy (Continued)

Step 3 Create your Captive Portal policies.

1. Select **Policies > Captive Portal**.
2. Click **Add** and enter a **Name** for the policy.
3. Define the matching criteria for the rule by completing the **Source**, **Destination**, and **Service/URL Category** tabs as appropriate to match the traffic you want to authenticate. The matching criteria on these tabs is the same as the criteria you define when creating a security policy. See [Set Up Basic Security Policies](#) for details.
4. Define the **Action** to take on traffic that matches the rule:
 - **no-captive-portal**—Allow traffic to pass without presenting a Captive Portal page for authentication.
 - **web-form**—Present a Captive Portal page for the user to explicitly enter authentication credentials or use client certificate authentication.
 - **browser-challenge**—Open an NTLM authentication request to the web browser of the user. The web browser will respond using the current login credentials of the user. If the login credentials are not available, the user will be prompted to supply them. If you select this action, you must enable [NT LAN Manager \(NTLM\)](#) authentication when you [Configure Captive Portal](#). If NTLM authentication fails, the firewall falls back to **web-form** authentication.
5. Click **OK**.

The following example shows a Captive Portal policy that instructs the firewall to present a web form to authenticate unknown users who send HTTP requests from the trust zone to the untrust zone.

Name	Tag	Source		Destination		Service	Action
		Zone	Address	Zone	Address		
all	none	trust	any	untrust	any	service-http	web-form

Step 4 Save your policy settings.

Click **Commit**.

Verify the User-ID Configuration

After you configure group mapping and user mapping and enable User-ID on your security policies and Captive Portal policies, you should verify that it is working properly.

Verify the User-ID Configuration																															
Step 1	Verify that group mapping is working.																														
	<p>From the CLI, enter the following command:</p> <pre>show user group-mapping statistics</pre>																														
Step 2	Verify that user mapping is working.																														
	<p>If you are using the on-device User-ID agent, you can verify this from the CLI using the following command:</p> <pre>show user ip-user-mapping-mp all</pre> <table border="1"> <thead> <tr> <th>IP (sec)</th> <th>Vsys</th> <th>From</th> <th>User</th> <th>Timeout</th> </tr> </thead> <tbody> <tr> <td>192.168.201.1</td> <td>vsys1</td> <td>UIA</td> <td>acme\george</td> <td>210</td> </tr> <tr> <td>192.168.201.11</td> <td>vsys1</td> <td>UIA</td> <td>acme\duane</td> <td>210</td> </tr> <tr> <td>192.168.201.50</td> <td>vsys1</td> <td>UIA</td> <td>acme\betsy</td> <td>210</td> </tr> <tr> <td>192.168.201.10</td> <td>vsys1</td> <td>UIA</td> <td>acme\administrator</td> <td>210</td> </tr> <tr> <td>192.168.201.100</td> <td>vsys1</td> <td>AD</td> <td>acme\administrator</td> <td>748</td> </tr> </tbody> </table> <p>Total: 5 users *: WMI probe succeeded</p>	IP (sec)	Vsys	From	User	Timeout	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
IP (sec)	Vsys	From	User	Timeout																											
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																											
Step 3	Test your security policy.																														
	<ul style="list-style-type: none"> From a machine in the zone where User-ID is enabled, attempt to access sites and applications to test the rules you defined in your policy and ensure that traffic is allowed and denied as expected. You can also use the test <code>security-policy-match</code> command to determine whether the policy is configured correctly. For example, suppose you have a rule that blocks user duane from playing World of Warcraft; you could test the policy as follows: <pre>test security-policy-match application worldofwarcraft source-user acme\duane source any destination any destination-port any protocol 6 "deny worldofwarcraft" { from corporate; source any; source-region any; to internet; destination any; destination-region any; user acme\duane; category any; application/service worldofwarcraft; action deny; terminal no; }</pre>																														

Verify the User-ID Configuration (Continued)

Step 4 Test your Captive Portal configuration.

1. From the same zone, go to a machine that is not a member of your directory, such as a Mac OS system, and try to ping to a system external to the zone. The ping should work without requiring authentication.
2. From the same machine, open a browser and navigate to a web site in a destination zone that matches a Captive Portal policy you defined. You should see the Captive Portal web form.

3. Log in using the correct credentials and confirm that you are redirected to the requested page.

4. You can also test your Captive Portal policy using the `test cp-policy-match` command as follows:

```
test cp-policy-match from corporate to internet
source 192.168.201.10 destination 8.8.8.8
```

```
Matched rule: 'captive portal' action: web-form
```

Step 5 Verify that user names are displayed in the log files (**Monitor > Logs**).

Receive Time	Category	URL	From Zone	To Zone	Source	Source User
12/18 15:16:17	computer-and-internet-info	*.urbanairship.com/	I3-trust	I3-untrust	10.31.32.18	acme\jdonaldson
12/18 15:16:01	social-networking	graph.facebook.com/	I3-trust	I3-untrust	10.31.32.18	acme\gbalfour
12/18 15:16:01	social-networking	graph.facebook.com/	I3-trust	I3-untrust	10.31.32.18	acme\jdonaldson
12/18 15:16:01	social-networking	orcart.facebook.com/	I3-trust	I3-untrust	10.31.32.18	acme\jparker
12/18 15:04:15	web-based-email	*.mg.mail.yahoo.com/	I3-trust	I3-untrust	10.31.32.18	acme\jreddick
12/18 14:51:13	search-engines	www.google.com/	I3-trust	I3-untrust	10.31.32.18	acme\jparker
12/18 14:49:06	search-engines	*.google.com/	I3-trust	I3-untrust	10.31.32.18	acme\jreddick
12/18 14:48:10	search-engines	www.google.com/	I3-trust	I3-untrust	10.31.32.18	acme\gbalfour
12/18 14:43:53	web-based-email	*.mg.mail.yahoo.com/	I3-trust	I3-untrust	10.31.32.18	acme\jreddick
12/18 14:43:51	internet-portals	android.register.push.mobile.yahoo.com/	I3-trust	I3-untrust	10.31.32.18	acme\jparker
12/18 14:43:51	web-based-email	*.mg.mail.yahoo.com/	I3-trust	I3-untrust	10.31.32.18	acme\jdonaldson
12/18 14:43:51	internet-portals	login.yahoo.com/	I3-trust	I3-untrust	10.31.32.18	acme\jparker
12/18 14:43:50	computer-and-internet-info	*.crittercism.com/	I3-trust	I3-untrust	10.31.32.18	acme\jreddick
12/18 14:36:35	computer-and-internet-info	mdmbeta.paloaltonetworks.com/	I3-trust	I3-untrust	10.31.32.18	acme\gbalfour
12/18 14:36:35	internet-portals	android.connector.push.mobile.yahoo.com/	I3-trust	I3-untrust	10.31.32.18	acme\jdonaldson
12/18 14:36:35	computer-and-internet-info	settings.crashlytics.com/	I3-trust	I3-untrust	10.31.32.18	acme\jparker

Verify the User-ID Configuration (Continued)

Step 6 Verify that user names are displayed in reports (**Monitor > Reports**). For example, when drilling down into the denied applications report, you should see a list of the users who attempted to access the applications as in the following example.

Application Information

Name: **zynga-games**

Description: This application can be used to identify and control the browser-based games created by Zynga that are available as widgets on social networking sites such as Facebook, Yahoo, and MySpace, as apps on iOS and Android mobile devices, and as independent web properties.

Standard Ports: **tcp/80,443**

Capable of File Transfer: **no**

Used by Malware: **no**

Excessive Bandwidth Use: **yes**

Evasive: **no**

Tunnels Other Applications: **no**

Depends on Applications: **facebook-apps, ssl, web-browsing**

Additional Information: [Wikipedia](#) [Zynga](#) [Google](#) [Yahoo!](#)

Category: **media**

Subcategory: **gaming**

Technology: **browser-based**

Risk: **2**

Widely Used: **yes**

Has Known Vulnerabilities: **yes**

Prone to Misuse: **no**

Session Timeout (seconds):

TCP Timeout (seconds):

UDP Timeout (seconds):

Top Applications

	Risk	Application	Sessions	Bytes
1	2	zynga-games	564	4.2 M

Top Sources

	Source Address	Source Host Name	Source User	Bytes	Sessions
1	10.154.63.185	10.154.63.185	pancademo\yamiro.js...	921.4 K	170
2	10.154.230.62	10.154.230.62	pancademo\david.joh...	488.4 K	132
3	10.154.62.201	10.154.62.201	pancademo\j.bolt	888.5 K	90
4	10.154.210.179	10.154.210.179	pancademo\snasi	302.0 K	76
5	10.154.6.160	10.154.6.160	pancademo\domingo.f...	1.2 M	61
6	10.154.110.194	10.154.110.194	pancademo\yan.ber...	351.1 K	24
7	10.154.7.121	10.154.7.121	pancademo\margaret....	30.4 K	9
8	10.154.12.28	10.154.12.28	pancademo\matthew....	90.4 K	2



App-ID

To safely enable applications on your network, the Palo Alto Networks next-generation firewalls provide both an application and web perspective—App-ID and URL Filtering—to protect against a full spectrum of legal, regulatory, productivity, and resource utilization risks.

App-ID enables visibility into the applications on the network, so you can learn how they work and understand their behavioral characteristics and their relative risk. This application knowledge allows you to create and enforce security policies to enable, inspect, and shape desired applications and block undesired applications. When you define policies to begin allowing traffic, App-ID begins to classify traffic without any additional configuration.

- ▲ [App-ID Overview](#)
- ▲ [Manage Custom or Unknown Applications](#)
- ▲ [Best Practices for Using App-ID in Policy](#)
- ▲ [Applications with Implicit Support](#)
- ▲ [Application Level Gateways](#)
- ▲ [Disable the SIP Application-level Gateway \(ALG\)](#)

App-ID Overview

App-ID, a patented traffic classification system only available in Palo Alto Networks firewalls, determines what the application is irrespective of port, protocol, encryption (SSH or SSL) or any other evasive tactic used by the application. It applies multiple classification mechanisms—application signatures, application protocol decoding, and heuristics—to your network traffic stream to accurately identify applications.

Here's how App-ID identifies applications traversing your network:

- Traffic is matched against policy to check whether it is allowed on the network.
- Signatures are then applied to allowed traffic to identify the application based on unique application properties and related transaction characteristics. The signature also determines if the application is being used on its default port or it is using a non-standard port. If the traffic is allowed by policy, the traffic is then scanned for threats and further analyzed for identifying the application more granularly.
- If App-ID determines that encryption (SSL or SSH) is in use, and a decryption policy is in place, the session is decrypted and application signatures are applied again on the decrypted flow.
- Decoders for known protocols are then used to apply additional context-based signatures to detect other applications that may be tunneling inside of the protocol (e.g., Yahoo! Instant Messenger used across HTTP). Decoders validate that the traffic conforms to the protocol specification and provide support for NAT traversal and opening dynamic pinholes for applications such as SIP and FTP.
- For applications that are particularly evasive and cannot be identified through advanced signature and protocol analysis, heuristics or behavioral analysis may be used to determine the identity of the application.

When the application is identified, the policy check determines how to treat the application, for example—block, or allow and scan for threats, inspect for unauthorized file transfer and data patterns, or shape using QoS.

Manage Custom or Unknown Applications

Palo Alto Networks provides weekly App-ID updates to identify new applications. By default, App-ID is always enabled on the firewall, and you don't need to enable a series of signatures to identify well-known applications. Typically, the only applications that are classified as *unknown* traffic—*tcp*, *udp* or *non-syn-tcp*—in the ACC and the traffic logs are commercially available applications that have not yet been added to App-ID, internal or custom applications on your network, or potential threats.

On occasion, the firewall may report an application as unknown for the following reasons:

- Incomplete data—A handshake took place, but no data packets were sent prior to the timeout.
- Insufficient data—A handshake took place followed by one or more data packets; however, not enough data packets were exchanged to identify the application.

The following choices are available to handle unknown applications:

- Create security policies to control unknown applications by unknown TCP, unknown UDP or by a combination of source zone, destination zone, and IP addresses.
- Request an App-ID from Palo Alto Networks—If you would like to inspect and control the applications that traverse your network, for any unknown traffic, you can record a packet capture. If the packet capture reveals that the application is a commercial application, you can submit this packet capture to Palo Alto Networks for App-ID development. If it is an internal application, you can create a custom App-ID and/or define an application override policy.
- Create a custom App-ID with a signature and attach it to a security policy, or create a custom App-ID and define an application override policy—A custom App-ID allows you to customize the definition of the internal application—its characteristics, category and sub-category, risk, port, timeout—and exercise granular policy control in order to minimize the range of unidentified traffic on your network. Creating a custom App-ID also allows you to correctly identify the application in the ACC and traffic logs and is useful in auditing/reporting on the applications on your network. For a custom application you can specify a signature and a pattern that uniquely identifies the application and attach it to a security policy that allows or denies the application.



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

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

The supported application protocol decoders that enable the firewall to detect applications that may be tunneling inside of the protocol include the following as of content update 424: HTTP, HTTPS, DNS, FTP, IMAP SMTP, Telnet, IRC (Internet Relay Chat), Oracle, RTMP, RTSP, SSH, GNU-Debugger, GIOP (Global Inter-ORB Protocol), Microsoft RPC, Microsoft SMB (also known as CIFS). Additionally, with the 4.0 release of PAN-OS, this custom App-ID capability has been expanded to include unknown TCP and unknown UDP.

Alternatively, if you would like the firewall to process the custom application using fast path (Layer-4 inspection instead of using App-ID for Layer-7 inspection), you can reference the custom App-ID in an application override policy. An application override with a custom application will prevent the session from being processed by the App-ID engine, which is a Layer-7 inspection. Instead it forces the firewall to handle the session as a regular stateful inspection firewall at Layer-4, and thereby saves application processing time.

For example, if you build a custom application that triggers on a host header *www.mywebsite.com*, the packets are first identified as *web-browsing* and then are matched as your custom application (whose parent application is web-browsing). Because the parent application is web-browsing, the custom application is inspected at Layer-7 and scanned for content and vulnerabilities.

If you define an application override, the firewall stops processing at Layer-4. The custom application name is assigned to the session to help identify it in the logs, and the traffic is not scanned for threats.

For more details, refer to the following articles:

- [Identifying Unknown Applications](#)
- [Video: How to Configure a Custom App-ID](#)
- [Custom Application Signatures](#)

Best Practices for Using App-ID in Policy

- ❑ Review the ACC for the list of applications on your network, and determine which applications to allow or block. If you are migrating from a firewall where you defined port-based rules, to get a list of applications that run on a given port, search by the port number in the application browser (**Objects > Applications**) on the Palo Alto Networks firewall or in [Applipedia](#).
- ❑ Use **application-default** for the **Service**. The firewall compares the port used with the list of default ports for that application. If the port used is not a default port for the application, the firewall drops the session and logs the message **appid policy lookup deny**. If you have an application that is accessed on many ports and you would like to limit the ports on which the application is used, specify it in **Service/Service Group** objects in policies.
- ❑ Use application filters to dynamically include new applications in existing policy rules. See an [example](#).

Applications with Implicit Support

When creating a policy to allow specific applications, you must also be sure that you are allowing any other applications on which the application depends. In many cases, you do not have to explicitly allow access to the dependent applications in order for the traffic to flow because the firewall is able to determine the dependencies and allow them implicitly. This implicit support also applies to [custom applications](#) that are based on HTTP, SSL, MS-RPC, or RTSP. Applications for which the firewall cannot determine dependent applications on time will require that you explicitly allow the dependent applications when defining your policies. You can determine application dependencies in [Applopedia](#).

The following table lists the applications for which the firewall has implicit support (as of [Content Update 484](#)).

Table: Applications with Implicit Support

Application	Implicitly Supports
360-safeguard-update	http
apple-update	http
apt-get	http
as2	http
avg-update	http
avira-antivir-update	http, ssl
blokus	rtmp
bugzilla	http
clubcooee	http
corba	http
cubby	http, ssl
dropbox	ssl
esignal	http
evernote	http, ssl
ezhelp	http
facebook	http, ssl
facebook-chat	jabber
facebook-social-plugin	http
fastviewer	http, ssl
forticlient-update	http
good-for-enterprise	http, ssl
google-cloud-print	http, ssl, jabber
google-desktop	http

Application	Implicitly Supports
google-drive-web	http
google-talk	jabber
google-update	http
gotomypc-desktop-sharing	citrix-jedi
gotomypc-file-transfer	citrix-jedi
gotomypc-printing	citrix-jedi
hipchat	http
iheartradio	ssl, http, rtmp
infront	http
instagram	http, ssl
issuu	http, ssl
java-update	http
jepptech-updates	http
kerberos	rpc
kik	http, ssl
lastpass	http, ssl
logmein	http, ssl
mcafee-update	http
megaupload	http
metatrader	http
mocha-rdp	t_120
mount	rpc
ms-frs	msrpc
ms-rdp	t_120
ms-scheduler	msrpc
ms-service-controller	msrpc
nfs	rpc
oovoo	http, ssl
paloalto-updates	ssl
panos-global-protect	http
panos-web-interface	http
pastebin	http

Application	Implicitly Supports
pastebin-posting	http
pinterest	http, ssl
portmapper	rpc
prezi	http, ssl
rdp2tcp	t_120
renren-im	jabber
roboform	http, ssl
salesforce	http
stumbleupon	http
supremo	http
symantec-av-update	http
trendmicro	http
trillian	http, ssl
twitter	http
whatsapp	http, ssl
xm-radio	rtsp

Application Level Gateways

The Palo Alto Networks firewall does not classify traffic by port and protocol; instead it identifies the application based on its unique properties and transaction characteristics using the App-ID technology. Some applications, however, require the firewall to dynamically open *pinholes* to establish the connection, determine the parameters for the session and negotiate the ports that will be used for the transfer of data; these applications use the application-layer payload to communicate the dynamic TCP or UDP ports on which the application opens data connections. For such applications, the firewall serves as an Application Level Gateway (ALG), and it opens a pinhole for a limited time and for exclusively transferring data or control traffic. The firewall also performs a NAT rewrite of the payload when necessary.

As of Content Release version 464, the Palo Alto Networks firewall provides NAT ALG support for the following protocols: FTP, H.225, H.248, MGCP, MySQL, Oracle/SQLNet/TNS, RPC, RTSP, SCCP, SIP, and UNIStim.



When the firewall serves as an ALG for the Session Initiation Protocol (SIP), by default it performs NAT on the payload and opens dynamic pinholes for media ports. In some cases, depending on the SIP applications in use in your environment, the SIP endpoints have NAT intelligence embedded in their clients. In such cases, you might need to disable the SIP ALG functionality to prevent the firewall from modifying the signaling sessions. When SIP ALG is disabled, if App-ID determines that a session is SIP, the payload is not translated and dynamic pinholes are not opened. See [Disable the SIP Application-level Gateway \(ALG\)](#).

Disable the SIP Application-level Gateway (ALG)

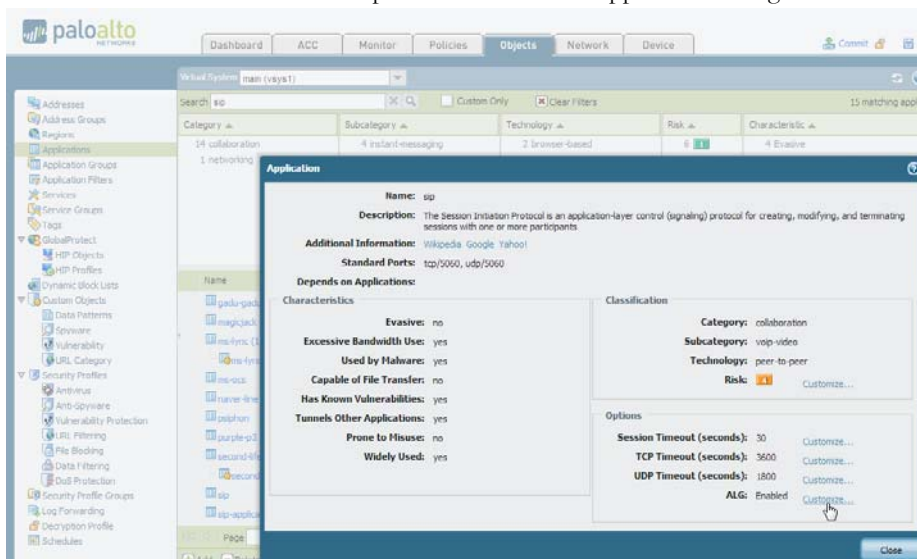
The Palo Alto Networks firewall uses the Session Initiation Protocol (SIP) application-level gateway (ALG) to open dynamic pinholes in the firewall where NAT is enabled. However, some applications—such as VoIP—have NAT intelligence embedded in the client application. In these cases, the SIP ALG on the firewall can interfere with the signaling sessions and cause the client application to stop working.

One solution to this problem is to define an Application Override Policy for SIP, but using this approach disables the App-ID and threat detection functionality. A better approach is to disable the SIP ALG, which does not disable App-ID or threat detection.

The following procedure describes how to disable the SIP ALG.

Disable the SIP ALG

1. Select **Objects > Applications**.
2. Select the **sip** application.
You can type **sip** in the **Search** box to help find the sip application.
3. Select **Customize...** for **ALG** in the Options section of the Application dialog box.



4. Select the **Disable ALG** check box in the Application - sip dialog box and click **OK**.



5. **Close** the Application dialog box and **Commit** the change.



Threat Prevention

The Palo Alto Networks next-generation firewall protects and defends your network from commodity threats and advanced persistent threats (APTs). The firewall's multi-pronged detection mechanisms include a signature-based (IPS/Command and Control/Antivirus) approach, heuristics-based (bot detection) approach, sandbox-based (WildFire) approach, and Layer 7 protocol analysis-based (App-ID) approach.

Commodity threats are exploits that are less sophisticated and more easily detected and prevented using a combination of the antivirus, anti-spyware, vulnerability protection and the URL filtering/Application identification capabilities on the firewall.

Advanced threats are perpetuated by organized cyber criminals or malicious groups that use sophisticated attack vectors to target your network, most commonly for intellectual property theft and financial data theft. These threats are more evasive and require intelligent monitoring mechanisms for detailed host and network forensics on malware. The Palo Alto Networks next-generation firewall in conjunction with [WildFire](#) and [Panorama](#) provides a comprehensive solution that intercepts and break the attack chain and provides visibility to prevent security infringement on your network—including mobile and virtualized—infrastructure.

- ▲ [Set Up Security Profiles and Policies](#)
- ▲ [Prevent Brute Force Attacks](#)
- ▲ [Customize the Action and Trigger Conditions for a Brute Force Signature](#)
- ▲ [Best Practices for Securing Your Network from Layer 4 and Layer 7 Evasions](#)
- ▲ [Enable Passive DNS Collection for Improved Threat Intelligence](#)
- ▲ [Use DNS Queries to Identify Infected Hosts on the Network](#)
- ▲ [Content Delivery Network Infrastructure for Dynamic Updates](#)
- ▲ [Threat Prevention Resources](#)

Set Up Security Profiles and Policies

The following sections provide basic threat prevention configuration examples:

- ▲ [Set Up Antivirus, Anti-spyware, and Vulnerability Protection](#)
- ▲ [Set Up Data Filtering](#)
- ▲ [Set Up File Blocking](#)

For information on controlling web access as part of your threat prevention strategy, see [URL Filtering](#).

Set Up Antivirus, Anti-spyware, and Vulnerability Protection

The following describes the steps needed to set up the default Antivirus, Anti-spyware, and Vulnerability Protection [Security Profiles](#).



All Anti-spyware and Vulnerability Protection signatures have a default action defined by Palo Alto Networks. You can view the default action by navigating to **Objects > Security Profiles > Anti-Spyware** or **Objects > Security Profiles > Vulnerability Protection** and then selecting a profile. Click the **Exceptions** tab and then click **Show all signatures** and you will see a list of the signatures with the default action in the Action column. To change the default action, you must create a new profile and then create rules with a non-default action, and/or add individual signature exceptions to **Exceptions** in the profile.

Set up Antivirus/Anti-spyware/Vulnerability Protection	
<p>Step 1 Verify that you have a Threat Prevention license.</p>	<ul style="list-style-type: none"> The Threat Prevention license bundles the Antivirus, Anti-spyware, and the Vulnerability Protection features in one license. Select Device > Licenses to verify that the Threat Prevention license is installed and check the expiration date.
<p>Step 2 Download the latest antivirus threat signatures.</p>	<ol style="list-style-type: none"> Select Device > Dynamic Updates and click Check Now at the bottom of the page to retrieve the latest signatures. <p>In the Actions column, click Download to install the latest Antivirus and Applications and Threats signatures.</p>
<p>Step 3 Schedule signature updates.</p>	<ol style="list-style-type: none"> From Device > Dynamic Updates, click the text to the right of Schedule to automatically retrieve signature updates for Antivirus and Applications and Threats. Specify the frequency and timing for the updates and whether the update will be downloaded and installed or only downloaded. If you select Download Only, you would need to manually go in and click the Install link in the Action column to install the signature. When you click OK, the update is scheduled. No commit is required. (Optional) You can also enter the number of hours in the Threshold field to indicate the minimum age of a signature before a download will occur. For example, if you entered 10, the signature must be at least 10 hours old before it will be downloaded, regardless of the schedule. In an HA configuration, you can also click the Sync To Peer option to synchronize the content update with the HA peer after download/install. This will not push the schedule settings to the peer device, you need to configure the schedule on each device.

Set up Antivirus/Anti-spyware/Vulnerability Protection (Continued)

Best Practices for Antivirus Schedules

The general recommendation for antivirus signature update schedules is to perform a **download-and-install** on a daily basis for antivirus and weekly for applications and vulnerabilities.

Recommendations for HA Configurations:

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

Step 4 Attach the security profiles to a security policy.

1. Select **Policies > Security**, select the desired policy to modify it and then click the **Actions** tab.
2. In **Profile Settings**, click the drop-down next to each security profile you would like to enable. In this example we choose default for **Antivirus, Vulnerability Protection, and Anti-spyware**.




If no security profiles have been previously defined, select **Profiles** from the **Profile Type** drop-down. You will then see the list of options to select the security profiles.

Step 5 Save the configuration.

Click **Commit**.

Set Up Data Filtering

The following describes the steps needed to configure a data filtering profile that will detect Social Security Numbers and a custom pattern identified in .doc and .docx documents.

Data Filtering Configuration Example	
<p>Step 1 Create a Data Filtering security profile.</p>	<ol style="list-style-type: none"> 1. Select Objects > Security Profiles > Data Filtering and click Add. 2. Enter a Name and a Description for the profile. In this example the name is <i>DF_Profile1</i> with the description <i>Detect Social Security Numbers</i>. 3. (Optional) If you want to collect data that is blocked by the filter, select the Data Capture check box.  You must set a password as described in Step 2 if you are using the data capture feature.
<p>Step 2 (Optional) Secure access to the data filtering logs to prevent other administrators from viewing sensitive data.</p> <p>When you enable this option, you will be prompted for the password when you view logs in Monitor > Logs > Data Filtering.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Content-ID. 2. Click Manage Data Protection in the Content-ID Features section. 3. Set the password that will be required to view the data filtering logs.

Data Filtering Configuration Example (Continued)

Step 3 Define the data pattern that will be used in the Data Filtering Profile.

In this example, we will use the keyword **confidential** and will set the option to search for SSN numbers with dashes (Example - 987-654-4320).



It is helpful to set the appropriate thresholds and define keywords within documents to reduce false positives.

1. From the Data Filtering Profile page click **Add** and select **New** from the **Data Pattern** drop-down. You can also configure data patterns from **Objects > Custom Signatures > Data Patterns**.
2. For this example, name the Data Pattern signature **Detect SS Numbers** and add the description **Data Pattern to detect Social Security numbers**.
3. In the **Weight** section for **SSN#** enter **3**. See [Weight and Threshold Values](#) for more details.

4. (Optional) You can also set **Custom Patterns** that will be subject to this profile. In this case, you specify a pattern in the custom patterns **Regex** field and set a weight. You can add multiple match expressions to the same data pattern profile. In this example, we will create a **Custom Pattern** named **SSN_Custom** with a custom pattern of **confidential** (the pattern is case sensitive) and use a weight of **20**. The reason we use the term **confidential** in this example is because we know that our social security Word docs contain this term, so we define that specifically.

Name	Regex	Weight
SSN_Custom	confidential	20

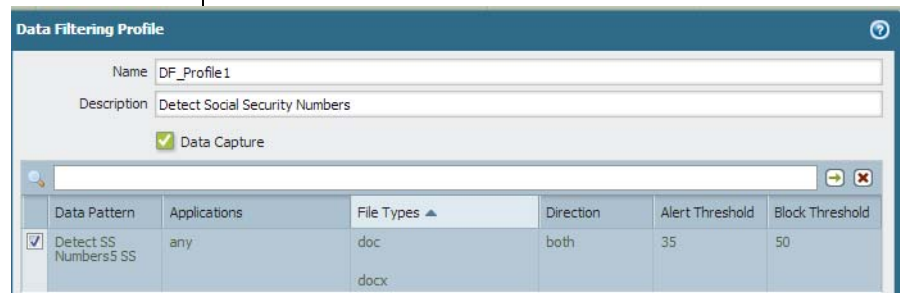
Data Filtering Configuration Example (Continued)

Step 4 Specify which applications to filter and the file types.

1. Set **Applications** to **Any**. This will detect any supported application such as: web-browsing, FTP, or SMTP. If you want to narrow down the application, you can select it from the list. For applications such as Microsoft Outlook Web App that uses SSL, you will need to enable decryption. Also make sure you understand the naming for each application. For example, Outlook Web App, which is the Microsoft name for this application is identified as the application **outlook-web** in the PAN-OS list of applications. You can check the logs for a given application to identify the name defined in PAN-OS.
2. Set **File Types** to **doc** and **docx** to only scan doc and docx files.

Step 5 Specify the direction of traffic to filter and the threshold values.

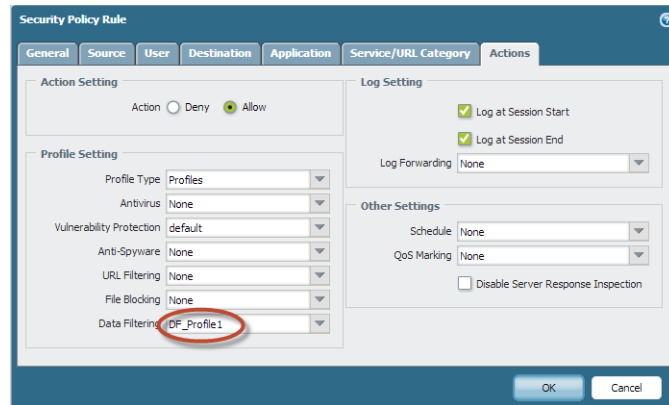
1. Set the **Direction** to **Both**. Files that are uploaded or downloaded will be scanned.
2. Set the **Alert Threshold** to **35**. In this case, an alert will be triggered if 5 instances of Social Security Numbers exist and 1 instance of the term **confidential** exists. The formula is 5 SSN instances with a weight of 3 = 15 plus 1 instance of the term confidential with a weight of 20 = 35.
3. Set the **Block Threshold** to **50**. The file will be blocked if the threshold of 50 instances of a SSN and/or the term confidential exists in the file. In this case, if the doc contained 1 instance of the word **confidential** with a weight of 20 that equals 20 toward the threshold, and the doc has 15 Social Security Numbers with a weight of 3 that equals 45. Add 20 and 45 and you have 65, which will exceed the block threshold of 50.



Data Filtering Configuration Example (Continued)

Step 6 Attach the Data Filtering profile to the security rule.

1. Select **Policies > Security** and select the security policy rule to which to apply the profile.
2. Click the security policy rule to modify it and then click the **Actions** tab. In the **Data Filtering** drop-down, select the new data filtering profile you created and then click **OK** to save. In this example, the data filtering rule name is **DF_Profile1**.



Step 7 **Commit** the configuration.

Step 8 Test the data filtering configuration.

If you have problems getting Data Filtering to work, you can check the Data Filtering log or the Traffic log to verify the application that you are testing with and make sure your test document has the appropriate number of unique Social Security Number instances. For example, an application such as Microsoft Outlook Web App may seem to be identified as web-browsing, but if you look at the logs, the application is **outlook-web**. Also increase the number of SSNs, or your custom pattern to make sure you are hitting the thresholds.

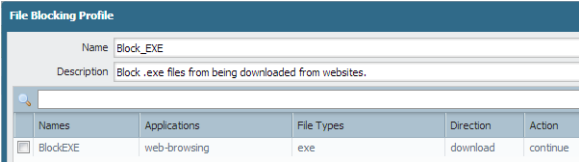
When testing, you must use real Social Security Numbers and each number must be unique. Also, when defining Custom Patterns as we did in this example with the word **confidential**, the pattern is case sensitive. To keep your test simple, you may want to just test using a data pattern first, then test the SSNs.

1. Access a client PC in the trust zone of the firewall and send an HTTP request to upload a .doc or .docx file that contains the exact information you defined for filtering.
2. Create a Microsoft Word document with one instance of the term confidential and five Social Security numbers with dashes.
3. Upload the file to a website. Use an HTTP site unless you have decryption configured, in which case you can use HTTPS.
4. Select **Monitoring > Logs > Data Filtering** logs.
5. Locate the log that corresponds to the file you just uploaded. To help filter the logs, use the source of your client PC and the destination of the web server. The action column in the log will show **reset-both**. You can now increase the number of Social Security Numbers in the document to test the block threshold.

Receive Time	File Name	Name	From Zone	To Zone	Source	Sou... User	Destination	To Port	Applicat...
10/25 14:03:06	DataFilter-Test.docx	Detect SS Numbers5 SS	13-vlan-trust	13-untrust	192.168.2.10		10.101.2.49	443	clearspac...
10/25 14:00:14	1351198789213	Detect SS Numbers5 SS	13-vlan-trust	13-untrust	192.168.2.10		10.101.2.49	443	clearspac...

Set Up File Blocking


This example will describe the basic steps needed to set up file blocking and forwarding. In this configuration, we will configure the options needed to prompt users to continue before downloading .exe files from websites. When testing this example, be aware that you may have other systems between you and the source that may be blocking content.

Configure File Blocking	
<p>Step 1 Create the file blocking profile.</p>	<ol style="list-style-type: none"> 1. Select Objects > Security Profiles > File Blocking and click Add. 2. Enter a Name for the file blocking profile, for example <i>Block_EXE</i>. Optionally enter a Description, such as <i>Block users from downloading exe files from websites</i>.
<p>Step 2 Configure the file blocking options.</p>	<ol style="list-style-type: none"> 1. Click Add to define the profile settings. 2. Enter a Name, such as <i>BlockEXE</i>. 3. Set the Applications for filtering, for example web-browsing. 4. Set File Types to exe. 5. Set the Direction to download. 6. Set the Action to continue. By choosing the continue option, users will be prompted with a response page prompting them to click continue before the file will be downloaded.  <ol style="list-style-type: none"> 7. Click OK to save the profile.
<p>Step 3 Apply the file blocking profile to a security policy.</p>	<ol style="list-style-type: none"> 1. Select Policies > Security and either select an existing policy or create a new policy as described in Set Up Basic Security Policies. 2. Click the Actions tab within the policy rule. 3. In the Profile Settings section, click the drop-down and select the file blocking profile you configured. In this case, the profile name is <i>Block_EXE</i>. 4. Commit the configuration. <p>If no security profiles have been previously defined, select the Profile Type drop-down and select Profiles. You will then see the list of options to select the security profiles.</p>

Configure File Blocking (Continued)

Step 4 To test your file blocking configuration, access a client PC in the trust zone of the firewall and attempt to download a .exe file from a website in the untrust zone. A response page should display. Click **Continue** to download the file. You can also set other actions, such as alert only, forward (which will forward to WildFire), or block, which will not provide a continue page to the user. The following shows the default response page for File Blocking:

Example: Default File Blocking Response Page



Step 5 (Optional) Define custom file blocking response pages (**Device > Response Pages**). This allows you to provide more information to users when they see a response page. You can include information such as company policy information and contact information for a Helpdesk.



When you create a file blocking profile with the action continue or continue-and-forward (used for WildFire forwarding), you can only choose the application web-browsing. If you choose any other application, traffic that matches the security policy will not flow through the firewall due to the fact that the users will not be prompted with a continue page. Also, if the website uses HTTPS, you will need to have a decryption policy in place.

You may want to check your logs to confirm what application is being used when testing this feature. For example, if you are using Microsoft Sharepoint to download files, even though you are using a web-browser to access the site, the application is actually `sharepoint-base`, or `sharepoint-document`. You may want to set the application type to **Any** for testing.

Prevent Brute Force Attacks

A brute force attack uses a large volume of requests/responses from the same source or destination IP address to break into a system. The attacker employs a trial-and-error method to guess the response to a challenge or a request.

The Vulnerability Protection profile on the firewall includes signatures to protect you from brute force attacks. Each signature has an ID, Threat Name, Severity and is triggered when a pattern is recorded. The pattern specifies the conditions and interval at which the traffic is identified as a brute-force attack; some signatures are associated with another child signature that is of a lower severity and specifies the pattern to match against. When a pattern matches against the signature or child signature, it triggers the default action for the signature.

To enforce protection:

- Attach the vulnerability profile to a security rule. See [Set Up Antivirus, Anti-spyware, and Vulnerability Protection](#).
- Install content updates that include new signatures to protect against emerging threats. See [Manage Content Updates](#).

Customize the Action and Trigger Conditions for a Brute Force Signature

The firewall includes two types of predefined brute force signatures—parent signature and child signature. A child signature is a single occurrence of a traffic pattern that matches the signature. A parent signature is associated with a child signature and is triggered when multiple events occur within a time interval and match the traffic pattern defined in the child signature.

Typically, a child signature is of default action *allow* because a single event is not indicative of an attack. In most cases, the action for a child signature is set to allow so that legitimate traffic is not blocked and threat logs are not generated for non-noteworthy events. Therefore, Palo Alto Networks recommends that you only change the default action after careful consideration.

In most cases, the brute force signature is a noteworthy event because of its recurrent pattern. If you would like to customize the action for a brute-force signature, you can do one of the following:

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

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



To effectively mitigate an attack, the **block-ip address** action is recommended over the drop or reset action for most brute force signatures.


Customize the Threshold and Action for a Signature

Step 1 Create a new Vulnerability Protection Profile.

1. Select **Objects > Security Profiles > Vulnerability Protection**.
2. Click **Add** and enter a **Name** for the Vulnerability Protection Profile.

Step 2 Create a rule that defines the action for all signatures in a category.

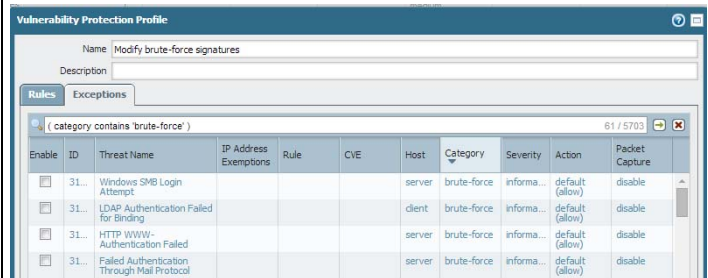
1. Select **Rules**, click **Add** and enter a **Name** for the rule.
2. Set the **Action**. In this example, it is set to **Block**.
3. Set **Category** to **brute-force**.
4. (Optional) If blocking, specify whether to block server or client, the default is any.
5. See [Step 3](#) to customize the action for a specific signature.
6. See [Step 4](#) to customize the trigger threshold for a parent signature.
7. Click **OK** to save the rule and the profile.



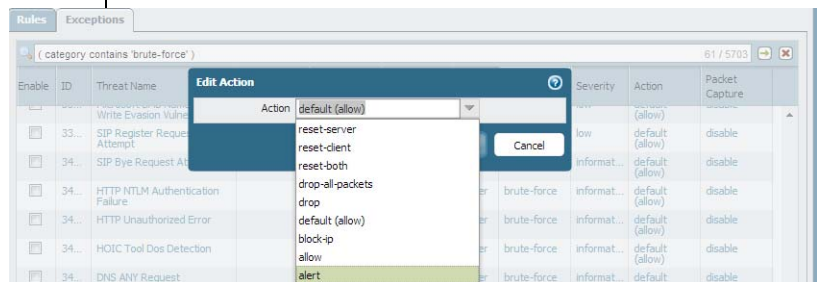
Customize the Threshold and Action for a Signature

Step 3 (Optional) Customize the action for a specific signature.

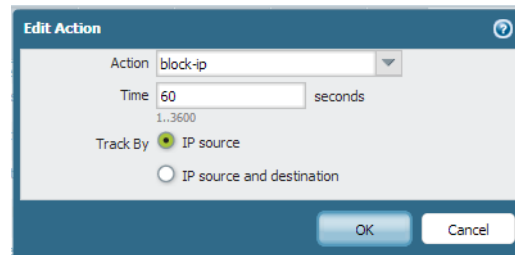
1. Select **Exceptions** and click **Show all signatures** to find the signature you want to modify.
To view all the signatures in the brute-force category, search for (category contains 'brute-force').
2. To edit a specific signature, click the predefined default action in the **Action** column.



3. Set the action to **allow**, **alert** or **block-ip**.




4. If you select **block-ip**, complete these additional tasks:
 - a. Specify the **Time** period (in seconds) after which to trigger the action.
 - b. In the **Track By** field, define whether to block the IP address by **IP source** or by **IP source and destination**.




5. Click **OK**
6. For each modified signature, select the check box in the **Enable** column.
7. Click **OK**.


Customize the Threshold and Action for a Signature

Step 4 Customize the trigger conditions for a parent signature.

A parent signature that can be edited is marked with this icon: .

In this example, the search criteria was brute force category and CVE-2008-1447.

1. Click  to edit the time attribute and the aggregation criteria for the signature.



- 2. To modify the trigger threshold specify the **Number of Hits per x seconds**.
- 3. Specify whether to aggregate the number of hits by **source**, **destination** or by **source and destination**.
- 4. Click **OK**.

Step 5 Attach this new profile to a security rule.

Step 6 Save your changes.

1. Click **Commit**.

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

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

- ❑ Upgrade to the most current PAN-OS software version and content release version to ensure that you have the latest security updates. See [Manage Content Updates](#) and [Install Software Updates](#).
- ❑ For servers, create security policy rules to only allow the application(s) that you sanction on each server. Verify that the standard port for the application matches the listening port on the server. For example, to ensure that only SMTP traffic is allowed to your email server set the Application to smtp and the Service to application-default. If your server uses only a subset of the standard ports (for example, if your SMTP server uses only port 587 while the smtp application has standard ports defined as 25 and 587), you should create a new custom service that only includes port 587 and use that new service in your security policy rule instead of using application-default. Additionally, make sure to restrict access to specific source and destinations zones and sets of IP addresses.
- ❑ Attach the following security profiles to your security policies to provide signature-based protection.
 - Create a vulnerability protection profile to block all vulnerabilities with severity low and higher.
 - Create an anti-spyware profile to block all spyware.
 - Create an antivirus profile to block all content that matches an antivirus signature.
- ❑ Block all unknown applications/traffic using security policy. Typically, the only applications that are classified as unknown traffic are internal or custom applications on your network, or potential threats. Because unknown traffic can be a non-compliant application or protocol that is anomalous or abnormal, or a known application that is using non-standard ports, unknown traffic should be blocked. See [Manage Custom or Unknown Applications](#).
- ❑ Create a file blocking profile that blocks Portable Executable (PE) file types for Internet-based SMB (Server Message Block) traffic from traversing the trust to untrust zones, (ms-ds-smb applications).
- ❑ Create a zone protection profile that is configured to protect against packet-based attacks:
 - Remove TCP timestamps on SYN packets before the firewall forwards the packet—When you remove the TCP timestamp option in a SYN packet, the TCP stack on both ends of the TCP connection will not support TCP timestamps. Therefore, by disabling the TCP timestamp for a SYN packet, you can prevent an attack that uses different timestamps on multiple packets for the same sequence number.
 - Drop malformed packets.
 - Drop mismatched and overlapping TCP segments—By deliberately constructing connections with overlapping but different data in them, attackers can attempt to cause misinterpretation of the intent of the connection. This can be used to deliberately induce false positives or false negatives. An attacker can use IP spoofing and sequence number prediction to intercept a user's connection and inject his/her own data into the connection. PAN-OS uses this field to discard such frames with mismatched and overlapping data. The scenarios where the received segment will be discarded are when the segment received is contained within another segment, the segment received overlaps with part of another segment, or the segment completely contains another segment.

- ❑ Verify that support for IPv6 is enabled, if you have configured IPv6 addresses on your network hosts. **(Network > Interfaces > Ethernet > IPv6)**

This allows access to IPv6 hosts and filters IPv6 packets that are encapsulated in IPv4 packets. Enabling support for IPv6 prevents IPv6 over IPv4 multicast addresses from being leveraged for network reconnaissance.

- ❑ Enable support for multicast traffic so that the firewall can enforce policy on multicast traffic. **(Network > Virtual Router > Multicast)**
- ❑ Enable the following CLI command to clear the URG bit flag in the TCP header and disallow out-of-band processing of packets.

The urgent pointer in the TCP header is used to promote a packet for immediate processing by removing it from the processing queue and expediting it through the TCP/IP stack on the host. This process is called out-of-band processing. Because the implementation of the urgent pointer varies by host, to eliminate ambiguity, use the following CLI command to disallow out-of-band processing; the out-of-band byte in the payload becomes part of the payload and the packet is not processed urgently. Making this change allows you to remove ambiguity in how the packet is processed on the firewall and the host, and the firewall sees the exact same stream in the protocol stack as the host for whom the packet is destined.

```
set deviceconfig setting tcp urgent-data clear
```

- ❑ If you configure the firewall to clear the URG bit flag and the packet has no other flags set in the TCP header, use the following CLI command to configure the firewall to drop packets with no flags:

```
set deviceconfig setting tcp drop-zero-flag yes
```

- ❑ Enable the following CLI command for disabling the bypass-exceed-queue.

The bypass exceed queue is required for out of order packets. This scenario is most common in an asymmetric environment where the firewall receives packets out of order. For identification of certain applications (App-ID) the firewall performs heuristic analysis. If the packets are received out of order, the data must be copied to a queue in order to complete the analysis for the application.

```
set deviceconfig setting application bypass-exceed-queue no
```

- ❑ Enable the following CLI commands for disabling the inspection of packets when the out-of-order packet limit is reached. The Palo Alto Networks firewall can collect up to 64 out-of-order packets per session. This counter identifies that packets have exceeded the 64-packet limit. When the bypass setting is set to **no**, the device drops the out-of-order packets that exceed the 64-packet limit. A commit is required.

```
set deviceconfig setting tcp bypass-exceed-oo-queue no
set deviceconfig setting ctd tcp-bypass-exceed-queue no
set deviceconfig setting ctd udp-bypass-exceed-queue no
```

- ❑ Enable the following CLI commands for checking the TCP timestamp. The TCP timestamp records when the segment was sent and allows the firewall to verify that the timestamp is valid for that session. Packets with invalid timestamps are dropped with this setting is enabled.

```
set deviceconfig setting tcp check-timestamp-option yes
```

- ❑ Disable the HTTP Range option. The HTTP Range option allows a client to fetch part of a file only. When a next-generation firewall in the path of a transfer identifies and drops a malicious file, it terminates the TCP session with a RST packet. If the web browser implements the HTTP Range option, it can start a new session to fetch only the remaining part of the file. This prevents the firewall from triggering the same signature again due to the lack of context into the initial session, while at the same time allowing the web browser to reassemble the file and deliver the malicious content. To prevent this, disable the HTTP Range option as follows:

```
set deviceconfig setting ctd skip-block-http-range no
```

Enable Passive DNS Collection for Improved Threat Intelligence

Passive DNS is an opt-in feature that enables the firewall to act as a passive DNS sensor and send select DNS information to Palo Alto Networks for analysis in order to improve threat intelligence and threat prevention capabilities. The data collected includes non-recursive (i.e. originating from the local recursive resolver, not individual clients) DNS query and response packet payloads. Data submitted via the Passive DNS Monitoring feature consists solely of mappings of domain names to IP addresses. Palo Alto Networks retains no record of the source of this data and does not have the ability to associate it with the submitter at a future date.

The Palo Alto Networks threat research team uses this information to gain insight into malware propagation and evasion techniques that abuse the DNS system. Information gathered through this data collection is used to improve accuracy and malware detection abilities within PAN-DB URL filtering, DNS-based command-and-control signatures, and WildFire.

DNS responses are only forwarded to the Palo Alto Networks and will only occur when the following requirements are met:

- DNS response bit is set
- DNS truncated bit is not set
- DNS recursive bit is not set
- DNS response code is 0 or 3 (NX)
- DNS question count bigger than 0
- DNS Answer RR count is bigger than 0 or if it is 0, the flags need to be 3 (NX)
- DNS query record type are A, NS, CNAME, AAAA, MX

Passive DNS monitoring is disabled by default, but it is recommended that you enable it to facilitate enhanced threat intelligence. Use the following procedure to enable Passive DNS:

Enable Passive DNS

1. Select **Objects > Security Profiles > Anti-Spyware**.

2. Select an existing profile to modify it or configure a new profile.



The Anti-Spyware profile must be attached to a security policy that governs your DNS server's external DNS traffic.

3. Select the **DNS Signatures** tab and click the **Enable Passive DNS Monitoring** check box.

4. Click **OK** and then **Commit**.

Use DNS Queries to Identify Infected Hosts on the Network

The DNS sinkhole action in Anti-Spyware profiles enables the firewall to forge a response to a DNS query for a known malicious domain, causing the malicious domain name to resolve to an IP address that you define. This allows you to identify hosts on your network that have been infected with malware. The following topics describe the DNS sinkhole action and provide instructions for enabling it and monitoring logs to identify infected hosts.

- ▲ [DNS Sinkholing](#)
- ▲ [Configure DNS Sinkholing](#)
- ▲ [Identify Infected Hosts](#)

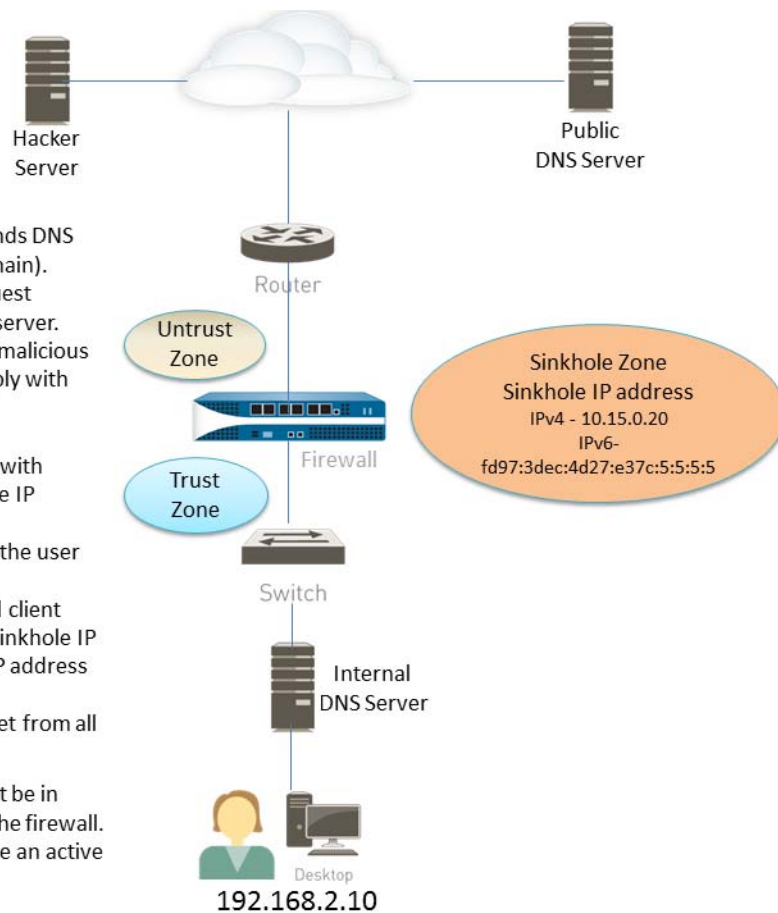
DNS Sinkholing

DNS sinkholing helps you to identify infected hosts on the protected network using DNS traffic in situations where the firewall cannot see the infected client's DNS query (that is, the firewall cannot see the originator of the DNS query). In a typical deployment where the firewall is north of the local DNS server, the threat log will identify the local DNS resolver as the source of the traffic rather than the actual infected host. Sinkholing malware DNS queries solves this visibility problem by forging responses to the client host queries directed at malicious domains, so that clients attempting to connect to malicious domains (for command-and-control, for example) will instead attempt to connect to a sinkhole IP address you define as illustrated in [Configure DNS Sinkholing](#). Infected hosts can then be easily identified in the traffic logs because any host that attempts to connect to the sinkhole IP address are most likely infected with malware.

Figure: DNS Sinkholing Example



1. Botnet on client host 192.168.2.10 sends DNS query for Hacker Server (malicious domain).
2. The internal DNS server relays the request through the firewall to the public DNS server.
3. The firewall DNS signature detects the malicious domain request and forges the DNS reply with the sinkhole IPv4 10.15.0.20 and IPv6 fd97:3dec:4d27:e37c:5:5:5:5.
4. Botnet then attempts to communicate with Hacker Server, but sends to the sinkhole IP address instead.
5. Session goes through the firewall from the user to the sinkhole address.
6. The security admin can then identify all client hosts trying to communicate with the sinkhole IP address by searching for the sinkhole IP address in the threat and traffic logs.
7. The Helpdesk then eradicates the botnet from all infected hosts.


Note: The client hosts and sinkhole IP must be in different zones, so sessions pass through the firewall. The sinkhole IP address does not have to be an active host, just an unused IP address.




Configure DNS Sinkholing

To enable DNS Sinkholing, you must enable the action in an Anti-spyware profile and attach the profile to a security rule. When a client host attempts to access a malicious domain, the firewall forges the destination IP address in the packet using the IP address you configure as the DNS sinkhole address.

Configure DNS Sinkholing	
<p>Step 1 Obtain both an IPv4 and IPv6 address to use as the sinkhole IP addresses.</p> <p>The DNS sinkhole address must be in a different zone than the client hosts to ensure that when an infected host attempts to start a session with the sinkhole IP address, it will be routed through the firewall. The reason both IPv4 and IPv6 are needed is because malicious software may perform DNS queries using one or both of these protocols.</p> <p> This sinkhole addresses must be reserved for this purpose and do not need to be assigned to a physical host. You can optionally use a honey-pot server as a physical host to further analyze the malicious traffic.</p>	<p>This configuration example uses the following DNS sinkhole addresses:</p> <p>IPv4 DNS sinkhole address—10.15.0.20</p> <p>IPv6 DNS sinkhole address—fd97:3dec:4d27:e37c:5:5:5:5</p>
<p>Step 2 Configure the sinkhole interface and zone.</p> <p>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.</p> <p> Use a dedicated zone for sinkhole traffic, because the infected host will be sending traffic to this zone.</p>	<ol style="list-style-type: none"> 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, use the IPv4 address 10.15.0.20 as the 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, use the IPv6 address fd97:3dec:4d27:e37c::/64 as the 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.

Configure DNS Sinkholing (Continued)	
<p>Step 3 Enable DNS sinkholing on the anti-spyware profile.</p>	<ol style="list-style-type: none"> 1. Select Objects > Security Profiles > Anti-Spyware. 2. Modify an existing profile, or select one of the existing defaults and clone it. 3. Name the profile and then select the DNS Signatures tab. 4. In the Action on DNS queries drop-down, select sinkhole. 5. In the Sinkhole IPv4 field enter the sinkhole IPv4 sinkhole address you configured in Step 2 (10.15.0.20 in this example). 6. In the Sinkhole IPv6 field enter the sinkhole IPv6 sinkhole address you configured in Step 2 (fd97:3dec:4d27:e37c:5:5:5:5 in this example). <ul style="list-style-type: none">  The default sinkhole address is the loopback address (127.0.0.1 for IPv4 and ::1 for IPv6). 7. (Optional) In the Packet Capture drop-down, select single-packet or extended-capture. The single-packet option will capture the first packet of the session or you can select extended to capture between 1-50 packets. You can then use the packet captures for further analysis. 8. Click OK to save the profile.
<p>Step 4 Edit the security policy rule that allows traffic from client hosts in the trust zone to the untrust zone to include the sinkhole zone as a destination and attach the anti-spyware profile.</p> <p>To ensure that you are identifying traffic from infected hosts, make these changes to the security rule(s) that allow traffic from client hosts in the trust zone to the untrust zone. By adding the sinkhole zone as a destination on the rule, you enable infected clients to send bogus DNS queries to the DNS sinkhole.</p>	<ol style="list-style-type: none"> 1. Select Policies > Security. 2. Select an existing rule that allows traffic from the client host zone to the untrust zone. 3. On the Destination tab, Add the Sinkhole zone. This allows client host traffic to flow to the sinkhole zone. 4. On the Actions tab, select the Log at Session Start check box to enable logging. This will ensure that traffic from client hosts in the Trust zone will be logged when accessing the Untrust or Sinkhole zones. 5. In the Profile Setting section, select the Anti-Spyware profile in which you enabled DNS sinkholing. 6. Click OK to save the security rule and then Commit.

Configure DNS Sinkholing (Continued)

<p>Step 5 To ensure that you will be able to identify infected hosts, verify that traffic going from the client host in the Trust zone to the new Sinkhole zone is being logged.</p> <p>In this example, the infected client host is 192.168.2.10 and the Sinkhole IPv4 address is 10.15.0.20.</p>	<ol style="list-style-type: none"> From a client host in the trust zone, open a command prompt and run the following command: <pre>C:\>ping <sinkhole address></pre> The following example output shows the ping request to the DNS sinkhole address at 10.15.0.2 and the result, which is <code>Request timed out</code> because in this example the sinkhole IP address is not assigned to a physical host: <pre>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)</pre> 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. <p> 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.</p>
<p>Step 6 Identify a malicious domain that you can use to verify that the DNS sinkhole functionality is configured properly.</p> <p>You must test this feature using a malicious domain that is included in the firewall's current antivirus signature database. The DNS Signatures used to identify malicious domains is only part of the full antivirus signature database, which contains hundreds of thousands of signatures.</p>	<p>To find a malicious domain for testing:</p> <ol style="list-style-type: none"> Select Device > Dynamic Updates and in the Antivirus section click the Release Notes link for the current antivirus DB that is installed. You can also find the antivirus release notes on the support site in Dynamic Updates. In most cases, the signature update is an incremental update, so only new viruses and DNS signatures are listed. There are many antivirus signatures and DNS signatures that will already be installed on the firewall. In the second column of the release note, locate a line item with a domain extension (for example, com, edu, or net). The left column will show the domain name. For example, in Antivirus release 1117-1560, there is an item in the left column named "tbsbana" and the right column lists "net". <p>The following shows the content in the release note for this line item:</p> <pre>conficker:tbsbanal variants: net</pre> <p>Because this domain shows up in the current database, it will work for testing.</p>

Configure DNS Sinkholing (Continued)**Step 7** Test the sinkhole action

This is similar to the action that would be performed if the client host was infected and the malicious application was attempting to reach a hacker server using DNS queries.

1. From the client host, open a command prompt.
2. Perform an NSLOOKUP to a URL that you identified as a known malicious domain in [Step 6](#).
For example, using the URL `track.bidtrk.com`:

```
C:\>nslookup track.bidtrk.com
Server: my-local-dns.local
Address: 10.0.0.222
Non-authoritative answer:
Name: track.bidtrk.com.org
Addresses: fd97:3dec:4d27:e37c:5:5:5:5
10.15.0.20
```

In the output, note that the NSLOOKUP to the malicious domain has been forged using the sinkhole IP addresses that we configured (10.15.0.20). Because the domain matched a malicious DNS signature, the sinkhole action was performed.
3. Select **Monitor > Logs > Threat** and locate the corresponding threat log entry to verify that the correct action was taken on the NSLOOKUP request.
4. Perform a ping to `track.bidtrk.com`, which will generate network traffic to the sinkhole address.

Identify Infected Hosts

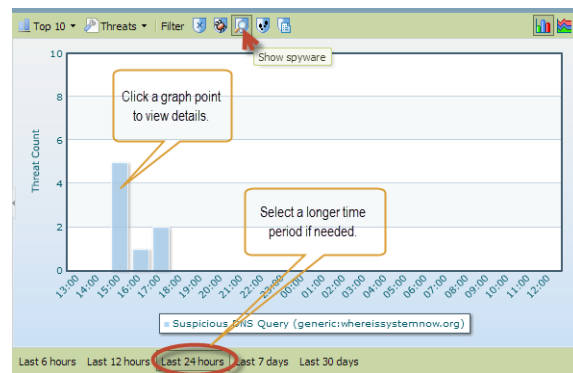
After you have configured DNS sinkholing and verified that traffic to a malicious domain goes to the sinkhole address, you should regularly monitor traffic to the sinkhole address, so that you can track down the infected hosts and eliminate the threat.

DNS Sinkhole Verification and Reporting

Step 1 Use App Scope to identify infected client hosts.

1. Select **Monitor > App Scope** and select **Threat Monitor**.
2. Click the **Show spyware** button along the top of the display page.
3. Select a time range.

The following screenshot shows three instances of Suspicious DNS queries, which were generated when the test client host performed an NSLOOKUP on a known malicious domain. Click the graph to see more details about the event.



DNS Sinkhole Verification and Reporting (Continued)

Step 2 Configure a custom report to identify all client hosts that have sent traffic to the sinkhole IP address, which is 10.15.0.20 in this example.

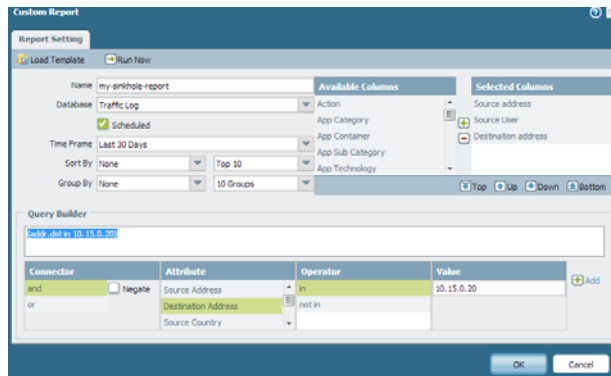


Forward to an SNMP manager, Syslog server and/or Panorama to enable alerts on these events.

In this example, the infected client host performed an NSLOOKUP to a known malicious domain that is listed in the Palo Alto Networks DNS Signature database. When this occurred, the query was sent to the local DNS server, which then forwarded the request through the firewall to an external DNS server. The firewall security policy with the Anti-Spyware profile configured matched the query to the DNS Signature database, which then forged the reply using the sinkhole address of 10.15.0.20 and fd97:3dec:4d27:e37c:5:5:5:5. The client attempts to start a session and the traffic log records the activity with the source host and the destination address, which is now directed to the forged sinkhole address.

Viewing the traffic log on the firewall allows you to identify any client host that is sending traffic to the sinkhole address. In this example, the logs show that the source address 192.168.2.10 sent the malicious DNS query. The host can then be found and cleaned. Without the DNS sinkhole option, the administrator would only see the local DNS server as the system that performed the query and would not see the client host that is infected. If you attempted to run a report on the threat log using the action “Sinkhole”, the log would show the local DNS server, not the infected host.

1. Select **Monitor > Manage Custom Reports**.
2. Click **Add** and **Name** the report.
3. Define a custom report that captures traffic to the sinkhole address as follows:
 - **Database**—Select **Traffic Log**.
 - **Scheduled**—Enable **Scheduled** and the report will run every night.
 - **Time Frame**—30 days
 - **Selected Columns**—Select **Source address** or **Source User** (if you have User-ID configured), which will identify the infected client host in the report, and **Destination address**, which will be the sinkhole address.
 - In the section at the bottom of the screen, create a custom query for traffic to the sinkhole address (10.15.0.20 in this example). You can either enter the destination address in the **Query Builder** window (**addr.dst in 10.15.0.20**) or select the following in each column and click **Add**: Connector = and, Attribute = Destination Address, Operator = in, and Value = 10.15.0.20. Click **Add** to add the query.



4. Click **Run Now** to run the report. The report will show all client hosts that have sent traffic to the sinkhole address, which indicates that they are most likely infected. You can now track down the hosts and check them for spyware.

Source	Source Host Name	Source User	Destination	Destination Host Name
1	192.168.2.10	192.168.2.10	10.15.0.20	10.15.0.20

5. To view scheduled reports that have run, select **Monitor > Reports**.

Content Delivery Network Infrastructure for Dynamic Updates

Palo Alto Networks maintains a Content Delivery Network (CDN) infrastructure for delivering content updates to the Palo Alto Networks devices. The devices access the web resources in the CDN to perform various App-ID and Content-ID functions. For enabling and scheduling the content updates, see [Manage Content Updates](#).

The following table lists the web resources that the firewall accesses for a feature or application:

Resource	URL	Static Addresses (If a static server is required)
Application Database	<ul style="list-style-type: none"> updates.paloaltonetworks.com:443 	staticupdates.paloaltonetworks.com or the IP address 199.167.52.15
Threat/Antivirus Database	<ul style="list-style-type: none"> updates.paloaltonetworks.com:443 downloads.paloaltonetworks.com:443 <p>As a best practice, set the update server to updates.paloaltonetworks.com. This allows the Palo Alto Networks device to receive content updates from the server closest to it in the CDN infrastructure.</p>	staticupdates.paloaltonetworks.com or the IP address 199.167.52.15
PAN-DB URL Filtering	<p>*.urlcloud.paloaltonetworks.com</p> <p>Resolves to the primary URL s0000.urlcloud.paloaltonetworks.com and is then redirected to the regional server that is closest:</p> <ul style="list-style-type: none"> 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	<ul style="list-style-type: none"> database.brightcloud.com:443/80 service.brightcloud.com:80 	Contact BrightCloud Customer Support.

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:

- [Applopedia](#)—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 provide the capability to decrypt and inspect traffic for visibility, control, and granular security. Decryption on a Palo Alto Networks firewall includes the capability to enforce security policies on encrypted traffic, where otherwise the encrypted traffic might not be blocked and shaped according to your configured security settings. Use decryption on a firewall to prevent malicious content from entering your network or sensitive content from leaving your network concealed as encrypted traffic. Enabling decryption on a Palo Alto Networks firewall can include preparing the keys and certificates required for decryption, creating a decryption policy, and configuring decryption port mirroring. See the following topics to learn about and configure decryption:

- ▲ [Decryption Overview](#)
- ▲ [Decryption Concepts](#)
- ▲ [Configure SSL Forward Proxy](#)
- ▲ [Configure SSL Inbound Inspection](#)
- ▲ [Configure SSH Proxy](#)
- ▲ [Configure Decryption Exceptions](#)
- ▲ [Enable Users to Opt Out of SSL Decryption](#)
- ▲ [Configure Decryption Port Mirroring](#)
- ▲ [Temporarily Disable SSL Decryption](#)

Decryption Overview

Secure Sockets Layer (SSL) and Secure Shell (SSH) are encryption protocols used to secure traffic between two entities, such as a web server and a client. SSL and SSH encapsulate traffic, encrypting data so that it is meaningless to entities other than the client and server with the keys to decode the data and the certificates to affirm trust between the devices. Traffic that has been encrypted using the protocols SSL and SSH can be decrypted to ensure that these protocols are being used for the intended purposes only, and not to conceal unwanted activity or malicious content.

Palo Alto Networks firewalls decrypt encrypted traffic by using keys to transform strings (passwords and shared secrets) from ciphertext to plaintext (decryption) and from plaintext back to ciphertext (re-encrypting traffic as it exits the device). Certificates are used to establish the firewall as a trusted third party and to create a secure connection. SSL decryption (both forward proxy and inbound inspection) requires certificates to establish trust between two entities in order to secure an SSL/TLS connection. Certificates can also be used when excluding servers from SSL decryption. You can integrate a hardware security module (HSM) with a firewall to enable enhanced security for the private keys used in SSL forward proxy and SSL inbound inspection decryption. To learn more about storing and generating keys using an HSM and integrating an HSM with your firewall, see [Secure Keys with a Hardware Security Module](#). SSH decryption does not require certificates.

Palo Alto Networks firewall decryption is policy-based, and can be used to decrypt, inspect, and control both inbound and outbound SSL and SSH connections. Decryption policies allow you to specify traffic for decryption according to destination, source, or URL category and in order to block or restrict the specified traffic according to your security settings. The firewall uses certificates and keys to decrypt the traffic specified by the policy to plaintext, and then enforces App-ID and security settings on the plaintext traffic, including Decryption, Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles. After traffic is decrypted and inspected on the firewall, the plaintext traffic is re-encrypted as it exits the firewall to ensure privacy and security. Use policy-based decryption on the firewall to achieve outcomes such as the following:

- Prevent malware concealed as encrypted traffic from being introduced into an corporate network.
- Prevent sensitive corporate information from moving outside the corporate network.
- Ensure the appropriate applications are running on a secure network.
- Selectively decrypt traffic; for example, exclude traffic for financial or healthcare sites from decryption by configuring a decryption exception.

The three decryption policies offered on the firewall, [SSL Forward Proxy](#), [SSL Inbound Inspection](#), and [SSH Proxy](#), all provide methods to specifically target and inspect SSL outbound traffic, SSL inbound traffic, and SSH traffic, respectively. The decryption policies provide the settings for you to specify what traffic to decrypt and decryption profiles can be selected when creating a policy, in order to apply more granular security settings to decrypted traffic, such as checks for server certificates, unsupported modes, and failures. This policy-based decryption on the firewall gives you visibility into and control of SSL and SSH encrypted traffic according to configurable parameters.

You can also choose to extend a decryption configuration on the firewall to include [Decryption Port Mirroring](#), which allows for decrypted traffic to be forwarded as plaintext to a third party solution for additional analysis and archiving.

Decryption Concepts

To learn about keys and certificates for decryption, decryption policies, and decryption port mirroring, see the following topics:

- ▲ [Keys and Certificates for Decryption Policies](#)
- ▲ [SSL Forward Proxy](#)
- ▲ [SSL Inbound Inspection](#)
- ▲ [SSH Proxy](#)
- ▲ [Decryption Exceptions](#)
- ▲ [Decryption Port Mirroring](#)

Keys and Certificates for Decryption Policies

Keys are strings of numbers that are typically generated using a mathematical operation involving random numbers and large primes. Keys are used to transform other strings—such as passwords and shared secrets—from plaintext to ciphertext (called *encryption*) and from ciphertext to plaintext (called *decryption*). Keys can be symmetric (the same key is used to encrypt and decrypt) or asymmetric (one key is used for encryption and a mathematically related key is used for decryption). Any system can generate a key.

X.509 certificates are used to establish trust between a client and a server in order to establish an SSL connection. A client attempting to authenticate a server (or a server authenticating a client) knows the structure of the X.509 certificate and therefore knows how to extract identifying information about the server from fields within the certificate, such as its FQDN or IP address (called a *common name* or *CN* within the certificate) or the name of the organization, department, or user to which the certificate was issued. All certificates must be issued by a certificate authority (CA). After the CA verifies a client or server, the CA issues the certificate and signs it using its private key.

With a decryption policy configured, an SSL/TLS session between the client and the server is established only if the firewall trusts the CA that signed the server's certificate. In order to establish trust, the firewall must have the server's root CA certificate in its certificate trust list (CTL) and use the public key contained in that root CA certificate to verify the signature. The firewall then presents a copy of the server certificate signed by the Forward Trust certificate for the client to authenticate. You can also configure the firewall to use an enterprise CA as a forward trust certificate for SSL Forward Proxy. If the firewall does not have the server's root CA certificate in its CTL, the firewall will present a copy of the server certificate signed by the Forward Untrust certificate to the client. The Forward Untrust certificate ensures that clients are prompted with a certificate warning when attempting to access sites hosted by a server with untrusted certificates.

For detailed information on certificates, see [Certificate Management](#).



To control the trusted CAs that your device trusts, use the **Device > Certificate Management > Certificates > Default Trusted Certificate Authorities** tab on the firewall web interface.

Table: Palo Alto Networks Device Keys and Certificates describes the different keys and certificates used by Palo Alto Networks devices for decryption. As a best practice, use different keys and certificates for each usage.

Table: Palo Alto Networks Device Keys and Certificates

Key/Certificate Usage	Description
Forward Trust	The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate that is signed by a CA that the firewall trusts. To configure a Forward Trust certificate on the firewall, see Step 2 in the Configure SSL Forward Proxy task. By default, the firewall determines the key size to use for the client certificate based on the key size of the destination server. However, you can also set a specific key size for the firewall to use. See Configure the Key Size for SSL Forward Proxy Server Certificates . For added security, store the forward trust certificate on a Hardware Security Module (HSM), see Store Private Keys on an HSM .
Forward Untrust	The certificate the firewall presents to clients during decryption if the site the client is attempting to connect to has a certificate that is signed by a CA that the firewall does not trust. To configure a Forward Untrust certificate on the firewall, see Step 3 in the Configure SSL Forward Proxy task.

Key/Certificate Usage	Description
SSL Exclude Certificate	Certificates for servers that you want to exclude from SSL decryption. For example, if you have SSL decryption enabled, but have certain servers that you do not want included in SSL decryption, such as the web services for your HR systems, you would import the corresponding certificates onto the firewall and configure them as SSL Exclude Certificates. See Exclude a Server From Decryption .
SSL Inbound Inspection	The certificate used to decrypt inbound SSL traffic for inspection and policy enforcement. For this application, you would import the server certificate for the servers for which you are performing SSL inbound inspection, or store them on an HSM (see Store Private Keys on an HSM).

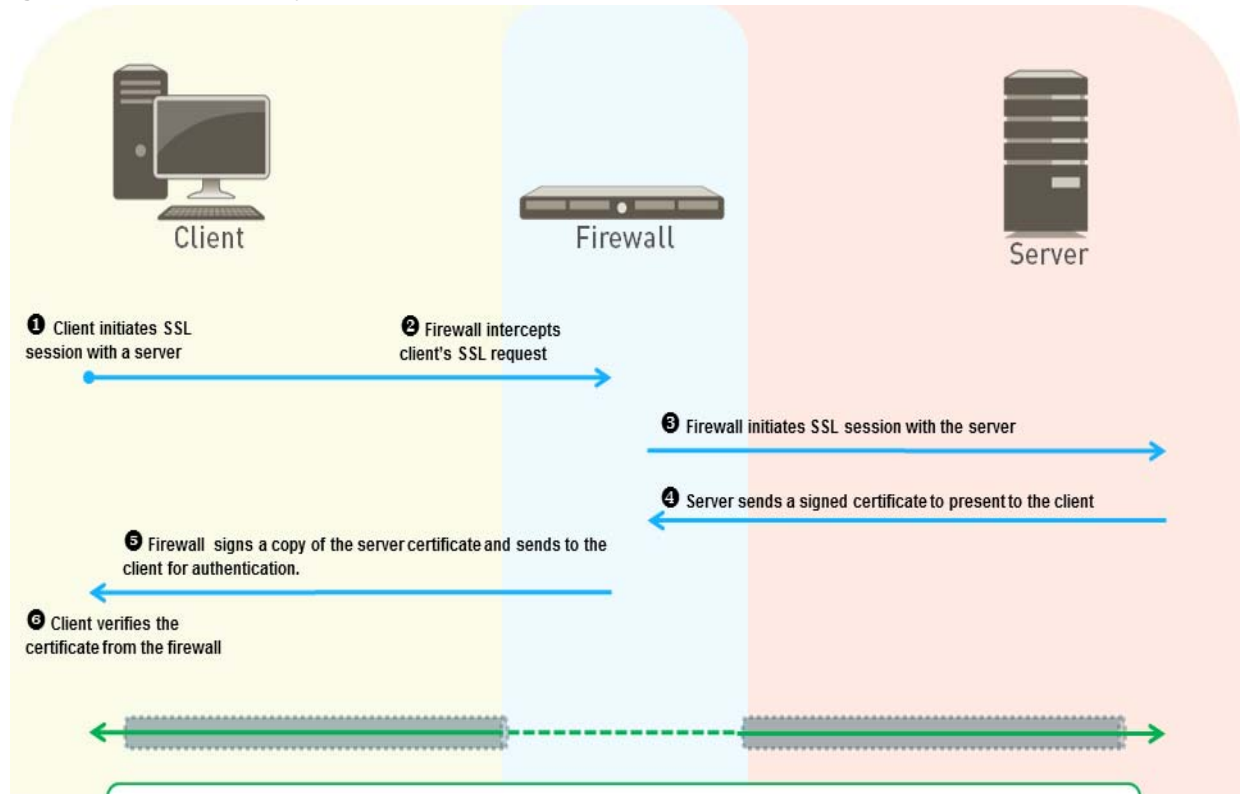
SSL Forward Proxy

Use an SSL Forward Proxy decryption policy to decrypt and inspect SSL/TLS traffic from internal users to the web. SSL Forward Proxy decryption prevents malware concealed as SSL encrypted traffic from being introduced to your corporate network; for example, if an employee is using her Gmail account from her corporate office and opens an email attachment that contains a virus, SSL Forward Proxy decryption will prevent the virus from infecting the client system and entering the corporate network.

With SSL Forward Proxy decryption, the firewall resides between the internal client and outside server. The firewall uses Forward Trust or Forward Untrust certificates to establish itself as a trusted third party to the session between the client and the server (For details on certificates, see [Keys and Certificates for Decryption Policies](#)). When the client initiates an SSL session with the server, the firewall intercepts the client's SSL request and forwards the SSL request to the server. The server sends a certificate intended for the client that is intercepted by the firewall. If the server's certificate is signed by a CA that the firewall trusts, the firewall creates a copy of the server's certificate signed by the Forward Trust certificate and sends the certificate to the client to authenticate. If the server's certificate is signed by a CA that the firewall does not trust, the firewall creates a copy of the server's certificate and signs it with the Forward Untrust certificate and sends it to the client. In this case, the client sees a block page warning that the site they're attempting to connect to is not trusted and the client can choose to proceed or terminate the session. When the client authenticates the certificate, the SSL session is established with the firewall functioning as a trusted forward proxy to the site that the client is accessing.

As the firewall continues to receive SSL traffic from the server that is destined for the client, it decrypts the SSL traffic into clear text traffic and applies security policies to the traffic. The traffic is then re-encrypted on the firewall and the firewall forwards the encrypted traffic to the client.

[Figure: SSL Forward Proxy](#) shows this process in detail.

Figure: SSL Forward Proxy

See [Configure SSL Forward Proxy](#) for details on configuring SSL Forward Proxy.

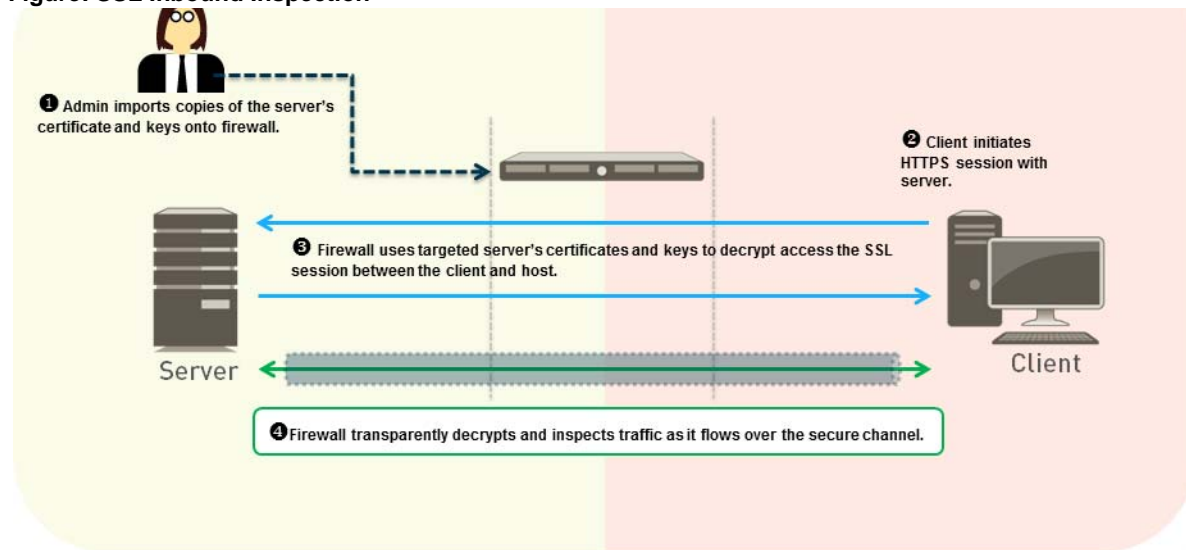
SSL Inbound Inspection

Use SSL Inbound Inspection to decrypt and inspect inbound SSL traffic from a client to a targeted server (any server you have the certificate for and can import it onto the firewall). For example, if an employee is remotely connected to a web server hosted on the company network and is attempting to add restricted internal documents to his Dropbox folder (which uses SSL for data transmission), SSL Inbound Inspection can be used to ensure that the sensitive data does not move outside the secure company network by blocking or restricting the session.

Configuring SSL Inbound Inspection includes importing the targeted server's certificate and key on to the firewall. Because the targeted server's certificate and key are imported on the firewall, the firewall is able to access the SSL session between the server and the client and decrypt and inspect traffic transparently, rather than functioning as a proxy. The firewall is able to apply security policies to the decrypted traffic, detecting malicious content and controlling applications running over this secure channel.

Figure: [SSL Inbound Inspection](#) shows this process in detail.

Figure: SSL Inbound Inspection



See [Configure SSL Inbound Inspection](#) for details on configuring SSL Inbound Inspection.

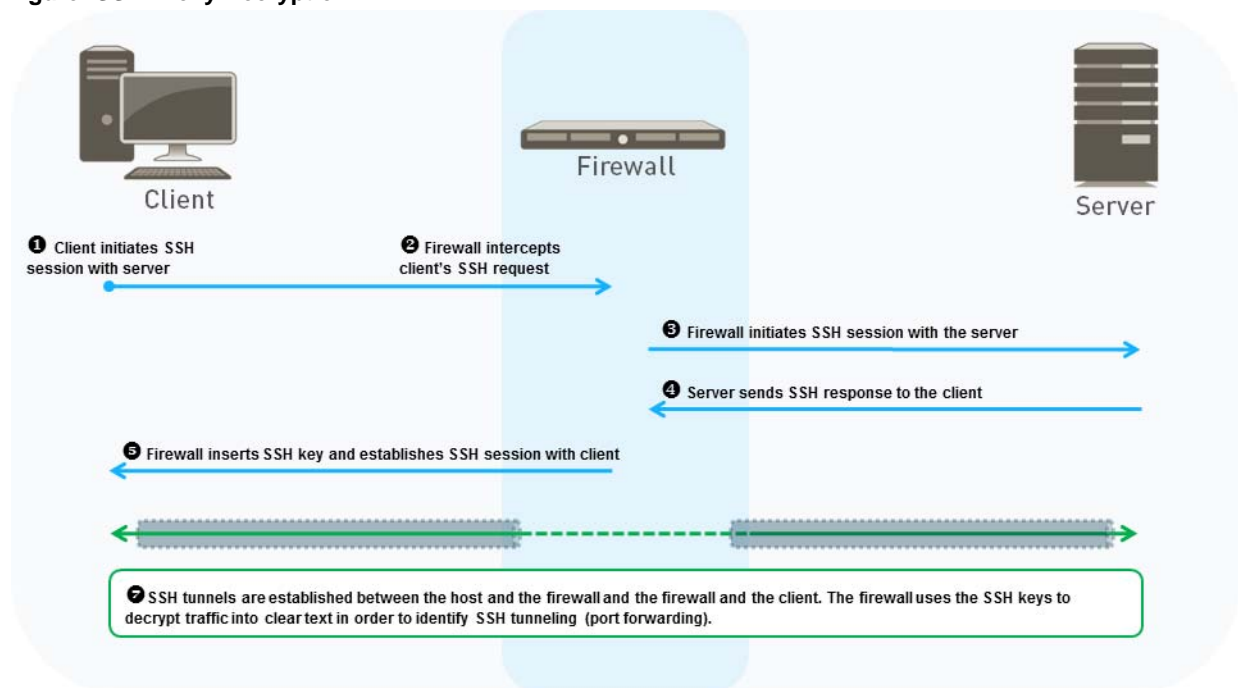
SSH Proxy

SSH Proxy provides the capability for the firewall to decrypt inbound and outbound SSH connections passing through the firewall, in order to ensure that SSH is not being used to tunnel unwanted applications and content. SSH decryption does not require any certificates and the key used for SSH decryption is automatically generated when the firewall boots up. During the boot up process, the firewall checks to see if there is an existing key. If not, a key is generated. This key is used for decrypting SSH sessions for all virtual systems configured on the device. The same key is also used for decrypting all SSH v2 sessions.

In an SSH Proxy configuration, the firewall resides between a client and a server. When the client sends an SSH request to the server, the firewall intercepts the request and forwards the SSH request to the server. The firewall then intercepts the server's response and forwards the response to the client, establishing an SSH tunnel between the firewall and the client and an SSH tunnel between the firewall and the server, with firewall functioning as a proxy. As traffic flows between the client and the server, the firewall is able to distinguish whether the SSH traffic is being routed normally or if it is using SSH tunneling (port forwarding). Content and threat inspections are not performed on SSH tunnels; however, if SSH tunnels are identified by the firewall, the SSH tunneled traffic is blocked and restricted according to configured security policies.

Figure: SSH Proxy Decryption shows this process in detail.

Figure: SSH Proxy Decryption



See [Configure SSH Proxy](#) for details on configuring an SSH Proxy policy.

Decryption Exceptions

Traffic can also be excluded from decryption according to matching criteria (using a decryption policy), a targeted server's traffic can be excluded from decryption (using certificates), and some applications are excluded from decryption by default.

Applications that do not function properly when decrypted by the firewall and are automatically excluded from SSL decryption. The applications that are excluded from SSL decryption by default are excluded because these applications often fail when decrypted due to the application looking for specific details in the certificate that might not be present in the certificate generated for SSL Forward Proxy. Refer to the KB article [List of Applications Excluded from SSL Decryption](#) for a current list of applications excluded by default from SSL decryption on the firewall.

You can configure decryption exceptions for certain URL categories or applications that either do not work properly with decryption enabled or for any other reason, including for legal or privacy purposes. You can use a decryption policy to exclude traffic from decryption based on source, destination, service, and URL category. For example, with SSL decryption enabled, you can exclude traffic that is categorized as financial or health-related from decryption, using the URL category selection.

You can also exclude servers from SSL decryption based on the Common Name (CN) in the server's certificate. For example, if you have SSL decryption enabled but have certain servers that you do not want included in SSL decryption, such as the web services for your HR systems, you can exclude those servers from decryption by importing the server certificate onto the firewall and modifying the certificate to be an **SSL Exclude Certificate**.

To exclude traffic from decryption based on source, destination, service, URL category or to exclude a specific server's traffic from decryption, see [Configure Decryption Exceptions](#).

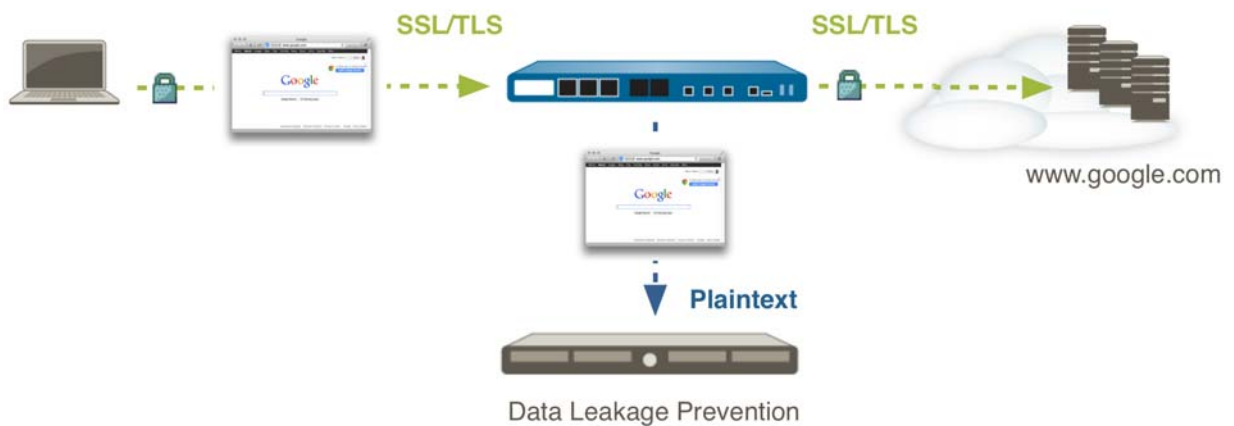
Decryption Port Mirroring

The Decryption Port mirror feature provides the capability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures—such as NetWitness or Solera—for archiving and analysis. This feature is necessary for organizations that require comprehensive data capture for forensic and historical purposes or data leak prevention (DLP) functionality. Decryption port mirroring is available on PA-7050, PA-5000 Series and PA-3000 Series platforms only and requires that a free license be installed to enable this feature.

Keep in mind that the decryption, storage, inspection, and/or use of SSL traffic is governed in certain countries and user consent might be required in order to use the decryption port mirror feature. Additionally, use of this feature could enable malicious users with administrative access to the firewall to harvest usernames, passwords, social security numbers, credit card numbers, or other sensitive information submitted using an encrypted channel. Palo Alto Networks recommends that you consult with your corporate council before activating and using this feature in a production environment.

[Figure: Decryption Port Mirroring](#) shows the process for decryption port mirroring and the section [Configure Decryption Port Mirroring](#) describes how to license and use this feature.

Figure: Decryption Port Mirroring



Configure SSL Forward Proxy

Configuring [SSL Forward Proxy](#) decryption on the firewall requires setting up the certificates needed for SSL Forward Proxy decryption and creating an SSL Forward Proxy decryption policy. The firewall can use self-signed certificates or certificates signed by an enterprise CA to perform SSL Forward Proxy decryption.



By default, the firewall determines the key size to use for the client certificates it generates based on the key size of the destination server certificate. You can optionally set a static key size to use regardless of the key size of the destination server certificate. See [Configure the Key Size for SSL Forward Proxy Server Certificates](#).

Use the following task to configure SSL Forward Proxy, including how to set up the certificates and create a decryption policy.

Configure SSL Forward Proxy	
<p>Step 1 Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces.</p>	<p>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.</p>

Configure SSL Forward Proxy

Step 2 Configure the forward trust certificate. Use either a self-signed certificate or a certificate signed by an enterprise CA.

Using self-signed certificates

When the certificate of the server that the client is connecting to is signed by a CA that is on the firewall's trusted CA list, the firewall signs a copy of the server's certificate with a self-signed forward trust certificate to present to the client for authentication. In this case, the self-signed certificate must be imported onto each client system so that the client recognizes the firewall as a trusted CA.

Use self-signed certificates for SSL Forward Proxy decryption if you do not use an enterprise CA or if you are only intended to perform decryption for a limited number of client systems (or if you are planning to use a centralized deployment).

To use a self-signed certificate:

1. Select **Device > Certificate Management > Certificates**.
2. Click **Generate** at the bottom of the window.
3. Enter a **Certificate Name**, such as *my-fwd-trust*.
4. Enter a **Common Name**, such as **192.168.2.1**. This should be the IP or FQDN that will appear in the certificate. In this case, we are using the IP of the trust interface. Avoid using spaces in this field.
5. Leave the **Signed By** field blank.
6. Click the **Certificate Authority** check box to enable the firewall to issue the certificate. Selecting this check box creates a certificate authority (CA) on the firewall that is imported to the client browsers, so clients trust the firewall as a CA.
7. Click **Generate** to generate the certificate.
8. Click the new certificate *my-fwd-trust* to modify it and enable the **Forward Trust Certificate** option.
9. Export the forward trust certificate for import into client systems by highlighting the certificate and clicking **Export** at the bottom of the window. Choose PEM format, and do not select the **Export private key** option. Because the certificate is self-signed, import it into the browser trusted root CA list on the client systems in order for the clients to trust it. When importing to the client browser, ensure the certificate is added to the Trusted Root Certification Authorities certificate store. On Windows systems, the default import location is the Personal certificate store. You can also simplify this process by using a centralized deployment, such as an Active Directory Group Policy Object (GPO).

 If the forward trust certificate is not imported on the client systems, users will see certificate warnings for each SSL site they visit.
10. Click **OK** to save.

Configure SSL Forward Proxy



Using an Enterprise CA



An enterprise CA can issue a signing certificate which the firewall can use to then sign the certificates for sites requiring SSL decryption. Send a Certificate Signing Request (CSR) for the enterprise CA to sign and validate. The firewall can then use the signed enterprise CA certificate for SSL Forward Proxy decryption. Because the enterprise CA is already trusted by the client systems, with this option, you do not need to distribute the certificate to client systems prior to configuring decryption.

To use an enterprise CA signed certificate, generate a CSR:

1. Select **Device > Certificate Management > Certificates** and click **Generate**.
2. Enter a **Certificate Name**, such as *my-fwd-proxy*.
3. In the **Signed By** drop-down, select **External Authority (CSR)**.
4. (Optional) If your enterprise CA requires it, add **Certificate Attributes** to further identify the firewall details, such as Country or Department.
5. Click **OK** to save the CSR. The pending certificate is now displayed on the **Device Certificates** tab.
6. Export the CSR:
 - a. Select the pending certificate displayed on the **Device Certificates** tab.
 - b. Click **Export** to download and save the certificate file.

 Leave **Export private key** unselected in order to ensure that the private key remains securely on the firewall.
 - c. Click **OK**.
 - d. Provide the certificate file to your enterprise CA. When you receive the signed enterprise CA certificate from your enterprise CA, save the signed enterprise CA certificate for import onto the firewall.
7. Import the signed enterprise CA onto the firewall:
 - a. Select **Device > Certificate Management > Certificates** and click **Import**.
 - b. Enter the pending **Certificate Name** exactly (in this case, *my-fwd-trust*). The **Certificate Name** that you enter must exactly match the pending certificate's name in order for the pending certificate to be validated.
 - c. Select the signed **Certificate File** that you received from your enterprise CA.
 - d. Click **OK**. The certificate is displayed as valid with the Key and CA check boxes selected.
 - e. Select the validated certificate, in this case, *my-fwd-proxy*, to enable it as a **Forward Trust Certificate** to be used for SSL Forward Proxy decryption.
 - f. Click **OK**.

Configure SSL Forward Proxy	
<p>Step 3 Configure the forward untrust certificate.</p> <p>With SSL Forward Proxy decryption, when the site the client is connecting to uses a certificate signed by a CA that is not in the firewall's trusted CA list, the firewall presents a forward untrust certificate to the client. The forward untrust certificate ensures that clients are prompted with a certificate warning when attempting to access sites with untrusted certificates.</p>	<ol style="list-style-type: none"> 1. Click Generate at the bottom of the certificates page. 2. Enter a Certificate Name, such as <i>my-fwd-untrust</i>. 3. Set the Common Name, for example 192.168.2.1. Leave Signed By blank. 4. Click the Certificate Authority check box to enable the firewall to issue the certificate. 5. Click Generate to generate the certificate. 6. Click OK to save. 7. Click the new <i>my-ssl-fw-untrust</i> certificate to modify it and enable the Forward Untrust Certificate option. <ul style="list-style-type: none">  Do not export the forward untrust certificate for import into client systems. If the forward trust certificate is imported on client systems, the users will not see certificate warnings for SSL sites with untrusted certificates. 8. Click OK to save.
<p>Step 4 (Optional) Set the key size of the SSL Forward Proxy certificates that the firewall presents to clients.</p> <p> Changing the key size clears the current certificate cache.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Session and, in the Decryption Settings section, click Forward Proxy Server Certificate Settings. 2. Select a Key Size: Defined by destination host (default), 1024-bit RSA, or 2048-bit RSA.
<p>Step 5 (Optional) Create a Decryption profile.</p> <p>Decryption profiles can be associated with a decryption policy, enabling the firewall to block and control various aspects of traffic that is being decrypted. An SSL Forward Proxy decryption profile can be used to perform checks for server certificates, unsupported modes, and failures and block or restrict traffic accordingly. For a complete list of checks that can be performed, navigate to Objects > Decryption Profiles on the firewall and click the help icon.</p>	<ol style="list-style-type: none"> 1. Select Objects > Decryption Profile and click Add. 2. Select the SSL Forward Proxy tab to block and control specific aspects of SSL tunneled traffic. For example, you can choose to terminate sessions if system resources are not available to process decryption by selecting Block sessions if resources not available. 3. Click OK to save the profile.



Configure SSL Forward Proxy	
<p>Step 6 Configure a decryption policy.</p>	<ol style="list-style-type: none"> 1. Select Policies > Decryption and click Add. 2. On the General tab, give the policy a descriptive Name. 3. On the Source and Destination tabs, select Any for the Source Zone and Destination Zone to decrypt all SSL traffic destined for an external server. If you want to specify traffic from or to certain sources or destinations for decryption, click Add. 4. In the URL Category tab, leave Any to decrypt all traffic. If you only want to apply this profile to certain website categories, click Add. <div style="margin-top: 10px;">  <p>Selecting a URL Category is useful when excluding certain sites from decryption. See Configure Decryption Exceptions.</p> </div> 5. On the Options tab, select Decrypt and select SSL Forward Proxy as the Type of decryption to perform. 6. (Optional) Select a Decryption Profile to apply additional settings to decrypted traffic (see Step 5). 7. Click OK to save.
<p>Step 7 Enable the firewall to forward decrypted SSL traffic for WildFire analysis.</p> <div style="margin-top: 10px;">  <p>This is a WildFire best practice. To forward portable executables (PEs) only, you do not need a WildFire license; however forwarding advanced file types requires an active WildFire license.</p> </div>	<p>On a firewall with no virtual systems configured:</p> <ol style="list-style-type: none"> 1. Select Device > Setup > Content-ID. 2. Edit the Content-ID settings and Allow Forwarding of Decrypted Content. 3. Click OK to save the changes. <p>On a firewall with multiple virtual systems configured:</p> <p>Select Device > Virtual Systems, select the virtual system you want to modify, and Allow Forwarding of Decrypted Content.</p>
<p>Step 8 Commit the configuration.</p>	<p>With an SSL Forward Proxy decryption policy enabled, all traffic identified by the policy is decrypted. Decrypted traffic is blocked and restricted according to the profiles configured on the firewall (including the decryption profiles associated with the policy and Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles). Traffic is re-encrypted as it exits the firewall.</p>

Configure SSL Inbound Inspection

Configuring [SSL Inbound Inspection](#) includes installing the targeted server's certificate on the firewall and creating an SSL Inbound Inspection decryption policy.

Use the following task to configure SSL Inbound Inspection.

Configure SSL Inbound Inspection	
<p>Step 1 Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces.</p>	<p>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.</p>
<p>Step 2 Ensure that the targeted server's certificate is installed on the firewall.</p>	<p>On the web interface, select Device > Certificate Management > Certificates > Device Certificates to view certificates installed on the firewall.</p> <p>To import the targeted server's certificate onto the firewall:</p> <ol style="list-style-type: none"> 1. On the Device Certificates tab, select Import. 2. Enter a descriptive Certificate Name. 3. Browse for and select the targeted server's Certificate File. 4. Click OK.
<p>Step 3 (Optional) Create a Decryption profile.</p> <p>Decryption profiles can be associated with a decryption policy, enabling the firewall to block and control various aspects of traffic that is being decrypted. An SSL Inbound Inspection decryption profile can be used to perform checks for unsupported modes and failures and block or restrict traffic accordingly. For a complete list of checks that can be performed, select Objects > Decryption Profiles and then click the help icon.</p>	<ol style="list-style-type: none"> 1. Select Objects > Decryption Profile and click Add. 2. Select the SSL Inbound Inspection tab to block and control specific aspects of SSL traffic. For example, you can choose to terminate sessions if system resources are not available to process decryption by selecting Block sessions if resources not available. 3. Click OK to save the profile.

Configure SSL Inbound Inspection	
<p>Step 4 Configure a decryption policy.</p>	<ol style="list-style-type: none"> 1. Select Policies > Decryption and click Add. 2. On the General tab, give the policy a descriptive Name. 3. On the Destination tab, Add the Destination Address of the targeted server. 4. In the URL Category tab, leave Any to decrypt all traffic. If you only want to apply this profile to certain website categories, click Add. <div style="margin-top: 10px;">  <p>Selecting a URL Category is useful when excluding certain sites from decryption. See Configure Decryption Exceptions.</p> </div> 5. On the Options tab, select Decrypt and select SSL Inbound Inspection as the Type of traffic to decrypt. Select the Certificate for the internal server that is the destination of the inbound SSL traffic. 6. (Optional) Select a Decryption Profile to apply additional settings to decrypted traffic. 7. Click OK to save.
<p>Step 5 Enable the firewall to forward decrypted SSL traffic for WildFire analysis.</p> <div style="margin-top: 10px;">  <p>This is a WildFire best practice. To forward portable executables (PEs) only, you do not need a WildFire license; however forwarding advanced file types requires an active WildFire license.</p> </div>	<p>On a firewall with no virtual systems configured:</p> <ol style="list-style-type: none"> 1. Select Device > Setup > Content-ID. 2. Edit the Content-ID settings and Allow Forwarding of Decrypted Content. 3. Click OK to save the changes. <p>On a firewall with multiple virtual systems configured:</p> <p>Select Device > Virtual Systems, select the virtual system you want to modify, and Allow Forwarding of Decrypted Content.</p>
<p>Step 6 Commit the configuration.</p>	<p>With an SSL Inbound Inspection decryption policy enabled, all SSL traffic identified by the policy is decrypted and inspected. Decrypted traffic is blocked and restricted according to the profiles configured on the firewall (including the decryption profiles associated with the policy and Antivirus, Vulnerability, Anti-Spyware, URL Filtering, and File-Blocking profiles). Traffic is re-encrypted as it exits the firewall.</p>

Configure SSH Proxy

Configuring **SSH Proxy** does not require certificates and the key used to decrypt SSH sessions is generated automatically on the firewall during boot up.

Use the following task to configure SSH Proxy decryption.

Configure SSH Proxy Decryption	
<p>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.</p>	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.
<p>Step 2 (Optional) Create a Decryption profile.</p> <p>Decryption profiles can be associated with a decryption policy, enabling the firewall to block and control various aspects of traffic that is being decrypted. The decryption profile can be used to perform checks for server certificates, unsupported modes, and failures and block or restrict traffic accordingly. For a complete list of checks that can be performed, navigate to Objects > Decryption Profiles on the firewall and then click the help icon.</p>	<ol style="list-style-type: none"> 1. Select Objects > Decryption Profile and click Add. 2. Select the SSH tab to block and control specific aspects of SSH tunneled traffic. For example, you can choose to terminate sessions if system resources are not available to process decryption by selecting Block sessions if resources not available. 3. Click OK to save the profile.
<p>Step 3 Configure a decryption policy.</p>	<ol style="list-style-type: none"> 1. Select Policies > Decryption and click Add. 2. On the General tab, give the policy a descriptive Name. 3. On the Source and Destination tabs, select Any to decrypt all SSH traffic. 4. On the URL Category tab, select Any to decrypt all SSH traffic. 5. On the Options tab, select Decrypt and select SSH Proxy as the Type of traffic to decrypt. 6. (Optional) Select a Decryption Profile to apply additional settings to decrypted traffic. 7. Click OK to save.
<p>Step 4 Commit the configuration.</p>	With the an SSH Proxy decryption policy enabled, all SSH traffic identified by the policy is decrypted and identified as either regular SSH traffic or as SSH tunneled traffic. SSH tunneled traffic is blocked and restricted according to the profiles configured on the firewall. Traffic is re-encrypted as it exits the firewall.

Configure Decryption Exceptions

You can purposefully exclude traffic from decryption based on matching criteria, such the traffic’s source, destination, URL category, or service. You can also exclude a specific server’s traffic from decryption. See the following topics to configure [Decryption Exceptions](#):

- ▲ [Exclude Traffic From Decryption](#)
- ▲ [Exclude a Server From Decryption](#)

Exclude Traffic From Decryption

To purposefully exclude applications or certain traffic from other existing SSL or SSH decryption policies, you can create a new decryption policy that defines the traffic to exclude from decryption with the **No Decrypt** action selected in the policy. You can define traffic for policy-based exclusion according to matching criteria, such as source, destination, URL categories, or the service (port or protocol). Make sure the decryption policy that excludes traffic from decryption is listed first in your decryption policy list by dragging and dropping the policy above the other decryption policies.

See the following procedure to configure a decryption policy that excludes traffic from SSL or SSH decryption.

Exclude Traffic from a Decryption Policy	
<p>Step 1 Create a decryption policy.</p> <p>Use a decryption policy to exclude traffic from decryption according to the traffic’s source and destination zones or addresses and URL categories. This example shows how to exclude traffic categorized as financial or health-related from SSL Forward Proxy decryption.</p>	<ol style="list-style-type: none"> 1. Go to Policies > Decryption and click Add. 2. Give the policy a descriptive Name, such as <i>No-Decrypt-Finance-Health</i>. 3. On the Source and Destination tabs, select Any for the Source Zone and Destination Zone to apply the <i>No-Decrypt-Finance-Health</i> rule to all SSL traffic destined for an external server. 4. On the URL Category tab, Add the URL categories financial-services and health-and-medicine to the policy, specifying that traffic that matches these categories will not be decrypted. 5. On the Options tab, select No Decrypt and select the Type of decryption policy you are excluding the traffic from. For example, to exclude traffic categorized as financial or health-related from a separately configured SSL Forward Proxy decryption policy, select SSL Forward Proxy as the Type. 6. Click OK to save the <i>No-Decrypt-Finance-Health</i> decryption policy.
<p>Step 2 Move the decryption policy to the top of the list of decryption policies.</p>	<p>On the Decryption > Policies page, select the policy <i>No-Decrypt-Finance-Health</i>, and click Move Up until it appears at the top of the list (or you can drag and drop).</p> <p>The order in which the decryption policies are listed is the order in which they are applied to network traffic. Moving the policy with the No Decrypt action applied to the top of the list ensures that the specified traffic is not decrypted according to another configured policy.</p>

Exclude Traffic from a Decryption Policy

Step 3 **Commit** the configuration.

A decryption policy with **No Decrypt** enabled ensures that the specified traffic is remains encrypted as it flows through the firewall, and that the traffic is not decrypted according to other decryption policies configured and listed on the **Policies > Decryption** page.

Exclude a Server From Decryption

You can exclude a targeted server's traffic from SSL decryption based on the Common Name (CN) in the server's certificate. For example, if you have SSL decryption enabled, you could configure a decryption exception for the server on your corporate network that hosts the web services for your HR systems. See the following procedure to configure a server's certificate so that the targeted server's traffic is excluded from decryption:

Exclude a Server from Decryption

Step 1 Import the targeted server's certificate onto the firewall:

1. On the **Device > Certificate Management > Certificates > Device Certificates** tab, select **Import**.
2. Enter a descriptive **Certificate Name**.
3. Browse for and select the targeted server's **Certificate File**.
4. Click **OK**.

Step 2 Select the targeted server's certificate on the **Device Certificates** tab and enable it as an SSL Exclude Certificate.

With the targeted server's certificate imported on the firewall and designated as an SSL Exclude Certificate, the server's traffic is not decrypted as it passes through the firewall.

Enable Users to Opt Out of SSL Decryption

In some cases, you might need to alert your users to the fact that the firewall is decrypting certain web traffic and allow them to terminate sessions that they do not want inspected. With SSL Opt Out enabled, the first time a user attempts to browse to an HTTPS site or application that matches your decryption policy, the firewall displays a response page notifying the user that it will decrypt the session. Users can either click **Yes** to allow decryption and continue to the site or click **No** to opt out of decryption and terminate the session. The choice to allow decryption applies to all HTTPS sites that users try to access for the next 24 hours, after which the firewall 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.

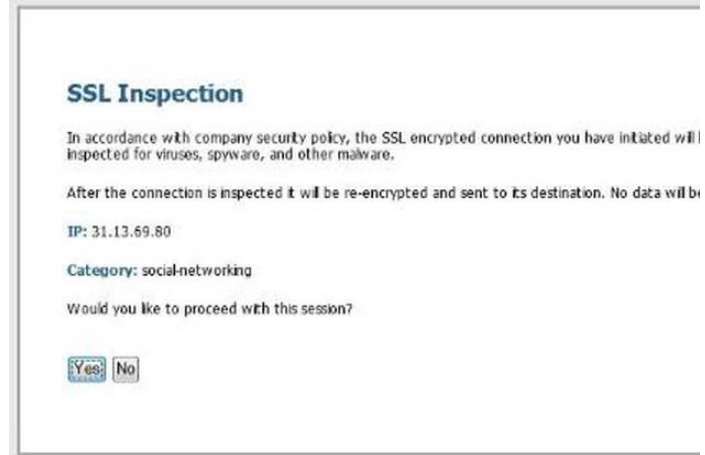
Enable Users to Opt Out of SSL Decryption	
<p>Step 1 (Optional) Customize the SSL Decryption Opt-out Page.</p>	<ol style="list-style-type: none"> 1. Select Device > Response Pages. 2. Select the SSL Decryption Opt-out Page link. 3. Select the Predefined page and click Export. 4. Using the HTML text editor 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: <pre></pre> 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.
<p>Step 2 Enable SSL Decryption Opt Out.</p>	<ol style="list-style-type: none"> 1. On the Device > Response Pages page, click the Disabled link. 2. Select the Enable SSL Opt-out Page and click OK. <div data-bbox="737 1581 1247 1728" data-label="Image"> </div> 3. Commit the changes.

Enable Users to Opt Out of SSL Decryption

Step 3 Verify that the Opt Out page displays when you attempt to browse to a site.

From a browser, go to an encrypted site that matches your decryption policy.

Verify that the SSL Decryption Opt-out response page displays.



Configure Decryption Port Mirroring

Before you can enable decryption port mirroring, you must obtain and install a Decryption Port Mirror license. The license is free of charge and can be activated through the support portal as described in the following procedure. After you install the Decryption Port Mirror license and reboot the firewall, you can enable decryption port mirroring.

Enabling decryption port mirroring includes enabling the forwarding of decrypted traffic and configuring a decrypt mirror interface. You can then create a decryption profile that specifies the interface and attach it to a decryption policy. To learn more about implementing Decryption Port Mirroring, see [Decryption Port Mirroring](#).

Use the following procedure to obtain and install a Decryption Port Mirror license and configure Decryption Port Mirroring.

Configure Decryption Port Mirroring

Step 1 Request a license for each device on which you want to enable decryption port mirroring.

1. Log in to the [Palo Alto Networks Support](#) site and navigate to the **Assets** tab.
2. Select the device entry for the device you want to license and select **Actions**.
3. Select **Decryption Port Mirror**. A legal notice displays.
4. If you are clear about the potential legal implications and requirements, click **I understand and wish to proceed**.
5. Click **Activate**.

DEVICE LICENSES

DEVICE LICENSES

Serial Number: 0009C100103
 Model: PAN-PA-5050-B
 Device Name: PM Lab Firewall

Authorization Code: **Add** 

Feature Name	Authorization Code	Expiration Date	Actions
Threat Prevention	I4344239	01/06/2019	
PAN-DB URL Filtering	I9544847	01/06/2019	
Virtual Systems	I8729162	Perpetual	
Premium Support	I7480971	12/29/2015	

AVAILABLE FEATURE LICENSES

Decryption Port Mirror

Configure Decryption Port Mirroring (Continued)									
<p>Step 2 Install the Decryption Port Mirror license on firewall.</p>	<ol style="list-style-type: none"> From the firewall's web interface, select Device > Licenses. Click Retrieve license keys from license server. Verify that the license has been activated on the firewall. <div data-bbox="841 386 1279 527" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p style="text-align: center; background-color: #4f81bd; color: white; padding: 2px;">Decryption Port Mirror</p> <table style="width: 100%; border-collapse: collapse;"> <tr> <td style="padding: 2px;">Date Issued</td> <td style="padding: 2px;">January 06, 2014</td> </tr> <tr> <td style="padding: 2px;">Date Expires</td> <td style="padding: 2px;">Never</td> </tr> <tr> <td style="padding: 2px;">Description</td> <td style="padding: 2px;">Decryption Port Mirror</td> </tr> <tr> <td style="padding: 2px;">Active</td> <td style="padding: 2px;">Yes</td> </tr> </table> </div> Reboot the firewall (Device > Setup > Operations). This feature will not be available for configuration until PAN-OS reloads. 	Date Issued	January 06, 2014	Date Expires	Never	Description	Decryption Port Mirror	Active	Yes
Date Issued	January 06, 2014								
Date Expires	Never								
Description	Decryption Port Mirror								
Active	Yes								
<p>Step 3 Enable the ability to mirror decrypted traffic. Superuser permission is required to perform this step.</p>	<p>On a firewall with a single virtual system:</p> <ol style="list-style-type: none"> Select Device > Setup > Content - ID. Select the Allow forwarding of decrypted content check box. Click OK to save. <p>On a firewall with multiple virtual systems:</p> <ol style="list-style-type: none"> Select Device > Virtual System. Select a Virtual System to edit or create a new Virtual System by selecting Add. Select the Allow forwarding of decrypted content check box. Click OK to save. 								
<p>Step 4 Configure a decrypt mirror interface.</p>	<ol style="list-style-type: none"> Select Network > Interfaces > Ethernet. Select the Ethernet interface that you want to configure for decryption port mirroring. Select Decrypt Mirror as the Interface Type. This interface type will appear only if the Decryption Port Mirror license is installed. Click OK to save. 								

Configure Decryption Port Mirroring (Continued)	
<p>Step 5 Configure a Decryption Profile to enable decryption port mirroring.</p>	<ol style="list-style-type: none"> 1. Select Objects > Decryption Profile. 2. Select the Interface to use 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.
<p>Step 6 Set a decryption policy for decryption port mirroring.</p>	<ol style="list-style-type: none"> 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.
<p>Step 7 Save the configuration.</p>	<p>Click Commit.</p>

Temporarily Disable SSL Decryption

In some cases you may want to temporarily disable SSL decryption. For example, if your users are having problems accessing an encrypted site or application, you may want to disable SSL decryption in order to troubleshoot the issue. Although you could disable the associated decryption policies, modifying the policies is a configuration change that requires a Commit. Instead, use the following command to temporarily disable SSL decryption and then re-enable it after you finish troubleshooting. This command does not require a commit and it does not persist in your configuration after a reboot.

Temporarily Disable SSL Decryption	
• Disable SSL Decryption	<code>set system setting ssl-decrypt skip-ssl-decrypt yes</code>
• Re-enable SSL Decryption	<code>set system setting ssl-decrypt skip-ssl-decrypt no</code>



URL Filtering

The Palo Alto Networks URL filtering solution is a powerful PAN-OS feature that is used to monitor and control how users access the web over HTTP and HTTPS. The following topics provide an overview of URL filtering, configuration and troubleshooting information, and best practices for getting the most out of this feature:

- ▲ [URL Filtering Overview](#)
- ▲ [URL Filtering Concepts](#)
- ▲ [PAN-DB Categorization Workflow](#)
- ▲ [Enable a URL Filtering Vendor](#)
- ▲ [Determine URL Filtering Policy Requirements](#)
- ▲ [Monitor Web Activity](#)
- ▲ [Configure URL Filtering](#)
- ▲ [Log Client IP Addresses for Proxied HTTP/HTTPS Requests](#)
- ▲ [Customize the URL Filtering Response Pages](#)
- ▲ [Configure URL Admin Override](#)
- ▲ [Enable Safe Search Enforcement](#)
- ▲ [URL Filtering Use Case Examples](#)
- ▲ [Troubleshoot URL Filtering](#)

URL Filtering Overview

The Palo Alto Networks URL filtering feature compliments the [App-ID](#) feature by enabling you to configure your firewall to identify and control access to web (HTTP and HTTPS) traffic. By implementing URL filtering profiles in security policies and by using URL categories as a match criteria in policies (captive portal, decryption, security, and QoS), you will gain complete visibility and control of the traffic that traverses your firewall and will be able to safely enable and control how your users access the web.

The Palo Alto Networks URL filtering solution utilizes a URL filtering database that contains millions of websites and each website is placed in one of approximately 60 different categories. A URL filtering profile that contains the list of categories is then applied to a security policy that allows web traffic (HTTP/HTTPS) from the internal users to the Internet. After the URL filtering profile is applied and the alert or block action is set on a category, you will gain complete visibility into the websites that users access and can then decide which websites or website categories should be allowed, blocked, or logged. You can also define a list of URLs in the URL filtering profile that will always be blocked or allowed and you can create custom URL categories that contain a list of URLs that can be used the same way as the default category list. These same URL categories can also be used as a match criteria in other policies, such as captive portal, decryption, and QoS.

▲ URL Filtering Vendors

▲ Interaction Between App-ID and URL Categories

URL Filtering Vendors

Palo Alto Networks firewalls support two vendors for URL filtering purposes:

- **PAN-DB**—A Palo Alto Networks developed URL filtering database that is tightly integrated into PAN-OS by utilizing high-performance local caching to perform maximum inline performance for URL lookups while a distributed cloud architecture provides coverage for the latest websites. In addition, PAN-DB is tightly integrated with WildFire such that whenever WildFire deems a site malicious, it updates the corresponding PAN-DB URL category to malware, immediately blocking any future access to the site as long as you have a URL Filtering profile attached to the security policy rule. To view a list of PAN-DB URL filtering categories, refer to <https://urlfiltering.paloaltonetworks.com/CategoryList.aspx>.
- **BrightCloud**—A third-party URL database that is owned by Webroot, Inc. and is integrated into PAN-OS firewalls. For information on the BrightCloud URL database, visit <http://brightcloud.com>.

For instructions on configuring the firewall to use one of the supported URL Filtering vendors, see [Enable a URL Filtering Vendor](#).

Interaction Between App-ID and URL Categories

The Palo Alto Networks URL filtering feature in combination with Application Identification ([App-ID](#)) provides unprecedented protection against a full spectrum of legal, regulatory, productivity, and resource utilization risks. While App-ID gives you control over what applications users can access, URL filtering provides control over related web activity. When combined with User-ID, you can also apply these controls based on users and groups.

With today's application landscape and the way many applications use HTTP and HTTPS, you will need to determine when to use App-ID and when to use URL filtering in order to define comprehensive web access policies. In most cases, if an App-ID signature exists, you will want to use App-ID to control the content at the application level. For example, although you can control access to Facebook and/or LinkedIn using URL filtering, this would not block the use of all related applications, such as email, chat, as well as any new application that is introduced after you implement your policy.

In some cases, you will want to use both URL filtering and App-ID, but to ensure that conflicts do not occur in your policies, it is important to understand how these features work together. Palo Alto Networks generates signatures for many applications and those signatures can be very granular in regard to various features within the web-based applications—whereas URL filtering would only apply actions based on a specific website or URL category. For example, you may want to block social networking sites in general, but want to allow a few sites in that category to be accessible to specific departments and then control what features of the website are available to the allowed users. For more information, see [URL Filtering Use Case Examples](#).

URL Filtering Concepts

- ▲ [URL Categories](#)
- ▲ [URL Filtering Profile](#)
- ▲ [URL Filtering Profile Actions](#)
- ▲ [Block and Allow Lists](#)
- ▲ [Safe Search Enforcement](#)
- ▲ [Container Pages](#)
- ▲ [HTTP Header Logging](#)
- ▲ [URL Filtering Response Pages](#)
- ▲ [URL Category as Policy Match Criteria](#)

URL Categories

Each website defined in the URL filtering database is assigned one of approximately 60 different URL categories. There are two ways to make use of URL categorization on the firewall:

- Block or allow traffic based on URL category**—You can create a URL Filtering Profile that specifies an action for each URL category and attach the profile to a policy. Traffic that matches the policy would then be subject to the URL filtering settings in the profile. For example, to block all gaming websites you would set the block action for the URL category *games* in the URL profile and attach it to the security policy rule(s) that allow web access. See [Configure URL Filtering](#) for more information.
- Match traffic based on URL category for policy enforcement**—If you want a specific policy rule to apply only to web traffic to sites in a specific category, you would add the category as match criteria when you create the policy rule. For example, you could use the URL category *streaming-media* in a QoS policy to apply bandwidth controls to all websites that are categorized as streaming media. See [URL Category as Policy Match Criteria](#) for more information.

By grouping websites into categories, it makes it easy to define actions based on certain types of websites. In addition to the standard URL categories, there are three additional categories:

Categories	Description
Not-resolved	<p>Indicates that the website was not found in the local URL filtering database and the firewall was unable to connect to the cloud database to check the category. When a URL category lookup is performed, the firewall first checks the dataplane cache for the URL; if no match is found, it will then check the management plane cache, and if no match is found there, it queries the URL database in the cloud.</p> <p>When deciding on what action to take for traffic that is categorized as <i>not-resolved</i>, be aware that setting the action to block may be very disruptive to users.</p> <p>For more information on troubleshooting lookup issues, see Troubleshoot URL Filtering.</p>
Private-ip-addresses	<p>Indicates that the website is a single domain (no sub-domains), the IP address is in the private IP range, or the URL root domain is unknown to the cloud.</p>
Unknown	<p>The website has not yet been categorized, so it does not exist in the URL filtering database on the firewall or in the URL cloud database.</p> <p>When deciding on what action to take for traffic categorized as <i>unknown</i>, be aware that setting the action to block may be very disruptive to users because there could be a lot of valid sites that are not in the URL database yet. If you do want a very strict policy, you could block this category, so websites that do not exist in the URL database cannot be accessed.</p> <p>See Configure URL Filtering.</p>

URL Filtering Profile

A URL filtering profile is a collection of URL filtering controls that are applied to individual security policy rules to enforce your web access policy. The firewall comes with a default profile that is configured to block threat-prone categories, such as malware, phishing, and adult. You can use the default profile in a security policy, clone it to be used as a starting point for new URL filtering profiles, or add a new URL profile that will have all categories set to allow for visibility into the traffic on your network. You can then customize the newly added URL profiles and add lists of specific websites that should always be blocked or allowed, which provides more granular control over URL categories. For example, you may want to block social-networking sites, but allow some websites that are part of the social-networking category.

The following topics describe the components of the URL Filtering profile:



- ▲ [URL Filtering Profile Actions](#)
- ▲ [Block and Allow Lists](#)
- ▲ [Safe Search Enforcement](#)
- ▲ [Container Pages](#)
- ▲ [HTTP Header Logging](#)

URL Filtering Profile Actions

The URL Filtering profile specifies an action for each known URL category. By default, when you create a new URL Filtering profile, the action is set to **allow** for all categories. This means that the users will be able to browse to all sites freely and will not be logged.



The default URL Filtering profile is set to allow access to all URL categories except for the following threat-prone categories, which are blocked: abused-drugs, adult, gambling, hacking, malware, phishing, questionable, and weapons. As a best practice, when you need a custom URL Filtering profile, clone the default profile rather than creating a new one to preserve these settings.

Action	Description
alert	The website is allowed and a log entry is generated in the URL filtering log.
allow	The website is allowed and no log entry is generated.
block	The website is blocked and the user will see a response page and will not be able to continue to the website. A log entry is generated in the URL filtering log.
continue	<p>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.</p> <p> The Continue page will not be displayed properly on client machines that are configured to use a proxy server.</p>
override	<p>The user will see a response page indicating that a password is required to allow access to websites in the given category. With this option, the security admin or helpdesk person would provide a password that will grant temporary access to all websites in the given category. A log entry is generated in the URL filtering log. See Configure URL Admin Override.</p> <p> The Override page will not be displayed properly on client machines that are configured to use a proxy server.</p>
none	<p>The none action only applies to custom URL categories. Select none to ensure that if multiple URL profiles exist, the custom category will not have any impact on other profiles. For example, if you have two URL profiles and the custom URL category is set to block in one of the profiles, the other profile should have the action set to none if you do not want it to apply.</p> <p>Also, in order to delete a custom URL category, it must be set to none in any profile where it is used.</p>

Block and Allow Lists

Block and allow lists allow you to define specific URLs or IP addresses in the URL filtering profile that are always allowed or always blocked, regardless of the action defined for the URL category. When entering URLs in the *Block List* or *Allow List*, enter each URL or IP address in a new row separated by a new line. When using wildcards in the URLs, follow these rules:

- Do not include HTTP and HTTPS when defining URLs. For example, enter `www.paloaltonetworks.com` or `paloaltonetworks.com` instead of `https://www.paloaltonetworks.com`.
- Entries in the block list must be an exact match and are case-insensitive.
For example: If you want to prevent a user from accessing any website within the domain `paloaltonetworks.com`, you would also enter `*.paloaltonetworks.com`, so whatever domain prefix (`http://`, `www`, or a sub-domain prefix such as `mail.paloaltonetworks.com`) is added to the address, the specified action will be taken. The same applies to the sub-domain suffix; if you want to block `paloaltonetworks.com/en/US`, you would need to add `paloaltonetworks.com/*` as well.
Further, if you want to specify that the match is for `paloaltonetworks.com` and not `paloaltonetworks.com.au`, you must add a `/` after `.com` to specify the scope for the match to just be `.com` and not for the domain suffix `.com.au`. In this case, you need to add the entry as `*.paloaltonetworks.com/`

Block and allow lists support wildcard patterns. Use the following characters as separators:

.
/
?
&
=
;
+

Every substring that is separated by the characters listed above is considered a token. A token can be any number of ASCII characters that does not contain any separator character or `*`. For example, the following patterns are valid:

`*.yahoo.com` (tokens are: `"*"`, `"yahoo"` and `"com"`)

`www*.com` (tokens are: `"www"`, `"*"` and `"com"`)

`www.yahoo.com/search=*` (tokens are: `"www"`, `"yahoo"`, `"com"`, `"search"`, `"*"`)

The following patterns are invalid because the character `"*"` is not the only character in the token.

`ww*.yahoo.com`

`www.y*.com`

Safe Search Enforcement

Many search engines have a safe search setting that filters out adult images and videos in search query return traffic. On the firewall, you can [Enable Safe Search Enforcement](#) so that the firewall will block search results if the end user is not using the strictest safe search settings in the search query. The firewall can enforce safe search for the following search providers: Google, Yahoo, Bing, Yandex, and YouTube. This is a best-effort setting and is not guaranteed by the search providers to work with every website.

To use this feature you must enable the **Safe Search Enforcement** option in a URL filtering profile and attach it to a security policy. The firewall will then block any matching search query return traffic that is not using the strictest safe search settings. There are two methods for blocking the search results:

- [Block Search Results that are not Using Strict Safe Search Settings](#)—When an end user attempts to perform a search without first enabling the strictest safe search settings, the firewall blocks the search query results and displays the URL Filtering Safe Search Block Page. By default, this page will provide a URL to the search provider settings for configuring safe search.
- [Enable Transparent Safe Search Enforcement](#)—When an end user attempts to perform a search without first enabling the strict safe search settings, the firewall blocks the search results with an HTTP 503 status code and redirects the search query to a URL that includes the safe search parameters. You enable this functionality by importing a new URL Filtering Safe Search Block Page containing the Javascript for rewriting the search URL to include the strict safe search parameters. In this configuration, users will not see the block page, but will instead be automatically redirected to a search query that enforces the strictest safe search options. This safe search enforcement method requires Content Release version 475 or later and is only supported for Google, Yahoo, and Bing searches.


Also, because most search providers now use SSL to return search results, you must also configure a [Decryption](#) policy for the search traffic to enable the firewall to inspect the search traffic and enforce safe search.



Safe search enforcement enhancements and support for new search providers is periodically added in content releases. This information is detailed in the Application and Threat Content Release Notes. How sites are judged to be safe or unsafe is performed by each search provider, not by Palo Alto Networks.

Safe search settings differ by search provider as detailed in [Table: Search Provider Safe Search Settings](#).

Table: Search Provider Safe Search Settings

Search Provider	Safe Search Setting Description
<p>Google/YouTube</p>	<p>Offers safe search on individual computers or network-wide through Google’s safe search virtual IP address:</p> <p>Safe Search Enforcement for Google Searches on Individual Computers</p> <p>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.</p> <p>Appending <code>safe=active</code> to a Google search query URL also enables the strictest safe search settings.</p> <p>Safe Search Enforcement for Google and YouTube Searches using a Virtual IP Address</p> <p>Alternatively, Google provides servers that Lock SafeSearch (<code>forcesafesearch.google.com</code>) settings in every Google and YouTube search. By adding a DNS entry for <code>www.google.com</code> and <code>www.youtube.com</code> (and other relevant Google and YouTube country subdomains) that includes a CNAME record pointing to <code>forcesafesearch.google.com</code> 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.</p> <p> 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 <code>www.google.com</code> and <code>www.youtube.com</code>, using the local IP address for the <code>forcesafesearch.google.com</code> server.</p>
<p>Yahoo</p>	<p>Offers safe search on individual computers only. The Yahoo Search Preferences include three SafeSearch settings: Strict, Moderate, or Off. When enabled, the setting is stored in a browser cookie as <code>vm=</code> and passed to the server each time the user performs a Yahoo search.</p> <p>Appending <code>vm=r</code> to a Yahoo search query URL also enables the strictest safe search settings.</p> <p> When performing a search on Yahoo Japan (<code>yahoo.co.jp</code>) while logged into a Yahoo account, end users must also enable the SafeSearch Lock option.</p>
<p>Bing</p>	<p>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>adlt=</code> and passed to the server each time the user performs a Bing search.</p> <p>Appending <code>adlt=strict</code> to a Bing search query URL also enables the strictest safe search settings.</p> <p>The Bing SSL search engine does not enforce the safe search URL parameters and you should therefore should consider blocking Bing over SSL for full safe search enforcement.</p>

Container Pages

A container page is the main page that a user accesses when visiting a website, but additional websites may be loaded within the main page. If the **Log Container page only** option is enabled in the URL filtering profile, only the main container page will be logged, not subsequent pages that may be loaded within the container page. Because URL filtering can potentially generate a lot of log entries, you may want to turn on this option, so log entries will only contain those URIs where the requested page file name matches the specific mime-types. The default set includes the following mime-types:

- application/pdf
- application/soap+xml
- application/xhtml+xml
- text/html
- text/plain
- text/xml



If you have enabled the **Log container page only** option, there may not always be a correlated URL log entry for threats detected by Antivirus or Vulnerability Protection.

HTTP Header Logging


URL filtering provides visibility and control over web traffic on your network. For improved visibility into web content, you can configure the URL Filtering profile to log HTTP header attributes included in a web request. When a client requests a web page, the HTTP header includes the user agent, referer, and x-forwarded-for fields as attribute-value pairs and forwards them to the web server. When enabled for logging HTTP headers, the firewall logs the following attribute-value pairs in the URL Filtering logs:

Attribute	Description
User-Agent	The web browser that the user used to access the URL, for example, Internet Explorer. This information is sent in the HTTP request to the server.
Referer	The URL of the web page that linked the user to another web page; it is the source that redirected (referred) the user to the web page that is being requested.
X-Forwarded-For	The header field option that preserves the IP address of the user who requested the web page. It allows you to identify the IP address of the user particularly if you have a proxy server on your network, where all requests might seem to originate from the proxy server's IP address.

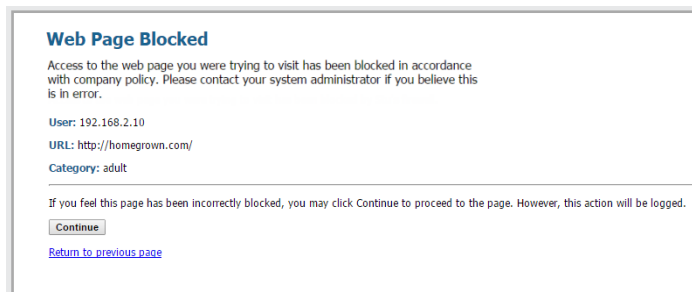
URL Filtering Response Pages

The firewall provides three predefined response pages that display by default when a user attempts to browse to a site in a category that is configured with one of the block actions in the [URL Filtering Profile](#) (block, continue, or override) or when [Safe Search Enforcement](#) is enabled:

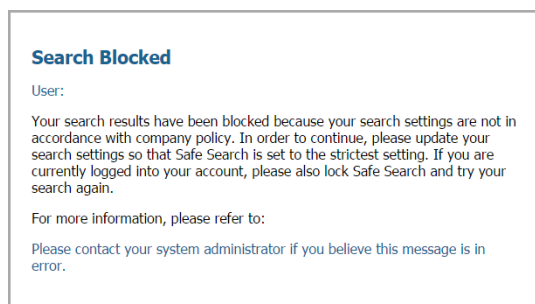
- **URL Filtering and Category Match Block Page**—Access blocked by a URL filtering profile or because the URL category is blocked by a security policy.



- **URL Filtering Continue and Override Page**—Page with initial block policy that allows users to bypass the block. With the override page, after clicking **Continue**, the user must supply a password to override the policy that blocks the URL.



- **URL Filtering Safe Search Block Page**—Access blocked by a security policy with a URL filtering profile that has the Safe Search Enforcement option enabled. The user will see this page if a search is performed using Google, Bing, Yahoo, or Yandex and their browser or search engine account setting for Safe Search is not set to strict.



You can either use the predefined pages, or you can [Customize the URL Filtering Response Pages](#) to communicate your specific acceptable use policies and/or corporate branding. In addition, you can use the [URL Filtering Response Page Variables](#) for substitution at the time of the block event or add one of the supported [Response Page References](#) to external images, sounds, or style sheets.

URL Filtering Response Page Variables

Variable	Usage
<user/>	The firewall replaces the variable with the username (if available via User-ID) or IP address of the user when displaying the response page.
<url/>	The firewall replaces the variable with the requested URL when displaying the response page.
<category/>	The firewall replaces the variable with the URL filtering category of the blocked request.
<pan_form/>	HTML code for displaying the continue button on the URL Filtering Continue and Override page.

You can also add code that triggers the firewall to display different messages using depending on what URL category the user is attempting to access. For example, the following code snippet from a response page specifies to display Message 1 if the URL category is games, Message 2 if the category as travel, or Message 3 if the category is kids:

```
var cat = "<category/>";
switch(cat)
{
  case 'games':
    document.getElementById("warningText").innerHTML = "Message 1";
    break;
  case 'travel':
    document.getElementById("warningText").innerHTML = "Message 2";
    break;
  case 'kids':
    document.getElementById("warningText").innerHTML = "Message 3";
    break;
}
```

Only a single HTML page can be loaded into each virtual system for each type of block page. However, other resources such as images, sounds, and cascading style sheets (CSS files) can be loaded from other servers at time the response page is displayed in the browser. All references must include a fully qualified URL.

Response Page References

Reference Type	Example HTML Code
Image	
Sound	<embed src="http://simplythebest.net/sounds/WAV/WAV_files/movie_WAV_files/do_not_go.wav" volume="100" hidden="true" autostart="true">
Style Sheet	<link href="http://example.com/style.css" rel="stylesheet" type="text/css" />
Hyperlink	View Corporate Policy

URL Category as Policy Match Criteria

URL categories can be used as a match criteria in a policy to provide more granularity in the policy. For example, you may have a decryption policy defined, but you would like specific websites to bypass decryption. To do this, you would configure a decryption policy with the *no-decrypt* action and a URL category would be defined as match criteria for the policy rule, so the policy would only match traffic flows to websites that are part of the specified category.

The following table describes the policy types that can utilize URL categories:

Policy Type	Description
Captive Portal	To ensure that users authenticate before being allowed access to a specific category, you can attach a URL category as a match criterion for the captive portal policy.
Decryption	Decryption policies can use URL categories as match criteria to determine if specified websites should be decrypted or not. For example, if you have a decryption policy with the action <i>decrypt</i> for all traffic between two zones, there may be specific website categories, such as <i>financial-services</i> and/or <i>health-and-medicine</i> , that should not be decrypted. In this case, you would create a new decryption policy with the action of <i>no-decrypt</i> that precedes the <i>decrypt</i> policy and then defines a list of URL categories as match criteria for the policy. By doing this, each URL category that is part of the <i>no-decrypt</i> policy will not be decrypted. You could also configure a custom URL category to define your own list of URLs that can then be used in the <i>no-decrypt</i> policy.
QoS	A QoS policy can use URL categories to determine throughput levels for specific website categories. For example, you may want to allow the streaming-media category, but limit throughput by adding the URL category as match criteria to the QoS policy.
Security	<p>In security policies you can use URL categories both as a match criteria in the Service/URL Category tab, and in URL filtering profiles that are attached in the Actions tab.</p> <p>If for example, the IT-security group in your company needs access to the hacking category, while all other users are denied access to the category, you must create the following rules:</p> <ul style="list-style-type: none"> • A security rule that allows the IT-Security group to access content categorized as hacking. The security rule references the <i>hacking</i> category in the Services/URL Category tab and IT-Security group in the Users tab. • Another security rule that allows general web access for all users. To this rule you attach a URL filtering profile that blocks the hacking category. <p>The policy that allows access to hacking must be listed before the policy that blocks hacking. This is because security policy rules are evaluated top down, so when a user who is part of the security group attempts to access a hacking site, the policy rule that allows access is evaluated first and will allow the user access to the hacking sites. Users from all other groups are evaluated against the general web access rule which blocks access to the hacking sites.</p>

PAN-DB Categorization Workflow

This section describes the PAN-DB components and describes the URL categorization resolution workflow that occurs as users access various URLs through the firewall.

- ▲ [PAN-DB URL Categorization Components](#)
- ▲ [PAN-DB URL Categorization Workflow](#)

PAN-DB URL Categorization Components

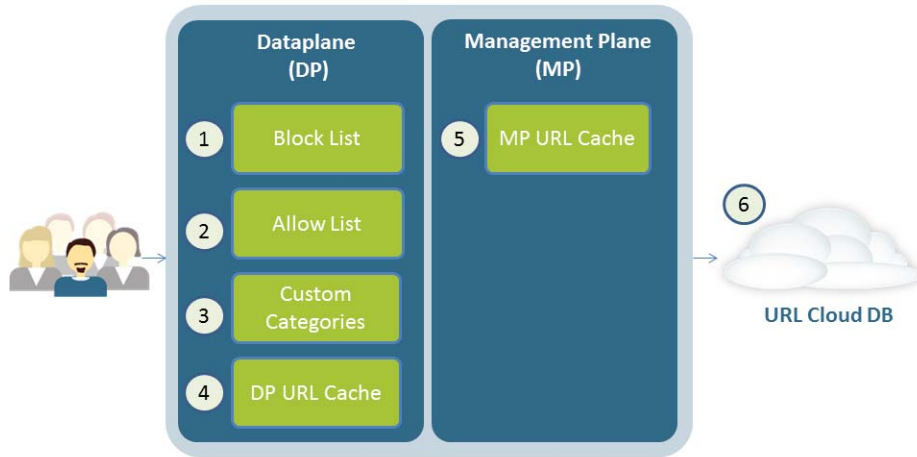
The following table describes the PAN-DB components in detail. The BrightCloud system works similarly, but does not use an initial seed database.

Component	Description
URL Filtering Seed Database	The initial seed database downloaded to the firewall is a small subset of the database that is maintained on the Palo Alto Networks URL cloud servers. The reason this is done is because the full database contains millions of URLs and many of these URLs may never be accessed by your users. When downloading the initial seed database, a region is selected (North America, Europe, APAC, Japan) and each region contains a subset of URLs most accessed for the given region. By doing this, the firewall will store a much smaller URL database, which greatly improves lookup performance. If a user accesses a website that is not in the local URL database, the full cloud database is queried and the firewall will then add the new URL to the local database. In other words, the local database on the firewall will be continually populated/customized based on user activity. Note that the local customized URL database will be cleared if PAN-DB seed database is re-downloaded or if you change the URL database vendor from PAN-DB to BrightCloud.
Cloud Service	<p>The PAN-DB cloud service is implemented using Amazon Web Services (AWS). AWS provides a distributed high-performance and stable environment for seed database downloads and URL lookups for Palo Alto Networks firewalls and communication is performed over SSL. The AWS cloud systems hold the entire PAN-DB and is updated as new URLs are identified. The PAN-DB cloud service supports an automated mechanism to update the firewall's local URL database if the version does not match. Each time the firewall queries the cloud servers for URL lookups, it will also check for critical updates. If there have been no queries to the cloud servers for more than 30 minutes, the firewall will check for updates on the cloud systems.</p> <p>The cloud system also provides a mechanism to submit URL category change requests. This is performed through the test-a-site service and is available directly from the device (URL filtering profile setup) and from the Palo Alto Networks Test A Site website. You can also submit a URL categorization change request directly from the URL filtering log on the firewall in the log details section.</p>

Component	Description
Management Plane (MP) URL Cache	<p>When PAN-DB is activated on the firewall, a seed database is downloaded from one of the PAN-DB cloud servers to initially populate the local cache, which is done to improve lookup performance. Each regional seed database contains the top URLs for the region and the size of the seed database (number of URL entries) also depends on the device platform. The URL MP cache is automatically written to the firewall's local drive every eight hours, before the firewall is rebooted, or when the cloud upgrades the URL database version on the firewall. After rebooting the firewall, the file that was saved to the local drive will be loaded to the MP cache. A least recently used (LRU) mechanism is also implemented in the URL MP cache in case the cache is full. If the cache becomes full, the URLs that have been accessed the least will be replaced by the newer URLs.</p>
Dataplane (DP) URL Cache	<p>This is a subset of the MP cache and is a customized, dynamic URL database that is stored in the dataplane (DP) and is used to improve URL lookup performance. The URL DP cache is cleared at each firewall reboot. The number of URLs that are stored in the URL DP cache varies by hardware platform and the current URLs stored in the TRIE (data structure). A least recently used (LRU) mechanism is implemented in the DP cache in case the cache is full. If the cache becomes full, the URLs that have been accessed the least will be replaced by the newer URLs. Entries in the URL DP cache expire after a specified period of time and the expiration period cannot be changed by the administrator.</p>

PAN-DB URL Categorization Workflow

When a user attempts to access a URL and the URL category needs to be determined, the firewall will compare the URL with the following components (in order) until a match has been found:



URL Categorization Resolution Process Flow

If a URL query matches an expired entry in the URL DP cache, the cache responds with the expired category, but also sends a URL categorization query to the management plane. This is done to avoid unnecessary delays in the DP, assuming that the frequency of changing categories is low. Similarly, in the URL MP cache, if a URL query from the DP matches an expired entry in the MP, the MP responds to the DP with the expired category and will also send a URL categorization request to the cloud service. Upon getting the response from the cloud, the firewall will resend the updated response to the DP.

As new URLs and categories are defined or if critical updates are needed, the cloud database will be updated. Each time the firewall queries the cloud for a URL lookup or if no cloud lookups have occurred for 30 minute, the database versions on the firewall be compared and if they do not match, an incremental update will be performed.

Enable a URL Filtering Vendor

To enable URL filtering on a firewall, you must purchase and activate a URL Filtering license for one of the supported [URL Filtering Vendors](#) and then install the database. Use one of the following procedures to enable a URL filtering license on the firewall set up the database. Only one URL filtering license can be active on a firewall.








Starting with PAN-OS 6.0, firewalls managed by Panorama do not need to be running the same URL filtering vendor that is configured on Panorama. For firewalls running PAN-OS 6.0 or later, when a mismatch is detected between the vendor enabled on the firewalls and what is enabled on Panorama, the firewalls can automatically migrate URL categories and/or URL profiles to (one or more) categories that align with that of the vendor enabled on it. For guidance on how to configure URL filtering Panorama if you are managing firewalls running different PAN-OS versions, refer to the [Panorama Administrator's Guide](#).


If you have valid licenses for both PAN-DB and BrightCloud, activating the PAN-DB license automatically deactivates the BrightCloud license (and vice versa).

- ▲ [Enable PAN-DB URL Filtering](#)
- ▲ [Enable BrightCloud URL Filtering](#)

Enable PAN-DB URL Filtering

Enable PAN-DB URL Filtering	
<p>Step 1 Obtain and install a PAN-DB URL filtering license and confirm that it is installed.</p> <p> If the license expires, PAN-DB URL filtering continues to work based on the URL category information that exists in the dataplane and management plane caches. However, URL cloud lookups and other cloud-based updates will not function until you install a valid license.</p>	<ol style="list-style-type: none"> 1. Select Device > Licenses and, in the License Management section, select the license installation method: <ul style="list-style-type: none"> • Activate feature using authorization code • Retrieve license keys from license server • Manually upload license key 2. After installing the license, confirm that the PAN-DB URL Filtering section, Date Expires field, displays a valid date.
<p>Step 2 Download the initial seed database and activate PAN-DB URL filtering.</p> <p> The firewall must have Internet access; you cannot manually upload the PAN-DB.</p>	<ol style="list-style-type: none"> 1. In the PAN-DB URL Filtering section, Download Status field, click Download Now. 2. Choose a region (North America, Europe, APAC, Japan) and then click OK to start the download. 3. After the download completes, click Activate. <p> If PAN-DB is already the active URL filtering vendor and you click Re-Download, this will reactivate PAN-DB by clearing the dataplane and management plane caches and replacing them with the contents of the new seed database. You should avoid doing this unless it is necessary, as you will lose your cache, which is customized based on the web traffic that has previously passed through the firewall based on user activity.</p>
<p>Step 3 Schedule the firewall to download dynamic updates for Applications and Threats.</p> <p> A Threat Prevention license is required to receive content updates, which covers Antivirus and Applications and Threats.</p>	<ol style="list-style-type: none"> 1. Select Device > Dynamic Updates. 2. In the Schedule field in the Applications and Threats section, click the None link to schedule periodic updates. <p> 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.</p> <p>The Applications and Threats updates might contain updates for URL filtering related to the Safe Search Enforcement option in the URL filtering profile (Objects > Security Profiles > URL Filtering). For example, if Palo Alto Networks adds support for a new search provider vendor or if the method used to detect the Safe Search setting for an existing vendor changes, the Application and Threats updates will include that update.</p>

Enable BrightCloud URL Filtering

Enable BrightCloud URL Filtering	
<p>Step 1 Obtain and install a BrightCloud URL filtering license and confirm that it is installed.</p> <p> 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.</p>	<ol style="list-style-type: none"> 1. Select Device > Licenses and, in the License Management section, select the license installation method: <ul style="list-style-type: none"> • 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.
<p>Step 2 Install the BrightCloud database.</p> <p>The way you do this depends on whether or not the firewall has direct Internet access.</p>	<p>Firewall with Direct Internet Access</p> <p>In the Device > Licenses page, BrightCloud URL Filtering section, Active field, click the Activate link to install the BrightCloud database. This operation automatically initiates a system reset.</p> <p>Firewall without Direct Internet Access</p> <ol style="list-style-type: none"> 1. Download the BrightCloud database to a host that has Internet access. The firewall must have access to the host: <ol style="list-style-type: none"> a. On a host with Internet access, go to the Palo Alto Support website (https://support.paloaltonetworks.com) and log in. b. In the Resources section, click Dynamic Updates. c. In the BrightCloud Database section, click Download and save the file to the host. 2. Upload the database to the firewall: <ol style="list-style-type: none"> a. Log in to the firewall, select Device > Dynamic Updates and click Upload. b. For the Type, select URL Filtering. c. Enter the path to the File on the host or click Browse to find it, then click OK. When the Status is <i>Completed</i>, click Close. 3. Install the database: <ol style="list-style-type: none"> a. In the Device > Dynamic Updates page, click Install From File. b. For the Type, select URL Filtering. The firewall automatically selects the file you just uploaded. c. Click OK and, when the Result is <i>Succeeded</i>, click Close.
<p>Step 3 Enable cloud lookups for dynamically categorizing a URL if the category is not available on the local BrightCloud database.</p>	<ol style="list-style-type: none"> 1. Access the firewall CLI. 2. Enter the following commands to enable dynamic URL filtering: <pre>configure set deviceconfig setting url dynamic-url yes commit</pre>

Enable BrightCloud URL Filtering (Continued)

Step 4 Schedule the firewall to download dynamic updates for Applications and Threats signatures and URL filtering.

You can only schedule dynamic updates if the firewall has direct Internet access.

The Applications and Threats updates might contain updates for URL filtering related to the **Safe Search Enforcement** option in the URL filtering profile. For example, if Palo Alto Networks adds support for a new search provider vendor or if the method used to detect the Safe Search setting for an existing vendor changes, the Application and Threats updates will include that update.

BrightCloud updates include a database of approximately 20 million websites that are stored on the firewall drive. You must schedule the URL filtering updates to receive database updates.



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 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.




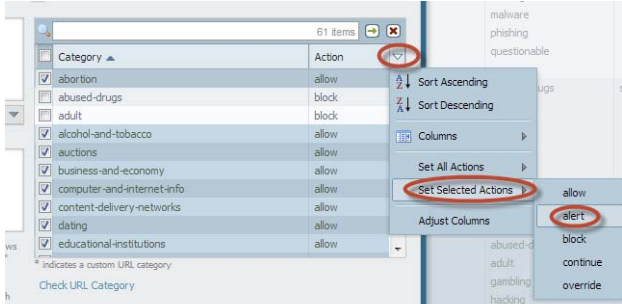
If updates are already scheduled in a section, the link text displays the schedule settings.

Determine URL Filtering Policy Requirements

The recommended practice for deploying URL filtering in your organization is to first start with a passive URL filtering profile that will alert on most categories. After setting the alert action, you can then monitor user web activity for a few days to determine patterns in web traffic. After doing so, you can then make decisions on the websites and website categories that should be controlled.

In the procedure that follows, threat-prone sites will be set to block and the other categories will be set to alert, which will cause all websites traffic to be logged. This may potentially create a large amount of log files, so it is best to do this for initial monitoring purposes to determine the types of websites your users are accessing. After determining the categories that your company approves of, those categories should then be set to allow, which will not generate logs. You can also reduce URL filtering logs by enabling the **Log container page only** option in the URL Filtering profile, so only the main page that matches the category will be logged, not subsequent pages/categories that may be loaded within the container page.

Configure and Apply a Passive URL Filtering Profile

<p>Step 1 Create a new URL Filtering profile.</p>	<ol style="list-style-type: none"> 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.
<p>Step 2 Configure the action for all categories to alert, except for threat-prone categories, which should remain blocked.</p> <p> To select all items in the category list from a Windows system, click the first category, then hold down the shift key and click the last category—this will select all categories. Hold the control key (ctrl) down and click items that should be deselected. On a Mac, do the same using the shift and command keys. You could also just set all categories to alert and manually change the recommended categories back to block.</p>	<ol style="list-style-type: none"> 1. In the section that lists all URL categories, select all categories. 2. To the right of the <i>Action</i> column heading, mouse over and select the down arrow and then select Set Selected Actions and choose alert.  <ol style="list-style-type: none"> 3. To ensure that you block access to threat-prone sites, select the following categories and then set the action to block: abused-drugs, adult, gambling, hacking, malware. phishing, questionable, weapons. 4. Click OK to save the profile.
<p>Step 3 Apply the URL Filtering profile to the security policy rule(s) that allows web traffic for users.</p>	<ol style="list-style-type: none"> 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.
<p>Step 4 Save the configuration.</p>	<p>Click Commit.</p>

Configure and Apply a Passive URL Filtering Profile (Continued)

<p>Step 5 View the URL filtering logs to determine all of the website categories that your users are accessing. In this example, some categories are set to block, so those categories will also appear in the logs.</p> <p>For information on viewing the logs and generating reports, see Monitor Web Activity.</p>	<p>Select Monitor > Logs > URL Filtering. A log entry will be created for any website that exists in the URL filtering database that is in a category that is set to any action other than allow.</p>
--	---

Monitor Web Activity

URL filtering logs and reports show all user web activity for URL categories that are set to **alert**, **block**, **continue**, or **override**. By monitoring the logs, you can gain a better understanding of the web activity of your user base to determine a web access policy.

The following topics describe how to monitor web activity:

- ▲ [Interpret the URL Filtering Logs](#)
- ▲ [Use the ACC to Monitor Web Activity](#)
- ▲ [View URL Filtering Reports](#)
- ▲ [Configure Custom URL Filtering Reports](#)

Interpret the URL Filtering Logs

The following bullet points show examples of the URL filtering logs (**Monitor > Logs > URL filtering**).

- **Alert log**—In this log, the category is shopping and the action is alert.

Receive Time	Category	URL	Source	Destination	Destin... Country	Application	Action	Decrypted	Source User
11/25 16:05:51	shopping	www.amazon.com/	192.168.2.10	72.21.215.232	US	web-browsing	alert	no	bsimpson

- **Block log**—In this log, the category alcohol-and-tobacco was set to block, so the action is block-url and the user will see a response page indicating that the website was blocked.

Receive Time	Category	URL	Source	Destination	Destin... Country	Application	Action	Decrypted	Source User
11/25 16:11:47	alcohol-and-tobacco	www.bevmo.com/	192.168.2.10	12.24.44.212	US	web-browsing	block-url	no	bsimpson

- **Alert log on encrypted website**—In this example, the category is social-networking and the application is facebook-base, which is required to access the Facebook website and other Facebook applications. Because facebook.com is always encrypted using SSL, the traffic was decrypted by the firewall, which allows the website to be recognized and controlled if needed.

Receive Time	Category	URL	Source	Destination	Destin... Country	Application	Action	Decrypted	Source User
11/25 16:13:44	social-networking	www.facebook.com/	192.168.2.10	69.171.237.20	US	facebook-base	alert	yes	bsimpson

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.

The screenshot shows a log table with columns: Destin... Country, Application, Action, and Decrypted. A dropdown menu is open over the 'Action' column, showing 'Columns' and 'Adjust Columns'. The 'Adjust Columns' menu is open, listing various attributes with checkboxes. The following table represents the visible items in the 'Adjust Columns' menu:

Attribute	Checked
Receive Time	Yes
Category	Yes
Content Type	No
URL	Yes
From Zone	No
To Zone	No
Source	Yes
Source Country	No
NAT Source IP	No
Destination	Yes
Destination User	No
Destination Country	Yes
NAT Dest IP	No
From Port	No
NAT Source Port	No
To Port	No
NAT Destination Port	No

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.

The screenshot shows the Palo Alto Networks log details interface. On the left, a list of logs is shown with a red circle around the 'Log Details' icon for the log received at 11/25 16:13:44. The main panel displays details for this log:

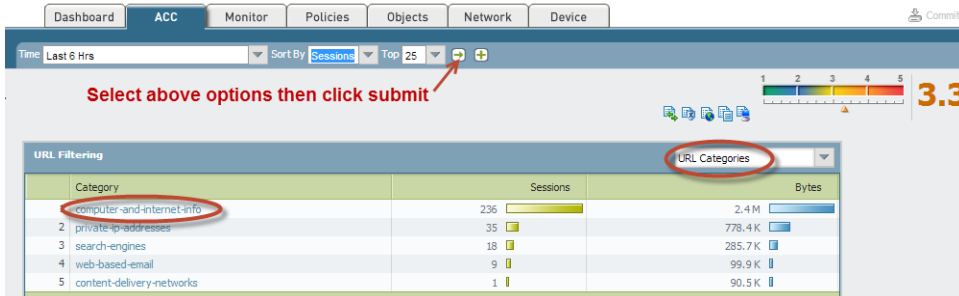
- General:** Session ID 35278, Action alert, Application facebook-base, Rule Rule1, Virtual System vsys1, Device SN, IP Protocol tcp, Log Action, Category social-networking, Generated Time, Receive Time 2013/11/25 16:13:44.
- Source:** User, Address 192.168.2.10, Country 192.168.0.0-192.168.255.255, Port 58779, Zone D-untrust, Interface vlan.1, NAT IP 10.43.14.64, NAT Port 32124.
- Destination:** User, Address 69.171.237.20, Country US, Port 443, Zone D-untrust, Interface ethernet1/1, NAT IP 69.171.237.20, NAT Port 443.
- URL Details:** Severity informational, Repeat Count 1, URL http://www.facebook.com/, and a link 'Request Categorization Change' circled in red.
- Flags:** Captive Portal, Proxy Transaction, Decrypted, Packet Capture, Client to Server, Server to Client.

At the bottom, a table shows related logs:

Receive Time	Log	Type	Application	Action	Rule	Bytes	Packets	Severity	Category	URL / Filename
11/25 16:13:42	traffic	start	sd	allow	Rule1	493	4			
11/25 16:13:43	traffic	start	social-networking	allow	Rule1	1,871	8			

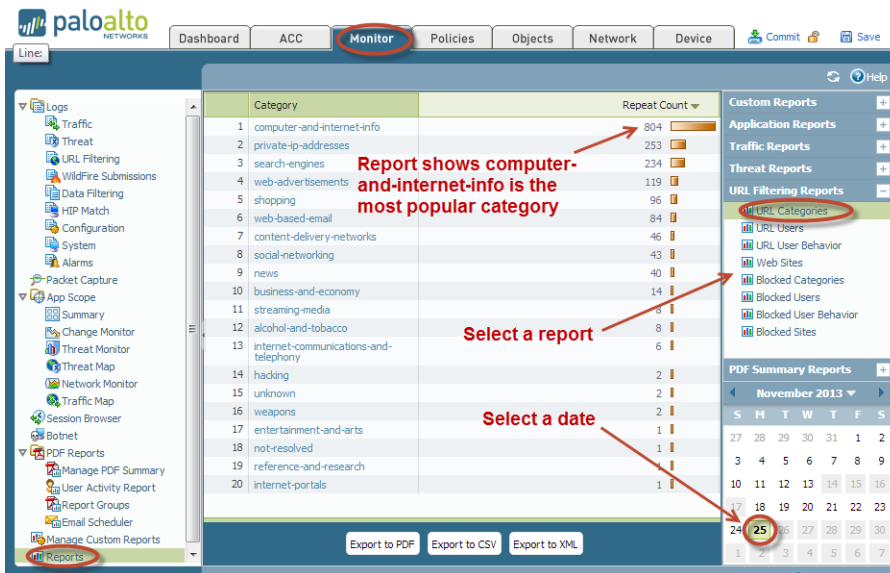
Use the ACC to Monitor Web Activity

For a quick view of the most common categories being accessed in your environment, select the **ACC** tab and scroll down to the **URL Filtering** section. Along the top of this window, you can also set the time range, sort by option, and define how many results will appear. Here you will see the most popular categories that are accessed by your users, sorted by the most popular at the top of the list. In this example, computer-and-internet-info is the most accessed category, followed by private-ip-addresses (internal servers), and search-engines. In the drop-down in the upper right of the statistics, you can also choose to list by URL Categories, Blocked URL Categories, and Blocked URLs.




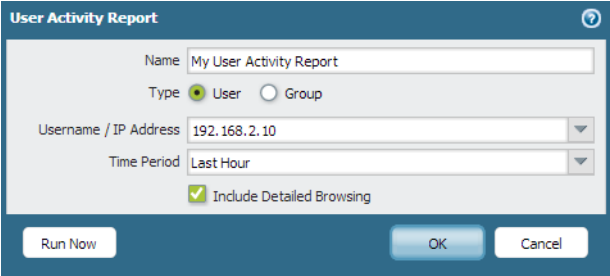
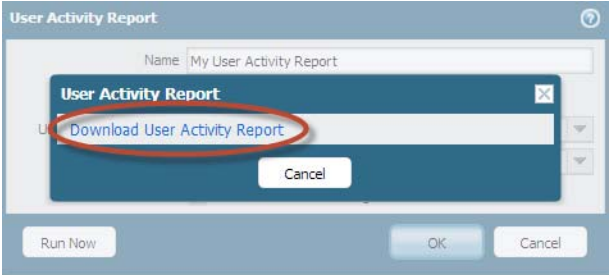
View URL Filtering Reports

To view the default URL filtering reports, select **Monitor > Reports** and under the **URL Filtering Reports** section, choose one of the reports. You can generate reports on URL Categories, URL users, Web Sites accessed, Blocked Categories, and more. The reports are based on a 24-hour period and the day is selected by choosing a day in the calendar section. You can also export the report to PDF, CSV, or XML.



View the User Activity Report

This report provides a quick method of viewing user or group activity and also provides an option to view browse time activity.

Generate a User Activity Report	
<p>Step 1 Configure a User Activity Report</p>	<ol style="list-style-type: none"> Select Monitor > PDF Reports > User Activity Report. Enter a report Name and select the report type. Select User to generate a report for one person, or select Group for a group of users. <div style="margin-top: 10px;">  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. </div> Enter the Username/IP address for a user report or enter the group name for a user group report. Select the time period. You can select an existing time period, or select Custom. Select the Include Detailed Browsing check box, so browsing information is included in the report. <div style="text-align: center; margin-top: 10px;">  </div>
<p>Step 2 Run the user activity report and then download the report.</p>	<ol style="list-style-type: none"> Click Run Now. After the report is generated, click the Download User Activity Report link. <div style="text-align: center; margin-top: 10px;">  </div> After the report is downloaded, click Cancel and then click OK to save the report.


Generate a User Activity Report (Continued)

Step 3 View the user activity report by opening the PDF file that was downloaded. The top of the report will contain a table of contents similar to the following:

User Activity Report for 192.168.2.10
Tuesday, December 31, 2013 09:35:47 - 10:35:46

Application Usage	2
Traffic Summary by URL Category	3
Browsing Summary by URL Category	4
Browsing Summary by Website	5
Blocked Browsing Summary by Website	6
Detailed Web Browsing Activity	7

Step 4 Click an item in the table of contents to view details. For example, click *Traffic Summary by URL Category* to view statistics for the selected user or group.

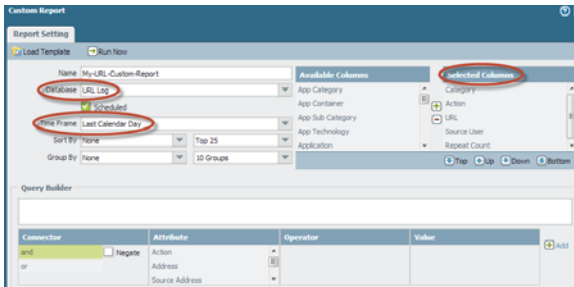


Traffic Summary by URL Category

Category	Count	Bytes
web-advertisements	49	426.8k
business-and-economy	36	693.6k
computer-and-internet-info	35	224.7k
news	24	499.4k
content-delivery-networks	20	6.2M
private-ip-addresses	11	206.2k
search-engines	6	62.2k
web-based-email	2	26.0k
social-networking	1	685


Configure Custom URL Filtering Reports



To generate a detailed report that can also be scheduled, you can configure a custom report and select from a list of all available URL filtering log fields.

Configure a Custom URL Filtering Report	
<p>Step 1 Add a new custom report.</p>	<ol style="list-style-type: none"> 1. Select Monitor > Manage Custom Reports and click Add. 2. Enter a report Name, for example, My-URL-Custom-Report. 3. From the Database drop-down, select URL Log.
<p>Step 2 Configure report options.</p>	<ol style="list-style-type: none"> 1. Select the Time Frame drop-down and select a range. 2. (Optional) To customize how the report is sorted and grouped, select Sort By and chose the number of items to display (top 25 for example) and then select Group By and select an option such as Category, and then select how many groups will be defined. 3. In the Available Columns list, select the fields to include the report. The following columns are typically used for a URL report: <ul style="list-style-type: none"> • Action • Category • Destination Country • Source User • URL <div style="text-align: center;">  </div>
<p>Step 3 Run the report to check the results. If the results are satisfactory, set a schedule to run the report automatically.</p>	<ol style="list-style-type: none"> 1. Click the Run Now icon to immediately generate the report that will appear in a new tab. 2. (Optional) Click the Schedule check box to run the report once per day. This will generate a daily report that details web activity over the last 24 hours. To access the report, select Monitor > Report and then expand Custom Reports on the right column and select the report.
<p>Step 4 Save the configuration.</p>	<p>Click Commit.</p>

Configure URL Filtering

After you [Determine URL Filtering Policy Requirements](#), you should have a basic understanding of what types of websites and website categories your users are accessing. With this information, you are now ready to create custom URL filtering profiles and attach them to the security policy rule(s) that allow web access.

Configure Website Controls	
<p>Step 1 Create a URL Filtering profile or select an existing one.</p> <p> Because the default URL filtering profile blocks risky and threat-prone content, it is a best practice to clone this profile rather than creating a new profile to preserve these default settings.</p>	<ol style="list-style-type: none"> 1. Select Objects > Security Profiles > URL Filtering. Select the default profile and then click Clone. The new profile will be named default-1. 2. Select the new profile and rename it.
<p>Step 2 Define how to control access to web content.</p>	<p>In the Categories tab, for each category that you want visibility into or control over, select a value from the Action column as follows:</p> <ul style="list-style-type: none"> • If you do not care about traffic to a particular category (that is you neither want to block it nor log it), select allow. • For visibility into traffic to sites in a category, select alert. • To deny access to traffic that matches the category and to enable logging of the blocked traffic, select block. • To require users to click Continue to proceed to a questionable site, select continue. • To only allow access if users provide a configured password, select override. For more details on this setting, see Configure URL Admin Override.
<p>Step 3 Define websites that should always be blocked or allowed.</p> <p>For example, to reduce URL filtering logs, you may want add you corporate websites in the allow list, so no logs will be generated for those sites. Or, if there is a website this is being overly used and is not work related in any way, you can add it to the block list.</p> <p>Items in the block list will always be blocked regardless of the action for the associated category, and URLs in the allow list will always be allowed.</p> <p>For more information on the proper format and wildcards usage, see Block and Allow Lists.</p>	<ol style="list-style-type: none"> 1. In the URL filtering profile, enter URLs or IP addresses in the Block List and select an action: <ul style="list-style-type: none"> • block—Block the URL. • continue—Prompt users click Continue to proceed to the web page. • override—The user will be a prompted for a password to continue to the website. • alert—Allow the user to access the website and add an alert log entry in the URL log. 2. For the Allow list, enter IP addresses or URLs that should always be allowed. Each row must be separated by a new line. 3. (Optional) Enable Safe Search Enforcement.

Configure Website Controls	
<p>Step 4 Modify the setting to log Container Pages only.</p>	<p>The Log container page only option is enabled by default so that only the main page that matches the category is logged, not subsequent pages/categories that may be loaded within the container page. To enable logging for all pages/categories, clear the Log container page only check box.</p>
<p>Step 5 Enable HTTP Header Logging for one or more of the supported HTTP header fields.</p>	<p>To log an HTTP header field, select one or more of the following fields to log:</p> <ul style="list-style-type: none"> • User-Agent • Referer • X-Forwarded-For
<p>Step 6 Save the URL filtering profile.</p>	<p>Click OK.</p>
<p>Step 7 (Optional) Enable response pages on the ingress interface (the interface that first receives traffic for your users).</p> <p>This option is required if you enable the continue or override action for any URL category.</p> <p> You can also Customize the URL Filtering Response Pages.</p>	<ol style="list-style-type: none"> 1. Select Network > Network Profiles > Interface Mgmt and either add a new profile or edit an existing profile. 2. Click the Response Pages check box to enable. 3. Click OK to save the profile. 4. Select Network > Interfaces and then edit the Layer 3 interface or VLAN interface that is your ingress interface. 5. On the Advanced tab, select the Interface Mgmt profile that has the response page option enabled. 6. Click OK to save the interface configuration.
<p>Step 8 Save the configuration.</p> <p> To test the URL filtering configuration, simply access a website in a category that is set to block or continue to see if the appropriate action is performed.</p>	<p>Click Commit.</p>

Log Client IP Addresses for Proxied HTTP/HTTPS Requests

If you have a proxy server deployed between the users on your network and the firewall, in HTTP/HTTPS requests the firewall might see the proxy server IP address as the source IP address in the traffic that the proxy forwards rather than the IP address of the client that requested the content. In many cases, the proxy server adds an X-Forwarded-For (XFF) header to traffic packets that includes the actual IPv4 or IPv6 address of the client that requested the content or from whom the request originated. In such cases, you can configure the firewall to read the XFF header values and determine the IP addresses of the client who requested the content. The firewall matches the XFF IP addresses with usernames that your policy rules reference so that those rules can control access for the associated users and groups. The firewall also uses the XFF-derived usernames to populate the source user fields of logs so you can monitor user access to web services.

You can also configure the firewall to add XFF values to URL Filtering logs. In these logs, an XFF value can be the client IP address, client username (if available), the IP address of the last proxy server traversed in a proxy chain, or any string of up to 128 characters that the XFF header stores.

XFF user identification applies only to HTTP or HTTPS traffic, and only if the proxy server supports the XFF header. If the header has an invalid IP address, the firewall uses that IP address as a username for group mapping references in policies. If the XFF header has multiple IP addresses, the firewall uses the first entry from the left.

▲ [Use XFF Values for Policies and Logging Source Users](#)

▲ [Add XFF Values to URL Filtering Logs](#)

Use XFF Values for Policies and Logging Source Users

You can configure the firewall to use XFF values in user-based policies and in the source user fields of logs. To use XFF values in policies, you must also [Map IP Addresses to Users](#), [Map Users to Groups](#) (if you have group-based policies), and [Enable User- and Group-Based Policy](#).



Logging XFF values doesn't populate the source IP address values of logs. When you view the logs, the source field displays the IP address of the proxy server if one is deployed between the user clients and the firewall. However, you can configure the firewall to [Add XFF Values to URL Filtering Logs](#) so that you can see user IP addresses in those logs.

To ensure that attackers can't read and exploit the XFF values in web request packets that exit the firewall to retrieve content from an external server, you can also configure the firewall to strip the XFF values from outgoing packets.

These options are not mutually exclusive: if you configure both, the firewall zeroes out XFF values only after using them in policies and logs.

Use XFF Values for Policies and Logging Source Users

<p>Step 1 Enable the firewall to use XFF values in policies and in the source user fields of logs.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Content-ID and edit the URL Filtering section. 2. Select x-forwarded-for.
---	---



Use XFF Values for Policies and Logging Source Users (Continued)	
<p>Step 2 Remove XFF values from outgoing web requests.</p>	<ol style="list-style-type: none"> 1. Select Strip x-forwarded-for. 2. Click OK and Commit.
<p>Step 3 Verify the firewall is populating the source user fields of logs.</p>	<ol style="list-style-type: none"> 1. Select a log type that has a source user field (for example, Monitor > Logs > Traffic). 2. Verify that the Source User column displays the usernames of users who access the web.

Add XFF Values to URL Filtering Logs

You can configure the firewall to add the XFF values from web requests to URL Filtering logs. The XFF values that the logs display can be client IP addresses, usernames if available, or any values of up to 128 characters that the XFF fields store.



This method of logging XFF values doesn't add usernames to the source user fields in URL Filtering logs. To populate the source user fields, see [Use XFF Values for Policies and Logging Source Users](#).

Add XFF Values to URL Filtering Logs	
<p>Step 1 Configure a URL Filtering profile.</p>	<ol style="list-style-type: none"> 1. Select Objects > Security Profiles > URL Filtering. 2. Select an existing profile or Add a new profile and enter a descriptive Name.  You can't enable XFF logging in the default URL Filtering profile. 3. In the Categories tab, Define how to control access to web content. 4. Select the Settings tab and select X-Forwarded-For. 5. Click OK to save the profile.
<p>Step 2 Attach the URL Filtering profile to a policy rule.</p>	<ol style="list-style-type: none"> 1. Select Policies > Security and click the rule. 2. Select the Actions tab, set the Profile Type to Profiles, and select the URL Filtering profile you just created. 3. Click OK and Commit.
<p>Step 3 Verify the firewall is logging XFF values.</p>	<ol style="list-style-type: none"> 1. Select Monitor > Logs > URL Filtering. 2. Display the XFF values in one of the following ways: <ul style="list-style-type: none"> • To display the XFF value for a single log—Click the  icon for the log to displays its details. The HTTP Headers section displays the X-Forwarded-For value. • To display the XFF values for all logs—Open the drop-down in any column header, select Columns, and select X-Forwarded-For. The page then displays an X-Forwarded-For column.

Customize the URL Filtering Response Pages

The firewall provides three predefined [URL Filtering Response Pages](#) that display by default when a user attempts to browse to a site in a category that is configured with one of the block actions in the [URL Filtering Profile](#) (block, continue, or override) or when [Safe Search Enforcement](#) blocks a search attempt. However, you can create your own custom response pages with your corporate branding, acceptable use policies, links to your internal resources as follows:

Customize the URL Filtering Response Pages	
<p>Step 1 Export the default response page(s).</p>	<ol style="list-style-type: none"> 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.
<p>Step 2 Edit the exported page.</p>	<ol style="list-style-type: none"> 1. Using the HTML text editor of your choice, edit the page: <ul style="list-style-type: none"> • If you want the response page to display custom information about the specific user, URL, or category that was blocked, add one or more of the supported URL Filtering Response Page Variables. • If you want to include custom images (such as your corporate logo), a sound, or style sheet, or link to another URL, for example to a document detailing your acceptable web use policy, include one or more of the supported Response Page References. 2. Save the edited page with a new filename. Make sure that the page retains its UTF-8 encoding. For example, in Notepad you would select UTF-8 from the Encoding drop-down in the Save As dialog.
<p>Step 3 Import the customized response page.</p>	<ol style="list-style-type: none"> 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.
<p>Step 4 Save the new response page(s).</p>	<p>Commit the changes.</p>
<p>Step 5 Verify that the new response page displays.</p>	<p>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.</p>

Configure URL Admin Override

In some cases there may be URL categories that you want to block, but allow certain individuals to browse to on occasion. In this case, you would set the category action to **override** and define a URL admin override password in the firewall Content-ID configuration. When users attempt to browse to the category, they will be required to provide the override password before they are allowed access to the site. Use the following procedure to configure URL admin override:

Configure URL Admin Override	
Step 1 Set the URL admin override password.	<ol style="list-style-type: none"> 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 a Server Certificate for the firewall to present to the user if the site with the override is an HTTPS site. 6. Select the Mode for prompting the user for the password: <ul style="list-style-type: none"> • Transparent—The firewall intercepts the browser traffic destined for site in a URL category you have set to override and impersonates the original destination URL, issuing an HTTP 401 to prompt for the password. Note that the client browser will display certificate errors if it does not trust the certificate. • Redirect—The firewall intercepts HTTP or HTTPS traffic to a URL category set to override and redirects the request to a Layer 3 interface on the firewall using an HTTP 302 redirect in order to prompt for the override password. If you select this option, you must provide the Address (IP address or DNS hostname) to which to redirect the traffic. 7. Click OK.
Step 2 (Optional) Set a custom override period.	<ol style="list-style-type: none"> 1. Edit the URL Filtering section. 2. To change the amount of time users can browse to a site in a category for which they have successfully entered the override password, enter a new value in the URL Admin Override Timeout field. By default, users can access sites within the category for 15 minutes without re-entering the password. 3. To change the amount of time users are blocked from accessing a site set to override after three failed attempts to enter the override password, enter a new value in the URL Admin Lockout Timeout field. By default, users are blocked for 30 minutes. 4. Click OK.

Configure URL Admin Override (Continued)	
<p>Step 3 (Redirect mode only) Create a Layer 3 interface to which to redirect web requests to sites in a category configured for override.</p>	<ol style="list-style-type: none"> 1. Create a management profile to enable the interface to display the URL Filtering Continue and Override Page response page: <ol style="list-style-type: none"> a. Select Network > Interface Mgmt and click Add. b. Enter a Name for the profile, select Response Pages, and then click OK. 2. Create the Layer 3 interface. Be sure to attach the management profile you just created (on the Advanced > Other Info tab of the Ethernet Interface dialog).
<p>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.</p>	<p>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:</p> <ol style="list-style-type: none"> 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).
<p>Step 5 Specify which URL categories require an override password to enable access.</p>	<ol style="list-style-type: none"> 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.
<p>Step 6 Apply the URL Filtering profile to the security policy rule(s) that allows access to the sites requiring password override for access.</p>	<ol style="list-style-type: none"> 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.
<p>Step 7 Save the configuration.</p>	<p>Click Commit.</p>

Enable Safe Search Enforcement

Many search engines have a safe search setting that filters out adult images and videos for search query return traffic. You can configure [Safe Search Enforcement](#) the Palo Alto Networks next-generation firewall to prevent search requests that do not have the strictest safe search settings enabled.



The [Safe Search Enforcement for Google and YouTube Searches using a Virtual IP Address](#) is not compatible with Safe Search Enforcement on the firewall.

There are two ways to enforce Safe Search on the firewall:

- ▲ [Block Search Results that are not Using Strict Safe Search Settings](#)
- ▲ [Enable Transparent Safe Search Enforcement](#)

Block Search Results that are not Using Strict Safe Search Settings

By default, when you enable safe search enforcement, when a user attempts to perform a search without using the strictest safe search settings, the firewall will block the search query results and display the URL Filtering Safe Search Block Page. This page provides a link to the search settings page for the corresponding search provider so that the end user can enable the safe search settings. If you plan to use this default method for enforcing safe search, you should communicate the policy to your end users prior to deploying the policy. See [Table: Search Provider Safe Search Settings](#) for details on how each search provider implements safe search. The default URL Filtering Safe Search Block Page provides a link to the search settings for the corresponding search provider. You can optionally [Customize the URL Filtering Response Pages](#).

Alternatively, to enable safe search enforcement so that it is transparent to your end users, configure the firewall to [Enable Transparent Safe Search Enforcement](#).

Enable Safe Search Enforcement	
<p>Step 1 Enable Safe Search Enforcement in the URL Filtering profile.</p>	<ol style="list-style-type: none"> 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: <ol style="list-style-type: none"> a. On the Categories tab, set the search-engines category to block. b. For each search engine that you want end users to be able to access, enter the web address in the Allow List text box. For example, to allow users access to Google and Bing searches only, you would enter the following: <pre>www.google.com</pre> <pre>www.bing.com</pre> 5. Configure other settings as necessary to: <ul style="list-style-type: none"> • Define how to control access to web content. • Define websites that should always be blocked or allowed. 6. Click OK to save the profile.
<p>Step 2 Add the URL Filtering profile to the security policy rule that allows traffic from clients in the trust zone to the Internet.</p>	<ol style="list-style-type: none"> 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.
<p>Step 3 Enable SSL Forward Proxy decryption.</p> <p>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.</p>	<ol style="list-style-type: none"> 1. Add a custom URL category for the search sites: <ol style="list-style-type: none"> a. Select Objects > Custom Objects > URL Category and Add a custom category. b. Enter a Name for the category, such as <code>SearchEngineDecryption</code>. c. Add the following to the Sites list: <pre>www.bing.*</pre> <pre>www.google.*</pre> <pre>search.yahoo.*</pre> d. Click OK to save the custom URL category object. 2. Follow the steps to Configure SSL Forward Proxy. 3. On the Service/URL Category tab in the Decryption policy rule, Add the custom URL category you just created and then click OK.

Enable Safe Search Enforcement (Continued)	
<p>Step 4 (Optional, but recommended) Block Bing search traffic running over SSL.</p> <p>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.</p>	<ol style="list-style-type: none"> 1. Add a custom URL category for Bing: <ol style="list-style-type: none"> a. Select Objects > Custom Objects > URL Category and Add a custom category. b. Enter a Name for the category, such as EnableBingSafeSearch. c. Add the following to the Sites list: www.bing.com/images/* www.bing.com/videos/* d. Click OK to save the custom URL category object. 2. Create another URL filtering profile to block the custom category you just created: <ol style="list-style-type: none"> a. Select Objects > Security Profiles > URL Filtering. b. Add a new profile and give it a descriptive Name. c. Locate the custom category in the Category list and set it to block. d. Click OK to save the URL filtering profile. 3. Add a security policy rule to block Bing SSL traffic: <ol style="list-style-type: none"> a. Select Policies > Security and Add a policy rule that allows traffic from your trust zone to the Internet. b. On the Actions tab, attach the URL filtering profile you just created to block the custom Bing category. c. On the Service/URL Category tab Add a New Service and give it a descriptive Name, such as bingssl. d. Select TCP as the Protocol and set the Destination Port to 443. e. Click OK to save the rule. f. Use the Move options to ensure that this rule is below the rule that has the URL filtering profile with safe search enforcement enabled.
<p>Step 5 Save the configuration.</p>	<p>Click Commit.</p>


Enable Safe Search Enforcement (Continued)

Step 6 Verify the Safe Search Enforcement configuration.



This verification step only works if you are using block pages to enforce safe search. If you are using transparent safe search enforcement, the firewall block page will invoke a URL rewrite with the safe search parameters in the query string.

1. From a computer that is behind the firewall, disable the strict search settings for one of the supported search providers. For example, on bing.com, click the **Preferences** icon on the Bing menu bar.



2. Set the **SafeSearch** option to **Moderate** or **Off** and click **Save**.
3. Perform a Bing search and verify that the URL Filtering Safe Search Block page displays instead of the search results:

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.

Enable Transparent Safe Search Enforcement

If you want to enforce filtering of search query results with the strictest safe search filters, but you don't want your end users to have to manually configure the settings, you can enable transparent [Safe Search Enforcement](#) as follows. This functionality is supported on Google, Yahoo, and Bing search engines only and requires Content Release version 475 or later.

Enable Transparent Safe Search Enforcement	
<p>Step 1 Make sure the firewall is running Content Release version 475 or later.</p>	<ol style="list-style-type: none"> 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.
<p>Step 1 Enable Safe Search Enforcement in the URL Filtering profile.</p>	<ol style="list-style-type: none"> 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: <ol style="list-style-type: none"> a. On the Categories tab, set the search-engines category to block. b. For each search engine that you want end users to be able to access, enter the web address in the Allow List text box. For example, to allow users access to Google and Bing searches only, you would enter the following: www.google.com www.bing.com 5. Configure other settings as necessary to: <ul style="list-style-type: none"> • Define how to control access to web content. • Define websites that should always be blocked or allowed. 6. Click OK to save the profile.
<p>Step 2 Add the URL Filtering profile to the security policy rule that allows traffic from clients in the trust zone to the Internet.</p>	<ol style="list-style-type: none"> 1. Select Policies > Security and select a rule to which to apply the URL filtering profile that you just enabled for Safe Search Enforcement. 2. On the Actions tab, select the URL Filtering profile. 3. Click OK to save the security policy rule.

Enable Transparent Safe Search Enforcement (Continued)

<p>Step 3 (Optional, but recommended) Block Bing search traffic running over SSL.</p> <p>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.</p>	<ol style="list-style-type: none"> 1. Add a custom URL category for Bing: <ol style="list-style-type: none"> a. Select Objects > Custom Objects > URL Category and Add a custom category. b. Enter a Name for the category, such as EnableBingSafeSearch. c. Add the following to the Sites list: <pre>www.bing.com/images/* www.bing.com/videos/*</pre> d. Click OK to save the custom URL category object. 2. Create another URL filtering profile to block the custom category you just created: <ol style="list-style-type: none"> a. Select Objects > Security Profiles > URL Filtering. b. Add a new profile and give it a descriptive Name. c. Locate the custom category you just created in the Category list and set it to block. d. Click OK to save the URL filtering profile. 3. Add a security policy rule to block Bing SSL traffic: <ol style="list-style-type: none"> a. Select Policies > Security and Add a policy rule that allows traffic from your trust zone to the Internet. b. On the Actions tab, attach the URL filtering profile you just created to block the custom Bing category. c. On the Service/URL Category tab Add a New Service and give it a descriptive Name, such as bingssl. d. Select TCP as the Protocol, set the Destination Port to 443. e. Click OK to save the rule. f. Use the Move options to ensure that this rule is below the rule that has the URL filtering profile with safe search enforcement enabled.
<p>Step 4 Edit the URL Filtering Safe Search Block Page, replacing the existing code with the Javascript for rewriting search query URLs to enforce safe search transparently.</p>	<ol style="list-style-type: none"> 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 this script. <p> Copy the script and paste it into the HTML editor, replacing the entire block page.</p>

Enable Transparent Safe Search Enforcement (Continued)	
<p>Step 5 Import the edited URL Filtering Safe Search Block page onto the firewall.</p>	<ol style="list-style-type: none"> 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.
<p>Step 6 Enable SSL Forward Proxy decryption.</p> <p>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.</p>	<ol style="list-style-type: none"> 1. Add a custom URL category for the search sites: <ol style="list-style-type: none"> a. Select Objects > Custom Objects > URL Category and Add a custom category. b. Enter a Name for the category, such as <code>SearchEngineDecryption</code>. c. Add the following to the Sites list: <ul style="list-style-type: none"> <code>www.bing.*</code> <code>www.google.*</code> <code>search.yahoo.*</code> d. Click OK to save the custom URL category object. 2. Follow the steps to Configure SSL Forward Proxy. 3. On the Service/URL Category tab in the Decryption policy rule, Add the custom URL category you just created and then click OK.
<p>Step 7 Save the configuration.</p>	<p>Click Commit.</p>

URL Filtering Use Case Examples

The following use cases show how to use App-ID to control a specific set of web-based applications and how to use URL categories as match criteria in a policy. When working with App-ID, it is important to understand that each App-ID signature may have dependencies that are required to fully control an application. For example, with Facebook applications, the App-ID facebook-base is required to access the Facebook website and to control other Facebook applications. For example, to configure the firewall to control Facebook email, you would have to allow the App-IDs facebook-base and facebook-mail. As another example, if you search [Applopedia](#) (the App-ID database) for LinkedIn, you will see that in order to control LinkedIn mail, you need to apply the same action to both App-IDs: linkedin-base and linkedin-mail. To determine application dependencies for App-ID signatures, visit [Applopedia](#), search for the given application, and then click the application for details.



The **User-ID** feature is required to implement policies based on users and groups and a **Decryption** policy is required to identify and control websites that are encrypted using SSL.

This section includes two uses cases:

- ▲ [Use Case: Control Web Access](#)
- ▲ [Use Case: Use URL Categories in Policy](#)

Use Case: Control Web Access

When using URL filtering to control user website access, there may be instances where granular control is required for a given website. In this use case, a URL filtering policy is applied to the security policy that allows web access for your users and the *social-networking* URL category is set to block, but the allow list in the URL profile is configured to allow the social networking site Facebook. To further control Facebook, the company policy also states that only marketing has full access to Facebook and all other users within the company can only read Facebook posts and cannot use any other Facebook applications, such as email, posting, chat, and file sharing. To accomplish this requirement, App-ID must be used to provide granular control over Facebook.

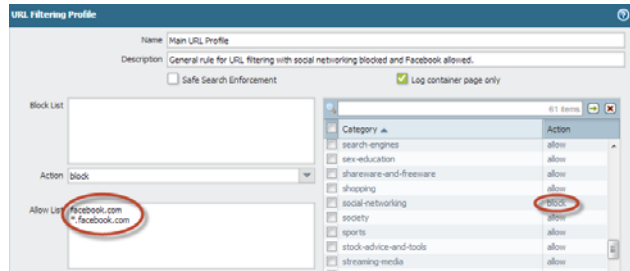
The first security rule will allow marketing to access the Facebook website as well as all Facebook applications. Because this allow rule will also allow access to the Internet, threat prevention profiles are applied to the rule, so traffic that matches the policy will be scanned for threats. This is important because the allow rule is terminal and will not continue to check other rules if there is a traffic match.

Control Web Access	
<p>Step 1 Confirm that URL filtering is licensed.</p>	<ol style="list-style-type: none"> 1. Select Device > Licenses and confirm that a valid date appears for the URL filtering database that will be used. This will either be PAN-DB or BrightCloud. 2. If a valid license is not installed, see Enable PAN-DB URL Filtering.
<p>Step 2 Confirm that User-ID is working. User-ID is required to create policies based on users and groups.</p>	<ol style="list-style-type: none"> 1. To check group mapping, from the CLI, enter the following command: show user group-mapping statistics 2. To check user mapping, from the CLI, enter the following command: show user ip-user-mapping-mp all 3. If statistics do not appear and/or IP to user mapping information is not displayed, see User-ID.
<p>Step 3 Set up a URL filtering profile by cloning the default profile.</p>	<ol style="list-style-type: none"> 1. Select Objects > Security Profiles > URL Filtering and select the default profile. 2. Click the Clone icon. A new profile should appear named default-1. 3. Select the new profile and rename it.

Control Web Access (Continued)

Step 4 Configure the URL filtering profile to block social-networking and allow Facebook.

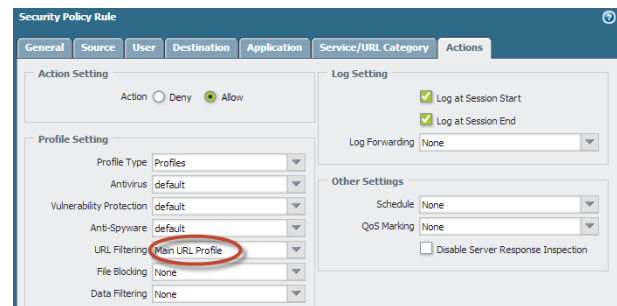
1. Modify the new URL filtering profile and in the **Category** list scroll to **social-networking** and in the **Action** column click on **allow** and change the action to **block**.
2. In the **Allow List** box, type `facebook.com`, press enter to start a new line and then type `*.facebook.com`. Both of these formats are required, so all URL variants a user may use will be identified, such as `facebook.com`, `www.facebook.com`, and `https://facebook.com`.



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

Step 5 Apply the new URL filtering profile to the security policy rule that allows web access from the user network to the Internet.

1. Select **Policies > Security** and click on the policy rule that allows web access.
2. On the **Actions** tab, select the URL profile you just created from the **URL Filtering** drop-down.



3. Click **OK** to save.

Control Web Access (Continued)

Step 6 Create the security policy that will allow marketing access the Facebook website and all Facebook applications.

This rule must precede other rules because it is more specific than the other policies and because it is an allow rule, which will terminate when a traffic match occurs.

1. Select **Policies > Security** and click **Add**.
2. Enter a **Name** and optionally a **Description** and **Tag(s)**.
3. On the **Source** tab add the zone where the users are connected.
4. On the **User** tab in the **Source User** section click **Add**.
5. Select the directory group that contains your *marketing* users.
6. On the **Destination** tab, select the zone that is connected to the Internet.
7. On the **Applications** tab, click **Add** and add the *facebook* App-ID signature.
8. On the **Actions** tab, add the default profiles for **Antivirus**, **Vulnerability Protection**, and **Anti-Spyware**.



9. Click **OK** to save the security profile.

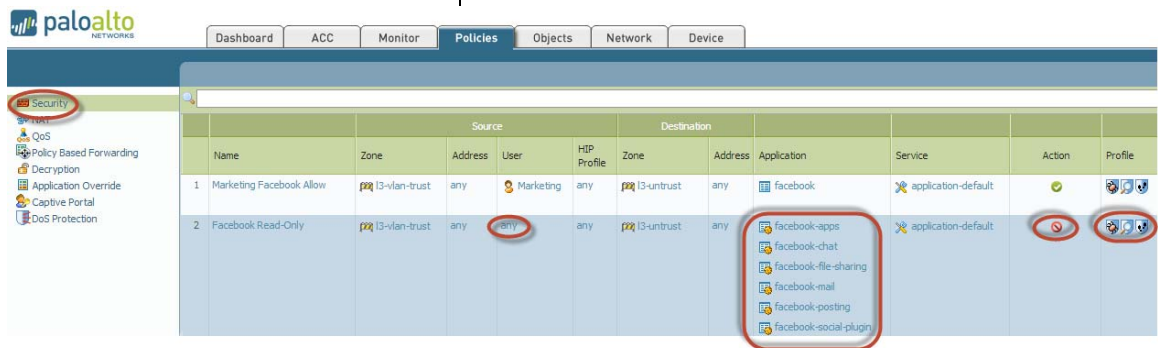
The *facebook* App-ID signature used in this policy encompasses all Facebook applications, such as facebook-base, facebook-chat, and facebook-mail, so this is the only App-ID signature required in this rule.

With this policy in place, when a marketing employee attempts to access the Facebook website or any Facebook application, the rule matches based on the user being part of the marketing group. For traffic from any user outside of marketing, the rule will be skipped because there would not be a traffic match and rule processing would continue.

Control Web Access (Continued)

Step 7 Configure the security policy to block all other users from using any Facebook applications other than simple web browsing. The easiest way to do this is to clone the marketing allow policy and then modify it.

1. From **Policies > Security** click the marketing Facebook allow policy you created earlier to highlight it and then click the **Clone** icon.
2. Enter a **Name** and optionally enter a **Description** and **Tag**(s).
3. On the **User** tab highlight the marketing group and delete it and in the drop-down select **any**.
4. On the **Applications** tab, click the *facebook* App-ID signature and delete it.
5. Click **Add** and add the following App-ID signatures:
 - facebook-apps
 - facebook-chat
 - facebook-file-sharing
 - facebook-mail
 - facebook-posting
 - facebook-social-plugin
6. On the **Actions** tab in the **Action Setting** section, select **Deny**. The profile settings should already be correct because this rule was cloned.



7. Click **OK** to save the security profile.
8. Ensure that this new deny rule is listed after the marketing allow rule, to ensure that rule processing occurs in the correct order to allow marketing users and then to deny/limit all other users.
9. Click **Commit** to save the configuration.

With these policies in place, any user who is part of the marketing group will have full access to all Facebook applications and any user that is not part of the marketing group will only have read-only access to the Facebook website and will not be able to use Facebook functions such as post, chat, email, and file sharing.

Use Case: Use URL Categories in Policy

URL categories can also be used as match criteria in the following policy types: captive portal, decryption, security, and QoS. In this use case, URL categories will be used in decryption policies to control which web categories should be decrypted or not decrypted. The first rule is a no-decrypt rule that will not decrypt user traffic if the website category is *financial-services* or *health-and-medicine* and the second rule will decrypt all other traffic. The decryption policy type is *ssl-forward-proxy*, which is used for controlling decryption for all outbound connections performed by users.

Configure a Decryption Policy Based on URL Category

Step 1 Create the no-decrypt rule that will be listed first in the decryption policies list. This will prevent any website that is in the *financial-services* or *health-and-medicine* URL categories from being decrypted.

1. Select **Policies > Decryption** and click **Add**.
2. Enter a **Name** and optionally enter a **Description** and **Tag(s)**.
3. On the **Source** tab, add the zone where the users are connected.
4. On the **Destination** tab, enter the zone that is connected to the Internet.
5. On the **URL Category** tab, click **Add** and select the *financial-services* and *health-and-medicine* URL categories.
6. On the **Options** tab, set the action to **No Decrypt** and the **Type** to **SSL Forward Proxy**.



7. Click **OK** to save the policy.

Configure a Decryption Policy Based on URL Category (Continued)

Step 2 Create the decryption policy that will decrypt all other traffic. This policy will be listed after the no-decrypt policy.

1. Select the no-decrypt policy you created previously and then click **Clone**.
2. Enter a **Name** and optionally enter a **Description** and **Tag(s)**.
3. On the **URL Category** tab, select *financial-services* and *health-and-medicine* and then click the **Delete** icon.
4. On the **Options** tab, set the action to **Decrypt** and the **Type** to **SSL Forward Proxy**.



5. Ensure that this new decryption rule is listed after the no-decrypt rule as shown in the previous screen capture. This will ensure that rule processing occurs in the correct order, so websites in the *financial-services* and *health-and-medicine* are not decrypted
6. Click **OK** to save the policy.

Step 3 (BrightCloud only) Enable cloud lookups for dynamically categorizing a URL when the category is not available on the local database on the firewall.

1. Access the CLI on the firewall.
2. Enter the following commands to enable Dynamic URL Filtering:
 - a. **configure**
 - b. **set deviceconfig setting url dynamic-url yes**
 - c. **commit**

Step 4 Save the configuration.

Click **Commit**.

With these two decrypt policies in place, any traffic destined for the *financial-services* or *health-and-medicine* URL categories will not be decrypted. All other traffic will be decrypted.

You can also define more granular control over decryption policies by defining decryption profiles, which are used to perform checks such as server certificate checks or blocking sessions with expired certificates. The profile is then added in the **Options** tab of the decryption policy. For a complete list of checks that can be performed, select **Objects > Decryption Profiles** from the firewall and then click the help icon.

Now that you have a basic understanding of the powerful features of URL filtering, App-ID, and User-ID, you can apply similar policies to your firewall to control any application in the Palo Alto Networks App-ID signature database and control any website contained in the URL filtering database.

For help in troubleshooting URL filtering issues, see [Troubleshoot URL Filtering](#).

Troubleshoot URL Filtering

The following topics provide troubleshooting guidelines for diagnosing and resolving common URL filtering problems.

- ▲ [Problems Activating PAN-DB](#)
- ▲ [PAN-DB Cloud Connectivity Issues](#)
- ▲ [URLs Classified as Not-Resolved](#)
- ▲ [Incorrect Categorization](#)
- ▲ [URL Database Out of Date](#)

Problems Activating PAN-DB

The following table describes procedures that you can use to resolve issues with activating PAN-DB.

Troubleshoot PAN-DB Activation Issues	
1.	Access the CLI on the firewall.
2.	Verify whether PAN-DB has been activated by running the following command: <pre>admin@PA-200> show system setting url-database</pre> If the response is <code>paloaltonetworks</code> , then PAN-DB is the active vendor.
3.	Verify that the firewall has a valid PAN-DB license by running the following command: <pre>admin@PA-200> request license info</pre> You should see the license entry <code>Feature: PAN_DB URL Filtering</code> . If the license is not installed, you will need to obtain and install a license. See Configure URL Filtering .
4.	After the license is installed, download a new PAN-DB seed database by running the following command: <pre>admin@PA-200> request url-filtering download paloaltonetworks region <region></pre>
5.	Check the download status by running the following command: <pre>admin@PA-200> request url-filtering download status vendor paloaltonetworks</pre> <ul style="list-style-type: none"> • If the message is different from <code>PAN-DB download: Finished successfully</code>, stop here; there may be a problem connecting to the cloud. Attempt to solve the connectivity issue by performing basic network troubleshooting between the firewall and the Internet. For more information, see PAN-DB Cloud Connectivity Issues. • If the message is <code>PAN-DB download: Finished successfully</code>, the firewall successfully downloaded the URL seed database. Try to enable PAN-DB again by running the following command: <pre>admin@PA-200> set system setting url-database paloaltonetworks</pre>
6.	If the problems persists, contact Palo Alto Networks support.

PAN-DB Cloud Connectivity Issues

To check cloud connectivity, run the following command:

```
admin@pa-200> show url-cloud status
```

If the cloud is accessible, the expected response is similar to the following:

```
admin@PA-200> show url-cloud status
PAN-DB URL Filtering
License :                               valid
Current cloud server :                   s0000.urlcloud.paloaltonetworks.com
Cloud connection :                        connected
URL database version - device :           2013.11.18.000
URL database version - cloud :            2013.11.18.000 ( last update time
2013/11/19
13:20:51 )
URL database status :                     good
URL protocol version - device :           pan/0.0.2
URL protocol version - cloud :           pan/0.0.2
Protocol compatibility status :           compatible
```

If the cloud is not accessible, the expected response is similar to the following:

```
admin@PA-200> show url-cloud status
PAN-DB URL Filtering
License :                               valid
Cloud connection :                        not connected
URL database version - device :           2013.11.18.000
URL database version - cloud :            2013.11.18.000 ( last update time
2013/11/19
13:20:51 )
URL database status :                     good
URL protocol version - device :           pan/0.0.2
URL protocol version - cloud :           pan/0.0.2
Protocol compatibility status :           compatible
```

The following table describes procedures that you can use to resolve issues based on the output of the `show url-cloud status` command, how to ping the URL cloud servers, and what to check if the firewall is in a High Availability (HA) configuration.

Troubleshoot Cloud Connectivity Issues

- PAN-DB URL Filtering license field shows invalid—Obtain and install a valid PAN-DB license.
- URL database status is out of date—Download a new seed database by running the following command:

```
admin@pa-200> request url-filtering download paloaltonetworks region <region>
```
- URL protocol version shows not compatible—Upgrade PAN-OS to the latest version.
- Attempt to ping the PAN-DB cloud server from the firewall by running the following command:

```
admin@pa-200> ping source <ip-address> host s0000.urlcloud.paloaltonetworks.com
```

For example, if your management interface IP address is 10.1.1.5, run the following command:

```
admin@pa-200> ping source 10.1.1.5 host s0000.urlcloud.paloaltonetworks.com
```
- If the firewall is in an HA configuration, verify that the HA state of the devices supports connectivity to the cloud systems. You can determine the HA state by running the following command:

```
admin@pa-200> show high-availability state
```

Connection to the cloud will be blocked if the firewall is not in one of the following states:
 - active
 - active-primary
 - active-secondary

If the problem persists, contact Palo Alto Networks support.

URLs Classified as Not-Resolved

The following table describes procedures you can use to resolve issues where some or all of the URLs being identified by PAN-DB are classified as *Not-resolved*:

Troubleshoot URLs Classified as Not-Resolved

1. Check the PAN-DB cloud connection by running the following command:

```
admin@PA-200> show url-cloud status
```

The Cloud connection: field should show *connected*. If you see anything other than *connected*, any URL that do not exist in the management plane cache will be categorized as *not-resolved*. To resolve this issue, see [PAN-DB Cloud Connectivity Issues](#).
2. If the cloud connection status shows *connected*, check the current utilization of the firewall. If the firewall's performance is spiking, URL requests may be dropped (may not reach the management plane), and will be categorized as *not-resolved*.
To view system resources, run the following command and view the *%CPU* and *%MEM* columns:

```
admin@PA-200> show system resources
```

You can also view system resources from the firewall's web interfaces by clicking the **Dashboard** tab and viewing the **System Resources** section.
3. If the problem persist, contact Palo Alto Networks support.

Incorrect Categorization

The following steps describe the procedures you can use if you identify a URL that does not have the correct categorization. For example, if the URL *paloaltonetworks.com* was categorized as *alcohol-and-tobacco*, the categorization is not correct; the category should be *computer-and-internet-info*.

Troubleshoot Incorrect Categorization Issues

1. Verify the category in the dataplane by running the following command:

```
admin@PA-200> show running url <URL>
```

For example, to view the category for the Palo Alto Networks website, run the following command:

```
admin@PA-200> show running url paloaltonetworks.com
```

If the URL stored in the dataplane cache has the correct category (*computer-and-internet-info* in this example), then the categorization is correct and no further action is required. If the category is not correct, continue to the next step.

2. Verify if the category in the management plane by running the command:

```
admin@PA-200> test url-info-host <URL>
```

For example:

```
admin@PA-200> test url-info-host paloaltonetworks.com
```

If the URL stored in the management plane cache has the correct category, remove the URL from the dataplane cache by running the following command:

```
admin@PA-200> clear url-cache url <URL>
```

The next time the device requests the category for this URL, the request will be forwarded to the management plane. This will resolve the issue and no further action is required. If this does not solve the issue, go to the next step to check the URL category on the cloud systems.

3. Verify the category in the cloud by running the following command:

```
admin@PA-200> test url-info-cloud <URL>
```

4. If the URL stored in the cloud has the correct category, remove the URL from the dataplane and the management plane caches.

Run the following command to delete a URL from the dataplane cache:

```
admin@PA-200> clear url-cache url <URL>
```

Run the following command to delete a URL from the management plane cache:

```
admin@PA-200> delete url-database url <URL>
```

The next time the device queries for the category of the given URL, the request will be forwarded to the management plane and then to the cloud. This should resolve the category lookup issue. If problems persist, see the next step to submit a categorization change request.

5. To submit a change request from the web interface, go to the URL log and select the log entry for the URL you would like to have changed.
-

Troubleshoot Incorrect Categorization Issues

6. Click the **Request Categorization** change link and follow instructions. You can also request a category change from the Palo Alto Networks [Test A Site](#) website by searching for the URL and then clicking the **Request Change** icon. To view a list of all available categories with descriptions of each category, refer to <https://urlfiltering.paloaltonetworks.com/CategoryList.aspx>.

If your change request is approved, you will receive an email notification. You then have two options to ensure that the URL category is updated on the firewall:

- Wait until the URL in the cache expires and the next time the URL is accessed by a user, the new categorization update will be put in the cache.
- Run the following command to force an update in the cache:

```
admin@PA-200> request url-filtering update url <URL>
```

URL Database Out of Date

If you have observed through the syslog or the CLI that PAN-DB is out-of-date, it means that the connection from the firewall to the URL Cloud is blocked. This usually occurs when the URL database on the firewall is too old (version difference is more than three months) and the cloud cannot update the firewall automatically. In order to resolve this issue, you will need to re-download an initial seed database from the cloud (this operation is not blocked). This will result in an automatic re-activation of PAN-DB.

To manually update the database, perform one of the following steps:

- From the web interface, select **Device > Licenses** and in the **PAN-DB URL Filtering** section click the **Re-Download** link.
- From the CLI, run the following command:

```
admin@PA-200> request url-filtering download paloaltonetworks region <region_name>
```



When the seed database is re-downloaded, the URL cache in the management plane and dataplane will be purged. The management plane cache will then be re-populated with the contents of the new seed database.

Quality of Service

Quality of Service (QoS) is a set of technologies that work on a network to guarantee its ability to dependably run high-priority applications and traffic under limited network capacity. QoS technologies accomplish this by providing differentiated handling and capacity allocation to specific flows in network traffic. This enables the network administrator to assign the order in which traffic is handled, and the amount of bandwidth afforded to traffic.

Palo Alto Networks Application Quality of Service (QoS) provides basic QoS applied to networks and extends it to provide QoS to applications and users.

Use the following topics to learn about and configure Palo Alto Networks Application QoS:

- ▲ [QoS Overview](#)
- ▲ [QoS Concepts](#)
- ▲ [Configure QoS](#)
- ▲ [Configure QoS for a Virtual System](#)
- ▲ [QoS Use Case Examples](#)



- You can use the Palo Alto Networks [product comparison tool](#) to view the QoS features supported on your firewall platform. Select two or more product platforms and click **Compare Now** to view QoS feature support for each platform (for example, you can check if your firewall platform supports QoS on subinterfaces and if so, the maximum number of subinterfaces that QoS can be enabled on).
- The PA-7050 firewall supports QoS on Aggregate Ethernet (AE) interfaces as of PAN-OS 6.0.0.

QoS Overview

Use QoS to prioritize and adjust quality aspects of network traffic. You can assign the order in which packets are handled and allot bandwidth, ensuring preferred treatment and optimal levels of performance are afforded to selected traffic, applications, and users.


Service quality measurements subject to a QoS implementation are bandwidth (maximum rate of transfer), throughput (actual rate of transfer), latency (delay), and jitter (variance in latency). The capability to shape and control these service quality measurements makes QoS of particular importance to high-bandwidth, real-time traffic such as voice over IP (VoIP), video conferencing, and video-on-demand that has a high sensitivity to latency and jitter. Additionally, use QoS to achieve outcomes such as the following:

- Prioritize network and application traffic, guaranteeing high priority to important traffic or limiting non-essential traffic.
- Achieve equal bandwidth sharing among different subnets, classes, or users in a network.
- Allocate bandwidth externally or internally or both, applying QoS to both upload and download traffic or to only upload or download traffic.
- Ensure low latency for customer and revenue-generating traffic in an enterprise environment.
- Perform traffic profiling of applications to ensure bandwidth usage.

QoS implementation on a Palo Alto Networks firewall begins with three primary configuration components that support a full QoS solution: a [QoS Profile](#), a [QoS Policy](#), and setting up the [QoS Egress Interface](#). Each of these options in the QoS configuration task facilitate a broader process that optimizes and prioritizes the traffic flow and allocates and ensures bandwidth according to configurable parameters.

The figure [Figure: QoS Traffic Flow](#) shows traffic as it flows from the source, is shaped by the firewall with QoS enabled, and is ultimately prioritized and delivered to its destination.

Figure: QoS Traffic Flow



The QoS configuration options allow you to control the traffic flow and define it at different points in the flow. The [Figure: QoS Traffic Flow](#) indicates where the configurable options define the traffic flow. Use the QoS Profile to define QoS classes and use the QoS Policy to associate QoS classes with selected traffic. Enable the QoS Profile on an interface to shape traffic according to the QoS configuration as it flows through the network.

You can configure a QoS Profile and QoS Policy individually or in any order, according to your preference. Each of the QoS configuration options has components that influence the definition of the other options and the QoS configuration options can be used to create a full and granular QoS policy or can be used sparingly with minimal administrator action.

Each firewall model supports a maximum number of ports that can be configured with QoS. Refer to the spec sheet for your [firewall model](#) or use the [product comparison tool](#) to view QoS feature support for two or more firewalls on a single page.

QoS Concepts

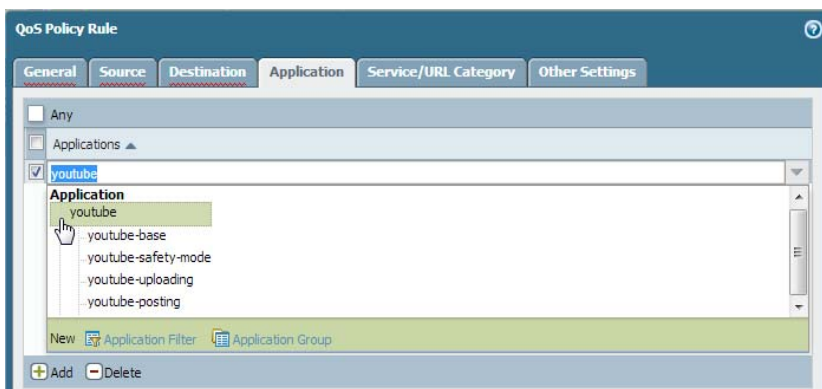
Use the following topics to learn about the different components and mechanisms of a QoS configuration on a Palo Alto Networks firewall:

- ▲ [QoS for Applications and Users](#)
- ▲ [QoS Profile](#)
- ▲ [QoS Classes](#)
- ▲ [QoS Policy](#)
- ▲ [QoS Egress Interface](#)
- ▲ [QoS Cleartext and Tunneled Traffic](#)


QoS for Applications and Users

A Palo Alto Networks firewall provides basic QoS, controlling traffic leaving the firewall according to network or subnet, and extends the power of QoS to also classify and shape traffic according to application and user. The Palo Alto Networks firewall provides this capability by integrating the features App-ID and User-ID with the QoS configuration. App-ID and User-ID entries that exist to identify specific applications and users in your network are available in the QoS configuration so that you can easily specify applications and users to apply QoS to.

You can use a QoS Policy in the web interface (**Policies > QoS**) to apply QoS specifically to an application's traffic:



Or to a user's traffic:



See [App-ID](#) and [User-ID](#) for more information on these features.

QoS Profile

Use a QoS profile to define values of up to eight QoS classes contained within that single profile (**Network > Network Profiles > QoS Profile**):

Name	Guaranteed Egress (Mbps)	Maximum Egress (Mbps)	Priority
restrict downloads to 15 Mbps			
class1			
class2			
class3			
class4			
class5			
class6			
class7			
class8			

Class	Priority	Egress Max	Egress Guaranteed
class1	real-time	0	0
class2	high	0	0
class3	high	0	0
class4	medium	0	0
class5	medium	0	0
class6	low	0	0
class7	low	0	0
class8	low	0	0

You enable QoS by applying a QoS profile to the egress interface for network, application or user traffic (or specifically, for cleartext or tunneled traffic). An interface configured with QoS shapes traffic according to the QoS profile class definitions and the traffic associated with those classes in the QoS policy.

A default QoS Profile is available on the firewall. The default profile and the classes defined in the profile do not have predefined maximum or guaranteed bandwidth limits.

You can set bandwidth limits for a QoS profile and/or set limits for individual QoS classes within the QoS profile. The total guaranteed bandwidth limits of all eight QoS classes in a QoS Profile cannot exceed the total bandwidth allocated to that QoS Profile. Enabling QoS on a physical interface includes setting the maximum bandwidth for traffic leaving the firewall through this interface. A QoS profile's guaranteed bandwidth (the **Egress Guaranteed** field) should not exceed the bandwidth allocated to the physical interface that QoS is enabled on.

For details, see [Create a QoS profile](#).

QoS Classes

A QoS class determines the priority and bandwidth for traffic it is assigned to. In the web interface, use the QoS profile to define QoS classes (**Network > Network Profiles > QoS Profile**):

Class	Priority	Egress Max	Egress Guaranteed
<input checked="" type="checkbox"/> class1	medium	0	0
<input type="checkbox"/> class2	high	0	0
<input type="checkbox"/> class3	high	0	0

Defining a QoS class includes setting the class's Priority, maximum bandwidth (Egress Max), and guaranteed bandwidth (Egress Guaranteed).

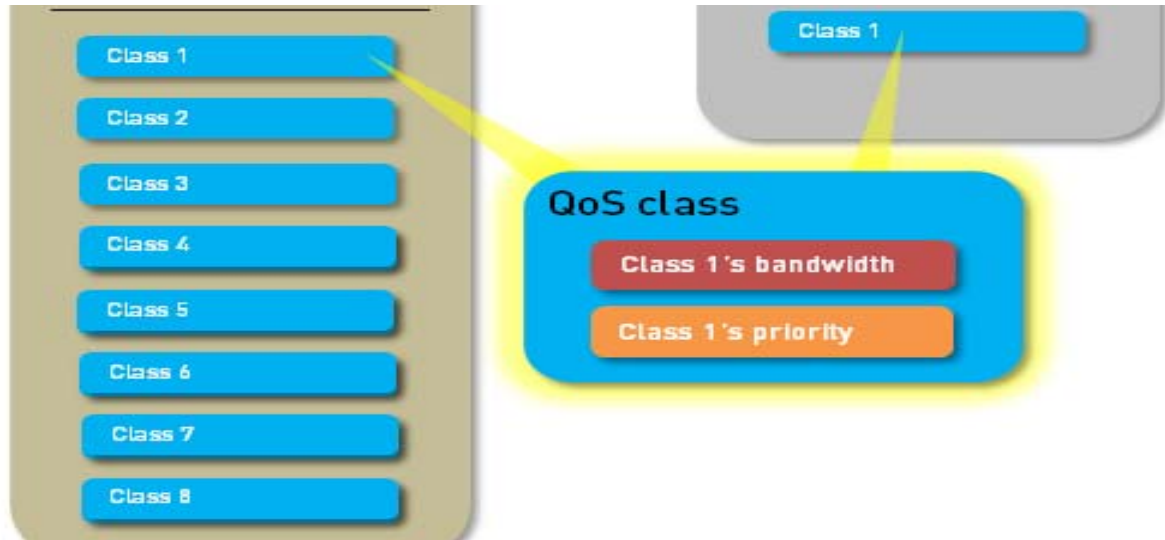


Real-time priority is typically used for applications that are particularly sensitive to latency, such as voice and video applications.

Use the QoS policy to assign a QoS class to specified traffic (**Policies > QoS**):

There are up to eight definable QoS classes in a single QoS profile. Unless otherwise configured, traffic that does not match a QoS class is assigned a class of 4.

QoS priority queuing and bandwidth management, the fundamental mechanisms of a QoS configuration, are configured within the QoS class definition (see [Step 3](#)). Queuing priority is determined by the priority set for a QoS class. Bandwidth management is determined according to the maximum and guaranteed bandwidths set for a QoS class.



The queuing and bandwidth management mechanisms determine the order of traffic and how traffic is handled upon entering or leaving a network:

- QoS Priority:** One of four QoS priorities can be defined in a QoS class: real-time, high, medium, and low. When a QoS class is associated with specific traffic, the priority defined in that QoS class is assigned to the traffic. Packets in the traffic flow are then queued according to their priority until the network is ready to process them. This method of priority queuing provides the capability to ensure that important traffic, applications, or users takes precedence.
- QoS Class Bandwidth Management:** QoS class bandwidth management provides the capability to control traffic flows on a network so that traffic does not exceed network capacity, resulting in network congestion, or to allocate specific bandwidth limits to traffic, applications, or users. You can set overall limits on bandwidth using the QoS profile or set limits for individual QoS classes. A QoS profile and QoS classes in the profile have guaranteed and maximum bandwidth limits. The guaranteed bandwidth limit (Egress Guaranteed) ensures that any amount of traffic up to that set bandwidth limit is processed. The maximum bandwidth limit (Egress Max) sets the total limit of bandwidth allocated to either the QoS Profile or QoS Class. Traffic in excess of the Maximum Bandwidth limit is dropped. The total bandwidth limits and guaranteed bandwidth limits of QoS classes in a QoS profile cannot exceed the bandwidth limit of the QoS profile.

QoS Policy

In a QoS configuration, the QoS policy identifies traffic requiring QoS treatment (either preferential treatment or bandwidth-limiting) using a defined parameter or multiple parameters and assigns it a class.

Use the QoS Policy, similar to a security policy, to set the criteria that identifies traffic:

- Applications and application groups.
- Source zones, source addresses, and source users.
- Destination zones and destination addresses.

- Services and service groups limited to specific TCP and/or UDP port numbers.
- URL categories, including custom URL categories.


The QoS Policy in the web interface (**Policies > QoS**) allows you to associate the criteria used to specify traffic with a QoS class.

The screenshot shows the 'QoS Policy Rule' configuration page. The 'General' tab is active, displaying a 'Class' dropdown menu set to '1' and a 'Schedule' dropdown menu set to 'None'. Other tabs include 'Source', 'Destination', 'Application', 'Service/URL Category', and 'Other Settings'.

QoS Egress Interface

Enabling a QoS profile on the egress interface of the traffic identified for QoS treatment completes a QoS configuration. The ingress interface for QoS traffic is the interface on which the traffic enters the firewall. The egress interface for QoS traffic is the interface that traffic leaves the firewall from. QoS is always enabled and enforced on a traffic flow's egress interface. The egress interface in a QoS configuration can either be the external- or internal-facing interface of the firewall, depending on the flow of the traffic receiving QoS treatment.

For example, in an enterprise network, if you are limiting employees' download traffic from a specific website, the egress interface in the QoS configuration is the firewall's internal interface, as the traffic flow is from the Internet, through the firewall, and to your company network. Alternatively, when limiting employees' upload traffic to the same website, the egress interface in the QoS configuration is the firewall's external interface, as the traffic you are limiting flows from your company network, through the firewall, and then to the Internet.



- 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.

See [Step 3](#) to learn how to [Identify the egress interface for applications that you identified as needing QoS treatment.](#)

QoS Cleartext and Tunneled Traffic

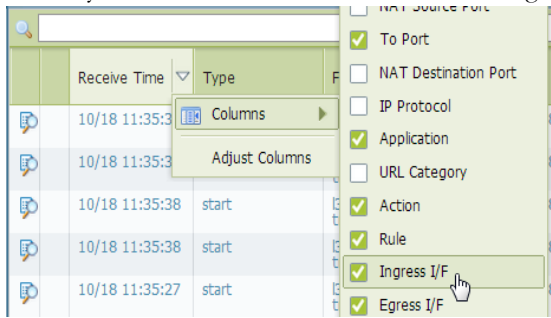

Within the QoS physical interface configuration, you can provide more granular QoS settings for cleartext traffic and tunneled traffic leaving through the interface. Individual tunnel interfaces can be assigned different QoS Profiles. Cleartext traffic can be assigned different QoS Profiles according to traffic's source interface and source subnet. In this case, a source interface and source subnet can be associated with a QoS Profile. If you choose not to select cleartext or tunneled traffic for unique QoS treatment, enabling QoS on an interface requires selecting a default QoS Profile to determine how to shape traffic for specific tunnel interfaces or, in the case of cleartext traffic, source interfaces and source subnets.



On Palo Alto Networks firewalls, the term “tunneled traffic” refers to tunnel interface traffic, specifically IPSec traffic in tunnel mode.

Configure QoS


Use the following task to configure Quality of Service (QoS), including how to create a QoS profile, create a QoS policy, and enable QoS on an interface.

Configure QoS	
<p>Step 1 Identify the traffic to which to apply QoS.</p> <p>This example shows how to use QoS to limit web browsing.</p>	<p>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.</p> <p>Click any application name to display detailed application information.</p>
<p>Step 2 Identify the egress interface for applications that you identified as needing QoS treatment.</p> <p>Tip: The egress interface for traffic depends on the traffic flow. If you are shaping incoming traffic, the egress interface is the internal-facing interface. If you are shaping outgoing traffic, the egress interface is the external-facing interface.</p>	<p>Select Monitor > Logs > Traffic to view the device's traffic logs.</p> <p>To filter and only show logs for a specific application:</p> <ul style="list-style-type: none"> • 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. <p>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:</p> <ul style="list-style-type: none"> • 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:  <p>In this example, the egress interface for web-browsing traffic is ethernet 1/1.</p>

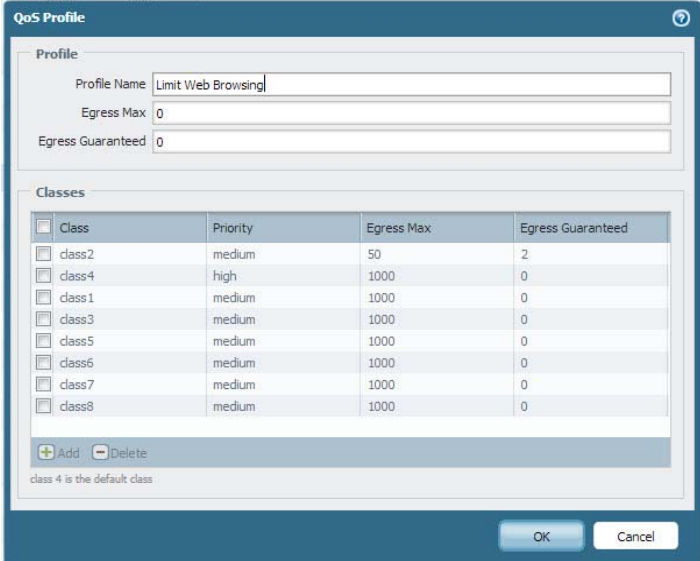
Configure QoS (Continued)**Step 3** Create a QoS profile.

You can edit any existing QoS profile, including the default, by clicking the QoS profile name.

1. Select **Network > Network Profiles > QoS Profile** and click **Add** to open the **QoS Profile** dialog.
2. Enter a descriptive **Profile Name**.
3. Enter an **Egress Max** to set the overall bandwidth allocation for the QoS Profile.
4. Enter an **Egress Guaranteed** to set the guaranteed bandwidth for the QoS Profile.

 Any traffic that exceeds the QoS Profile's egress guaranteed limit is best effort but is not guaranteed.
5. In the **Classes** section, specify how to treat up to eight individual QoS classes:
 - a. Click **Add** to add a class to the QoS Profile.
 - b. Select the **Priority** for the class.
 - c. Enter an **Egress Max** for a class to set the overall bandwidth limit for that individual class.
 - d. Enter an **Egress Guaranteed** for the class to set the guaranteed bandwidth for that individual class.
6. Click **OK** to save the QoS Profile.

In the following example, the QoS Profile named Limit Web Browsing limits traffic identified as Class 2 traffic to maximum bandwidth of 50 Mbps and a guaranteed bandwidth of 2 Mbps. Any traffic that is associated with class 2 in a the QoS policy ([Step 4](#)) is subject to these limits.



QoS Profile

Profile

Profile Name:

Egress Max:

Egress Guaranteed:

Classes

Class	Priority	Egress Max	Egress Guaranteed
<input type="checkbox"/> class2	medium	50	2
<input type="checkbox"/> class4	high	1000	0
<input type="checkbox"/> class1	medium	1000	0
<input type="checkbox"/> class3	medium	1000	0
<input type="checkbox"/> class5	medium	1000	0
<input type="checkbox"/> class6	medium	1000	0
<input type="checkbox"/> class7	medium	1000	0
<input type="checkbox"/> class8	medium	1000	0

class 4 is the default class

Configure QoS (Continued)

Step 4 Create a QoS policy.

1. Select **Policies > QoS** and click **Add** to open the QoS Policy Rule dialog.
2. On the **General** tab, give the QoS Policy Rule a descriptive **Name**.
3. Specify the traffic to which the QoS Policy Rule will apply. Use the **Source**, **Destination**, **Application**, and **Service/URL Category** tabs to define matching parameters for identifying traffic.


For example, select the **Application** tab, click **Add** and select web-browsing to apply the QoS Policy Rule to that application:



(Optional) Define additional parameters. For example, on the **Source** tab, click **Add** to limit a specific user's web-browsing, in this case, user1:



4. On the **Other Settings** tab, select a QoS Class to assign to the QoS Policy Rule. For example, assign Class 2 to the user1's web-browsing traffic:



5. Click **OK** to save the QoS Policy Rule.

Configure QoS (Continued)

Step 5 Enable the QoS Profile on a physical interface.

You can configure settings to select cleartext and tunneled traffic for unique QoS treatment, in addition to the QoS configuration on the physical interface:

To configure specific settings for cleartext traffic using traffic's source interface and source subnet as criteria for QoS identification and treatment, perform [Step 5 - 4](#).

To apply a QoS profile to a specific tunnel interface(s), perform [Step 5 - 5](#).

For more information, see [QoS Cleartext and Tunneled Traffic](#).



Check if the platform you're using supports enabling QoS on a sub-interface by reviewing a summary of the [Product Specifications](#).



It is a best practice to always define the **Egress Max** value for a QoS interface.

1. Select **Network > QoS** and click **Add** to open the QoS Interface dialog.
2. Enable QoS on the physical interface:
 - a. On the **Physical Interface** tab, select the **Interface Name** of the interface to apply the QoS Profile to.
In the example, Ethernet 1/1 is the egress interface for web-browsing traffic (see [Step 2](#)).
 - b. Select **Turn on QoS feature on this interface**.
3. On the **Physical Interface** tab, select a QoS profile to apply by default to all **Clear Text** traffic.
(Optional) Use the Tunnel Interface field to apply a QoS profile by default to all tunneled traffic.

For example, enable QoS on ethernet 1/1 and apply the QoS Profile named Limit Web Browsing as the default QoS Profile for clear text traffic.

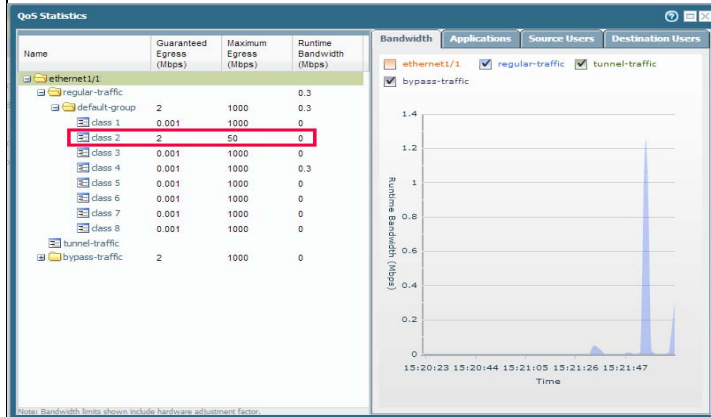
4. (Optional) On the **Clear Text Traffic** tab, configure more granular QoS settings for cleartext traffic:
 - Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.
 - Click **Add** to apply a QoS Profile to selected clear text traffic, further selecting the traffic for QoS treatment according to source interface and source subnet (creating a QoS node).
5. (Optional) On the **Tunneled Traffic** tab, configure more granular QoS settings for tunnel interfaces:
 - Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.
 - Click **Add** to associate a selected tunnel interface with a QoS Profile.
6. Click **OK** to save the QoS Profile.
7. **Commit** the changes to enable the QoS Profile on the interface.

Configure QoS (Continued)

Step 6 Verify QoS configuration.


Select **Network > QoS** to view the **QoS Policies** page and click the **Statistics** link to view QoS bandwidth, active sessions of a selected QoS node or class, and active applications for the selected QoS node or class.

For example, see the statistics for ethernet 1/1 with QoS enabled:



Class 2 traffic limited to 2 Mbps of guaranteed bandwidth and a maximum bandwidth of 50 Mbps.

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.

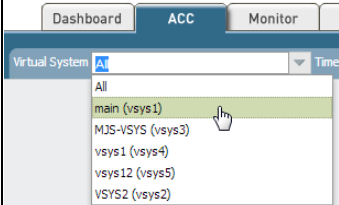
Configure QoS for a Virtual System

QoS can be configured for a single or several [Virtual Systems](#) within a Palo Alto Networks firewall. Because a virtual system is an independent firewall, QoS must be configured independently for a single virtual system to apply a QoS configuration to only that virtual system.

Configuring QoS for a virtual system is similar to configuring QoS on a physical firewall, with the exception that configuring QoS for a virtual system requires specifying the traffic flow's source and destination zones and source and destination interfaces. Because a virtual system exists without set physical boundaries (such as a physical interface) that traffic flows through, specifying source and destination zones and interfaces of a traffic flow allows you to control and shape traffic for that virtual system specifically, as a traffic flow spans more than one virtual system in a virtual environment.

The example below shows two virtual systems configured within a firewall. VSYS 1 (purple) and VSYS 2 (red) each have QoS configured to prioritize or limit two distinct traffic flows, indicated by their corresponding purple (VSYS 1) and red (VSYS 2) lines. The QoS nodes indicate the points at which QoS traffic is identified and then shaped in each virtual system.



Configure QoS in a Virtual System Environment	
<p>Step 1 Confirm that the appropriate interfaces, virtual routers, and security zones are associated with each virtual system.</p>	<ul style="list-style-type: none"> 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.
<p>Step 2 Identify traffic to apply QoS to.</p>	<p>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.</p> <p>To view information for a specific virtual system, select the virtual system from the Virtual System drop-down:</p>  <p>Click any application name to display detailed application information.</p>

Configure QoS in a Virtual System Environment

Step 3 Identify the egress interface for applications that you identified as needing QoS treatment.

In a virtual system environment, QoS is applied to traffic on the traffic's egress point on the virtual system. Depending on a virtual system's configuration and the QoS policy, the egress point of QoS traffic could be associated with a physical interface or could be a configured zone.

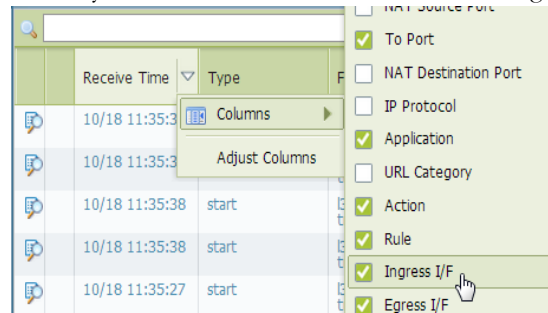
This example shows how to limit web-browsing traffic on vsys 1.

Select **Monitor > Logs > Traffic** to view the device's traffic logs. Each entry has the option to display columns with information necessary to configure QoS in a virtual system environment:

- virtual system
- egress interface
- ingress interface
- source zone
- destination zone

To display a column if it is not displayed by default:

- Click any column header to add a column to the log:



- Click the spyglass icon to the left of any entry to display a detailed log that includes the application's egress interface, as well as source and destination zones, in the **Source** and **Destination** sections:


Source		Destination	
User		User	
Address	10.22.2.68	Address	10.0.0.106
Country	10.0.0.0-10.255.255.255	Country	10.0.0.0-10.255.255.255
Port	61865	Port	80
Zone	trust	Zone	untrust
Interface	ethernet1/2	Interface	ethernet1/1

Session Counters		Flags	
Bytes	1980	Captive Portal	<input type="checkbox"/>
Bytes Received	858	Proxy Transaction	<input type="checkbox"/>
Bytes Sent	1122	Decrypted	<input type="checkbox"/>
Repeat Count	1	Packet Capture	<input type="checkbox"/>
Packets	12	Client to Server	<input type="checkbox"/>
Packets Received	4	Server to Client	<input type="checkbox"/>
Packets Sent	8	Symmetric Return	<input type="checkbox"/>
		Mirrored	<input type="checkbox"/>

For example, for web-browsing traffic from VSYS 1, the ingress interface is ethernet 1/2, the egress interface is ethernet 1/1, the source zone is *trust* and the destination zone is *untrust*.

Configure QoS in a Virtual System Environment**Step 4** Create a QoS Profile.

You can edit any existing QoS Profile, including the default, by clicking the profile name.

1. Select **Network > Network Profiles > QoS Profile** and click **Add** to open the QoS Profile dialog.
2. Enter a descriptive **Profile Name**.
3. Enter an **Egress Max** to set the overall bandwidth allocation for the QoS profile.
4. Enter an **Egress Guaranteed** to set the guaranteed bandwidth for the QoS profile.
 Any traffic that exceeds the QoS profile's egress guaranteed limit is best effort but is not guaranteed.
5. In the Classes section of the **QoS Profile**, specify how to treat up to eight individual QoS classes:
 - a. Click **Add** to add a class to the QoS Profile.
 - b. Select the **Priority** for the class.
 - c. Enter an **Egress Max** for a class to set the overall bandwidth limit for that individual class.
 - d. Enter an **Egress Guaranteed** for the class to set the guaranteed bandwidth for that individual class.
6. Click **OK** to save the QoS profile.

Configure QoS in a Virtual System Environment

Step 5 Create a QoS policy.


In a multi-VSYS environment, traffic can span more than one virtual system before the ingress point on the virtual system you are configuring QoS for. Specifying source and destination zones for QoS traffic ensures that the traffic is correctly identified as it flows through the specific virtual system (in this example, vsys 1) and QoS is applied to the traffic only for that designated virtual system (and not applied to traffic for other configured virtual systems).

1. Select **Policies > QoS** and click **Add** to open the QoS Policy Rule dialog.
2. On the **General** tab, give the QoS Policy Rule a descriptive **Name**.
3. Specify the traffic to which the QoS policy rule will apply. Use the **Source**, **Destination**, **Application**, and **Service/URL Category** tabs to define matching parameters for identifying traffic.


For example, select the **Application** tab, click **Add** and select web-browsing to apply the QoS Policy Rule to that application:




4. On the **Source** tab, click **Add** to select the source zone of vsys 1's web-browsing traffic.



5. On the **Destination** tab, click **Add** to select the destination zone of vsys 1's web-browsing traffic.



6. On the **Other Settings** tab, and select a **QoS Class** to assign to the QoS policy rule. For example, assign Class 2 to web-browsing traffic on vsys 1:



7. Click **OK** to save the QoS policy rule.

Configure QoS in a Virtual System Environment

Step 6 Enable the QoS Profile on a physical interface.



It is a best practice to always define the **Egress Max** value for a QoS interface.

1. Select **Network > QoS** and click **Add** to open the QoS Interface dialog.
2. Enable QoS on the physical interface:
 - a. On the **Physical Interface** tab, select the **Interface Name** of the interface to apply the QoS Profile to.

In this example, ethernet 1/1 is the egress interface for web-browsing traffic on vsys 1 (see [Step 2](#)).

- b. Select **Turn on QoS feature on this interface**.
3. On the **Physical Interface** tab, select the default QoS profile to apply to all **Clear Text** traffic.
(Optional) Use the **Tunnel Interface** field to apply a default QoS profile to all tunneled traffic.
4. (Optional) On the **Clear Text Traffic** tab, configure additional QoS settings for clear text traffic:
 - Set the **Egress Guaranteed** and **Egress Max** bandwidths for clear text traffic.
 - Click **Add** to apply a QoS Profile to selected clear text traffic, further selecting the traffic for QoS treatment according to source interface and source subnet (creating a QoS node).
5. (Optional) On the **Tunneled Traffic** tab, configure additional QoS settings for tunnel interfaces:
 - Set the **Egress Guaranteed** and **Egress Max** bandwidths for tunneled traffic.
 - Click **Add** to associate a selected tunnel interface with a QoS Profile.
6. Click **OK** to save changes.
7. **Commit** the changes.

Configure QoS in a Virtual System Environment

Step 7 Verify QoS configuration.

- Select **Network > QoS** to view the QoS Policies page. The **QoS Policies** page verifies that QoS is enabled and includes a **Statistics** link. Click the Statistics link to view QoS bandwidth, active sessions of a selected QoS node or class, and active applications for the selected QoS node or class.
- In a multi-VSYS environment, sessions cannot span multiple systems. Multiple sessions are created for one traffic flow if the traffic passes through more than one virtual system. To browse sessions running on the firewall and view applied QoS Rules and QoS Classes, select **Monitor > Session Browser**.

QoS Use Case Examples

The following use cases demonstrate how to use QoS in common scenarios:

- ▲ [Use Case: QoS for a Single User](#)
- ▲ [QoS for Voice and Video Applications](#)

Use Case: QoS for a Single User

A CEO finds that during periods of high network usage, she is unable to access enterprise applications to respond effectively to critical business communications. The IT admin wants to ensure that all traffic to and from the CEO receives preferential treatment over other employee traffic so that she is guaranteed not only access to, but high performance of, critical network resources.

Apply QoS to a Single User

Step 1 The admin creates the QoS profile *CEO_traffic* to define how traffic originating from the CEO will be treated and shaped as it flows out of the company network:

Class	Priority	Egress Max	Egress Guaranteed
<input checked="" type="checkbox"/> class1	high		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.

Apply QoS to a Single User (Continued)

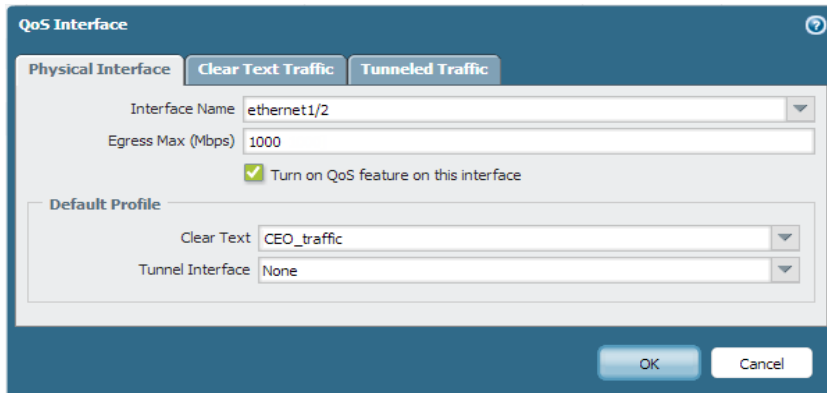
Step 2 The admin creates a QoS policy to identify the CEO’s traffic (**Policies > QoS**) and assigns it the class that he defined in the QoS profile (see **Step 1**). Because User-ID is configured, the admin uses the **Source** tab in the QoS policy to singularly identify the CEO’s traffic by her company network username. (If User-ID is not configured, the administrator could **Add** the CEO’s IP address under **Source Address**. See **User-ID**):



The admin associates the CEO’s traffic with Class 1 (**Other Settings** tab) and then continues to populate the remaining required policy fields; the admin gives the policy a descriptive **Name** (**General** tab) and selects **Any** for the **Source Zone** (**Source** tab) and **Destination Zone** (**Destination** tab):

Name	Tags	Zone	Source			Destination		Application	Service	Class	Schedule
			Address	User	Zone	Address					
1 Video	none	any	any	any	any	any	any	google-video http-video youtube	any	1	none
2 HTTPS	none	any	any	companynetwork\JoeAdmin	any	any	any	web-browsing	any	2	none
3 FTP	none	any	any	any	any	any	any	ftp	any	4	none
4 Guarantee CEO bandwidth	none	any	any	companynetwork\CEO	any	any	any	any	any	1	none

Step 3 Now that Class 1 is associated with the CEO’s traffic, the admin enables QoS by checking **Turn on QoS feature on interface** and selecting the traffic flow’s egress interface. The egress interface for the CEO’s traffic flow is the external-facing interface, in this case, ethernet 1/2:



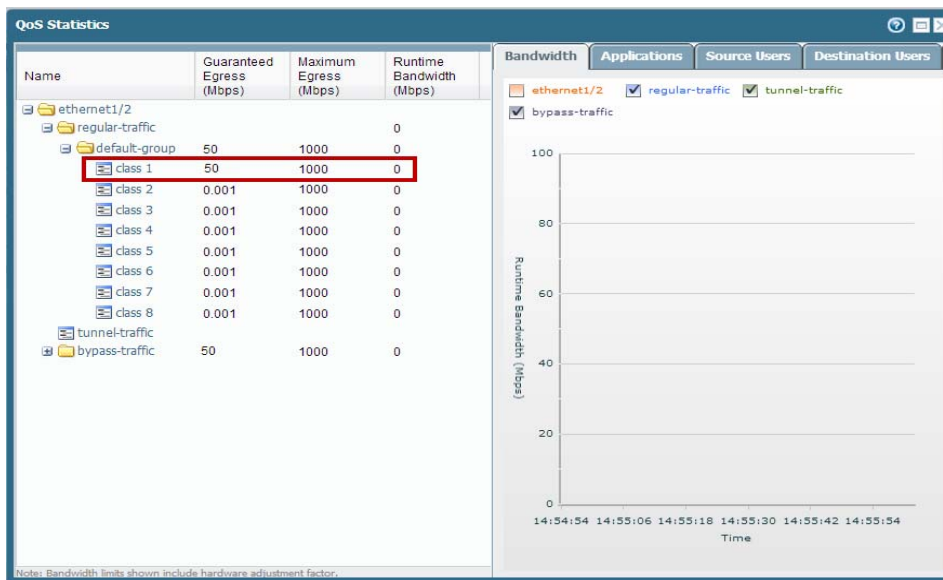
Because the admin wants to ensure that all traffic originating from the CEO is guaranteed by the QoS profile and associated QoS policy he created, he selects the *CEO_traffic* to apply to **Clear Text** traffic flowing from ethernet 1/2.

Apply QoS to a Single User (Continued)

Step 4 After committing the QoS configuration, the admin navigates to the **Network > QoS** page to confirm that the QoS profile *CEO_traffic* is enabled on the external-facing interface, ethernet 1/2:

Name	Guaranteed Egress (Mbps)	Maximum Egress (Mbps)	Profile	Enabled	
ethernet1/2		1,000.000		<input checked="" type="checkbox"/>	Statistics
Tunneled Traffic	0.000	0.000			
Clear Text Traffic	50.000	0.000	CEO_traffic		
ethernet1/18		0.000		<input checked="" type="checkbox"/>	Statistics
Tunneled Traffic	0.000	0.000			
Clear Text Traffic	0.000	0.000	Limit Facebook apps		
Facebook Apps (ethernet1/19 - any)			Limit Facebook apps		

He clicks **Statistics** to view how traffic originating with the CEO (Class 1) is being shaped as it flows from ethernet 1/2:



This case demonstrates how to apply QoS to traffic originating from a single source user. However, if you also wanted to guarantee or shape traffic to a destination user, you could configure a similar QoS setup. Instead of, or in addition to this work flow, create a QoS policy that specifies the user's IP address as the **Destination Address** on the **Policies > QoS** page (instead of specifying the user's source information, as shown in [Step 2](#)) and then enable QoS on the network's internal-facing interface on the **Network > QoS** page (instead of the external-facing interface, as shown in [Step 3](#).)

QoS for Voice and Video Applications

Voice and video traffic is particularly sensitive to measurements that the QoS feature shapes and controls, especially latency and jitter. For voice and video transmissions to be audible and clear, voice and video packets cannot be dropped, delayed, or delivered inconsistently. A best practice for voice and video applications, in addition to guaranteeing bandwidth, is to guarantee priority to voice and video traffic.

In this example, employees at a company branch office are experiencing difficulties and unreliability in using video conferencing and Voice-over-IP (VoIP) technologies to conduct business communications with other branch offices, with partners, and with customers. An IT admin intends to implement QoS in order to address these issues and ensure effective and reliable business communication for the branch employees. Because the admin wants to guarantee QoS to both incoming and outgoing network traffic, he will enable QoS on both the firewall's internal- and external-facing interfaces.

Ensure Quality for Voice and Video Applications

Step 1 The admin creates a QoS profile, defining Class 2 so that any traffic associated with Class 2 receives real-time priority and on an interface with a maximum bandwidth of 1000 Mbps, is guaranteed a bandwidth of 250 Mbps at all times, including peak periods of network usage.

Real-time priority is typically recommended for applications affected by latency, and is particularly useful in guaranteeing performance and quality of voice and video applications.

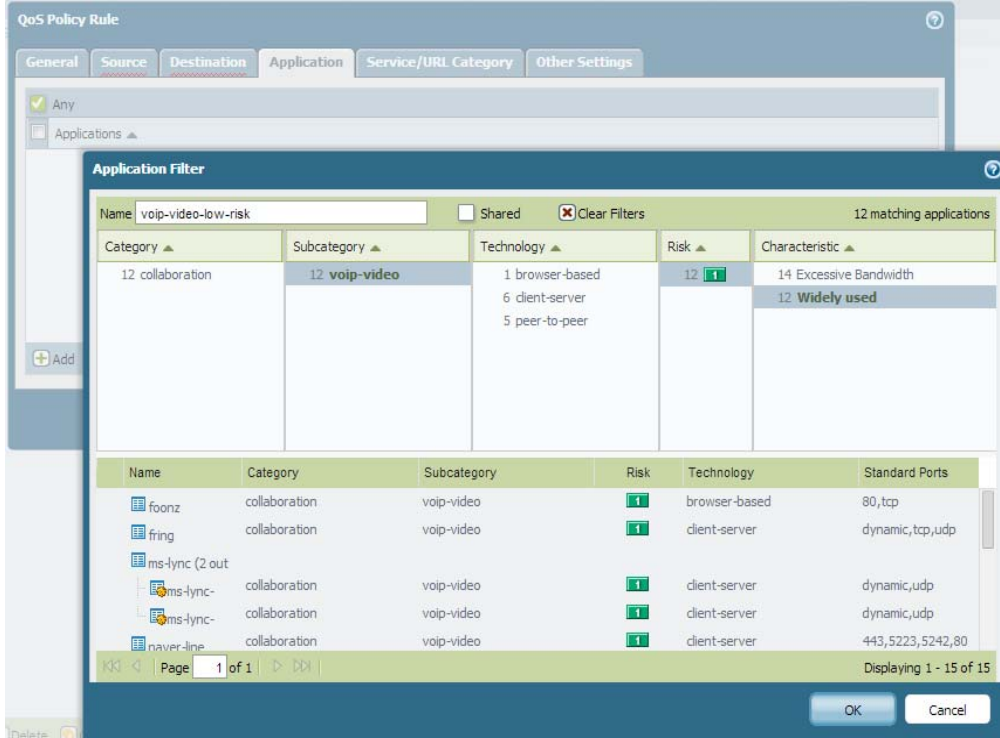
On the **Network > Network Profiles > QoS Profile** page, the admin clicks **Add**, enters the **Profile Name** *ensure voip-video traffic* and defines Class 2 traffic.

QoS Profile			
Profile			
Profile Name	ensure voip-video traffic		
Egress Max	1000		
Egress Guaranteed	250		
Classes			
Class	Priority	Egress Max	Egress Guaranteed
<input checked="" type="checkbox"/> class2	real-time	1000	250


Ensure Quality for Voice and Video Applications (Continued)

Step 2 The admin creates a QoS policy to identify voice and video traffic. Because the company does not have one standard voice and video application, the admin wants to ensure QoS is applied to a few applications that are widely and regularly used by employees to communicate with other offices, with partners, and with customers. On the **Policies > QoS > QoS Policy Rule > Applications** tab, the admin clicks **Add** and opens the **Application Filter** window. The admin continues by selecting criteria to filter the applications he wants to apply QoS to, choosing the Subcategory *voip-video*, and narrowing that down by specifying only voip-video applications that are both low-risk and widely-used.

The application filter is a dynamic tool that, when used to filter applications in the QoS policy, allows QoS to be applied to all applications that meet the criteria of *voip-video*, *low risk*, and *widely used* at any given time.



The admin names the **Application Filter** *voip-video-low-risk* and includes it in the QoS policy:



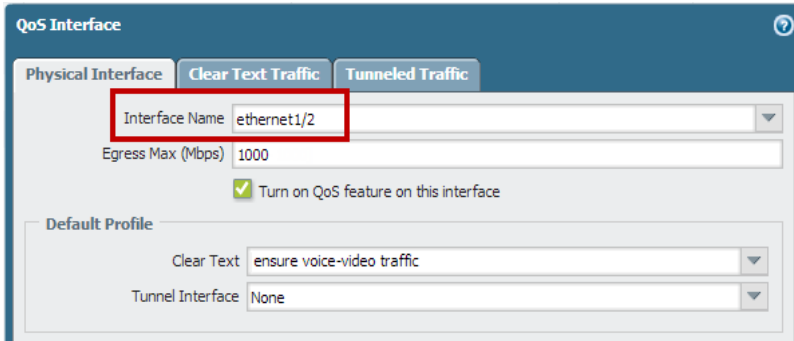
The admin names the QoS policy *Voice-Video* and associates the *voip-video-low-risk* application filter with Class 2 traffic (as he defined it in [Step 1](#)). He is going to use the *Voice-Video* QoS policy for both incoming and outgoing QoS traffic, so he sets **Source** and **Destination** information to **Any**:

2	HTTPS	none	any	any	companynet...	any	any	web-browsing	any	2	none
3	FTP	none	any	any	any	any	any	ftp	any	4	none
4	Voice-Video	none	any	any	any	any	any	voip-video-lo...	any	2	none

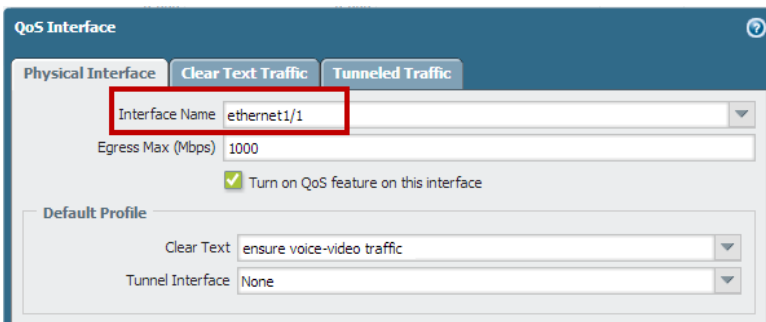
Ensure Quality for Voice and Video Applications (Continued)

Step 3 Because the admin wants to ensure QoS for both incoming and outgoing voice and video communications, he enables QoS on the network’s external-facing interface (to apply QoS to outgoing communications) and to the internal-facing interface (to apply QoS to incoming communications).

The admin begins by enabling the QoS profile he created in **Step 1**, *ensure voice-video traffic* (Class 1 in this profile is associated with policy created in **Step 2**, *Voice-Video*) on the external-facing interface, in this case, ethernet 1/2.



He then enables the same QoS profile *ensure voip-video traffic* on the internal-facing interface, in this case, ethernet 1/1.



Step 4 The admin confirms that QoS is enabled for both incoming and outgoing voice and video traffic:

Name	Guaranteed Egress (Mbps)	Maximum Egress (Mbps)	Profile	Enabled	
ethernet1/1		1,000.000		<input checked="" type="checkbox"/>	Statistics
Tunneled Traffic	0.000	0.000			
Clear Text Traffic	250.000	0.000	ensure voice-video traffic		
ethernet1/2		1,000.000		<input checked="" type="checkbox"/>	Statistics
Tunneled Traffic	0.000	0.000			
Clear Text Traffic	250.000	0.000	ensure voice-video traffic		

The admin has successfully enabled QoS on both the network’s internal- and external-facing interfaces. Real-time priority is now ensured for voice and video application traffic as it flows both into and out of the network, ensuring that these communications, which are particularly sensitive to latency and jitter, can be used reliably and effectively to perform both internal and external business communications.



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.

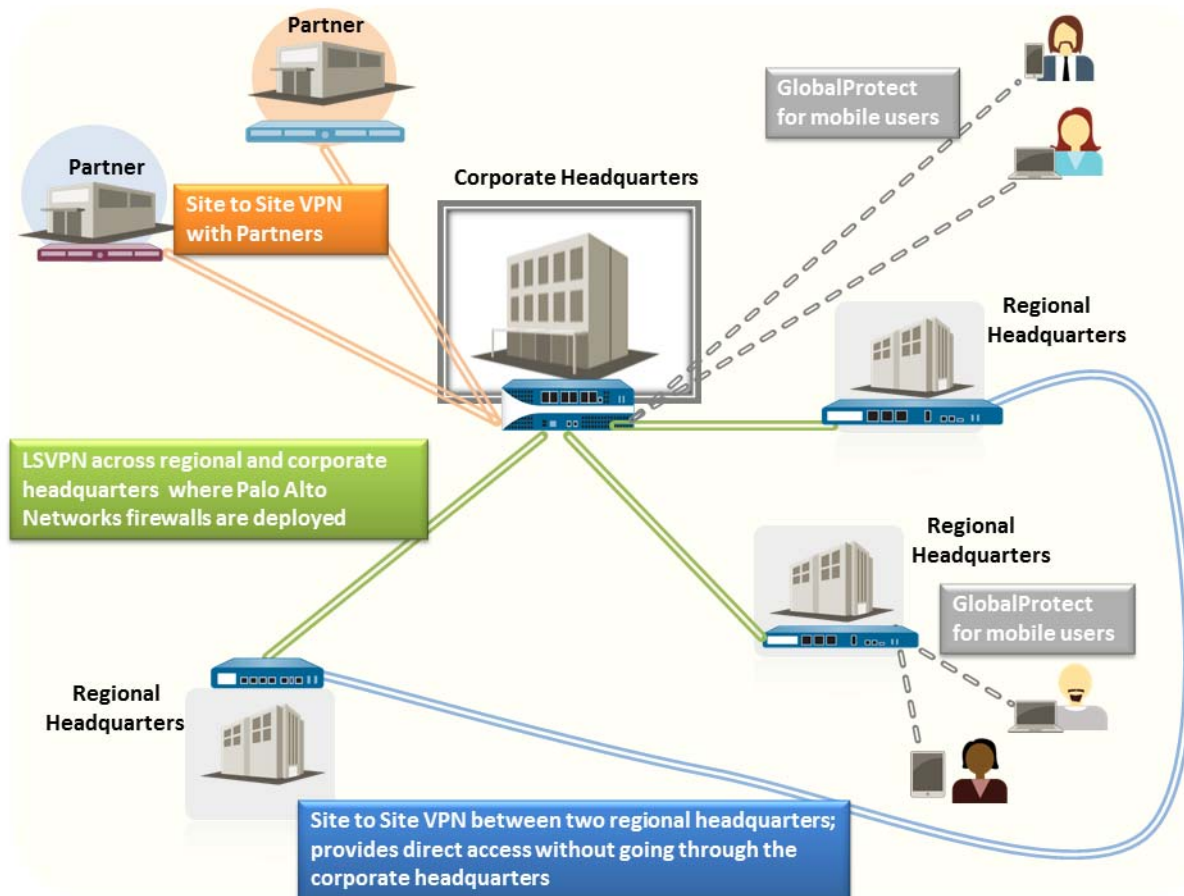
- ▲ [VPN Deployments](#)
- ▲ [Site-to-Site VPN Overview](#)
- ▲ [Site-to-Site VPN Concepts](#)
- ▲ [Set Up Site-to-Site VPN](#)
- ▲ [Site-to-Site VPN Quick Configs](#)

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 1024 satellite offices. The solution requires Palo Alto Networks firewalls to be deployed at the hub and at every spoke. It uses certificates for device authentication, SSL for securing communication between all components, and IPSec to secure data. See [Large Scale VPN \(LSVPN\)](#).

Figure: VPN Deployments



Site-to-Site VPN Overview

A VPN connection that allows you to connect two Local Area Networks (LANs) is called a site-to-site VPN. You can configure route-based VPNs to connect Palo Alto Networks firewalls located at two sites or to connect a Palo Alto Networks firewall with a third-party security device at another location. The firewall can also interoperate with third-party policy-based VPN devices; the Palo Alto Networks firewall supports route-based VPN.

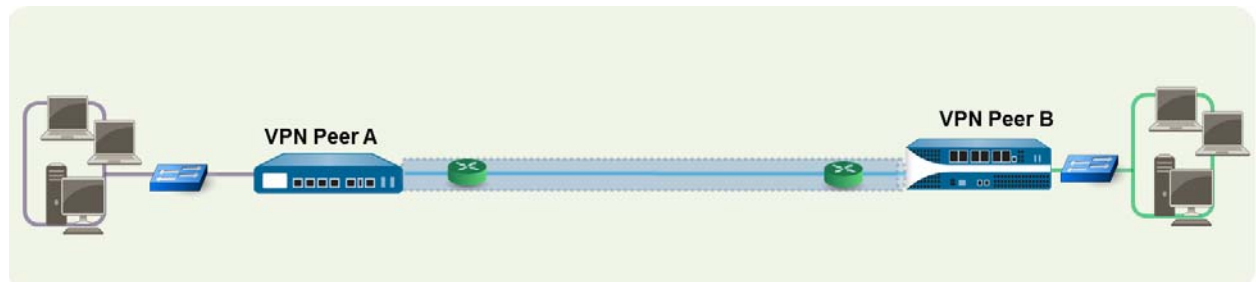
The Palo Alto Networks firewall sets up a route-based VPN, where the firewall makes a routing decision based on the destination IP address. If traffic is routed to a specific destination through a VPN tunnel, then it is handled as VPN traffic.

The IP Security (IPSec) set of protocols is used to set up a secure tunnel for the VPN traffic, and the information in the TCP/IP packet is secured (and encrypted if the tunnel type is ESP). The IP packet (header and payload) is embedded in another IP payload, and a new header is applied and then sent through the IPSec tunnel. The source IP address in the new header is that of the local VPN peer and the destination IP address is that of the VPN peer on the far end of the tunnel. When the packet reaches the remote VPN peer (the firewall at the far end of the tunnel), the outer header is removed and the original packet is sent to its destination.

In order to set up the VPN tunnel, first the peers need to be authenticated. After successful authentication, the peers negotiate the encryption mechanism and algorithms to secure the communication. The Internet Key Exchange (IKE) process is used to authenticate the VPN peers, and IPSec Security Associations (SAs) are defined at each end of the tunnel to secure the VPN communication. IKE uses digital certificates or preshared keys, and the Diffie Hellman keys to set up the SAs for the IPSec tunnel. The SAs specify all of the parameters that are required for secure transmission—including the security parameter index (SPI), security protocol, cryptographic keys, and the destination IP address—encryption, data authentication, data integrity, and endpoint authentication.

The following figure shows a VPN tunnel between two sites. When a client that is secured by VPN Peer A needs content from a server located at the other site, VPN Peer A initiates a connection request to VPN Peer B. If the security policy permits the connection, VPN Peer A uses the IKE Crypto profile parameters (IKE phase 1) to establish a secure connection and authenticate VPN Peer B. Then, VPN Peer A establishes the VPN tunnel using the IPSec Crypto profile, which defines the IKE phase 2 parameters to allow the secure transfer of data between the two sites.

Figure: Site-to-Site VPN



Site-to-Site VPN Concepts

A VPN connection provides secure access to information between two or more sites. In order to provide secure access to resources and reliable connectivity, a VPN connection needs the following components:

- ▲ [IKE Gateway](#)
- ▲ [Tunnel Interface](#)
- ▲ [Tunnel Monitoring](#)
- ▲ [Internet Key Exchange \(IKE\) for VPN](#)

IKE Gateway

The Palo Alto Networks firewalls or a firewall and another security device that initiate and terminate VPN connections across the two networks are called the IKE Gateways. To set up the VPN tunnel and send traffic between the IKE Gateways, each peer must have an IP address—static or dynamic—or FQDN. The VPN peers use preshared keys or certificates to mutually authenticate each other.

The peers must also negotiate the mode—main or aggressive—for setting up the VPN tunnel and the SA lifetime in IKE Phase 1. Main mode protects the identity of the peers and is more secure because more packets are exchanged when setting up the tunnel. Main mode is the recommended mode for IKE negotiation if both peers support it. Aggressive mode uses fewer packets to set up the VPN tunnel and is hence faster but a less secure option for setting up the VPN tunnel.

See [Set up an IKE Gateway](#) for configuration details.

Tunnel Interface

To set up a VPN tunnel, the Layer 3 interface at each end must have a logical *tunnel* interface for the firewall to connect to and establish a VPN tunnel. A tunnel interface is a logical (virtual) interface that is used to deliver traffic between two endpoints. 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 ProxyIDs. Each Proxy ID counts towards the IPsec VPN tunnel capacity of the firewall, and the tunnel capacity varies by the firewall model.

See [Set up an IPsec Tunnel](#) for configuration details.

Tunnel Monitoring

For a VPN tunnel, you can check connectivity to a destination IP address across the tunnel. The network monitoring profile on the firewall allows you to verify connectivity (using ICMP) to a destination IP address or a next hop at a specified polling interval, and to specify an action on failure to access the monitored IP address.

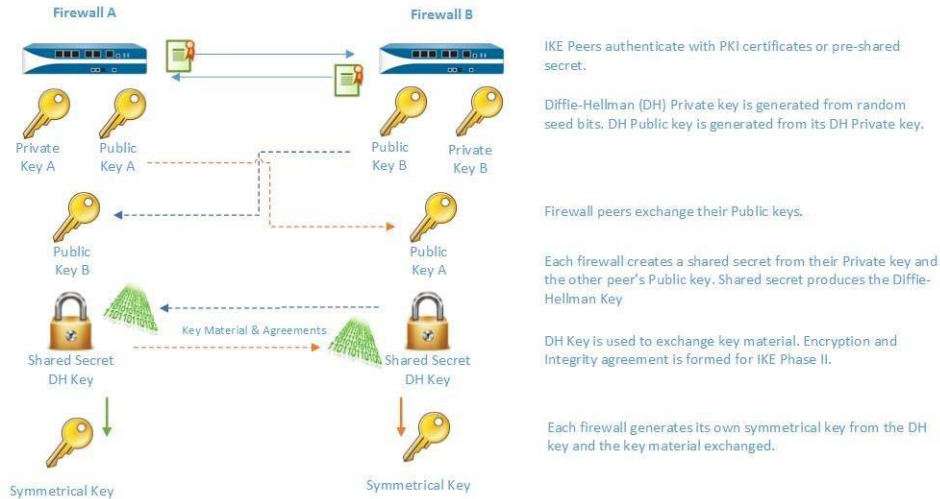
If the destination IP is unreachable, you either configure the firewall to wait for the tunnel to recover or configure automatic failover to another tunnel. In either case, the firewall generates a system log that alerts you to a tunnel failure and renegotiates the IPsec keys to accelerate recovery.

The default monitoring profile is configured to wait for the tunnel to recover; the polling interval is 3 seconds and the failure threshold is 5.

See [Set up Tunnel Monitoring](#) for configuration details.

Internet Key Exchange (IKE) for VPN

The IKE process allows the VPN peers at both ends of the tunnel to encrypt and decrypt packets using mutually agreed-upon keys or certificate and method of encryption. The IKE process occurs in two phases: [IKE Phase 1](#) and [IKE Phase 2](#). Each of these phases use keys and encryption algorithms that are defined using cryptographic profiles— IKE crypto profile and IPsec crypto profile—and the result of the IKE negotiation is a Security Association (SA). An SA is a set of mutually agreed-upon keys and algorithms that are used by both VPN peers to allow the flow of data across the VPN tunnel. The following illustration depicts the key exchange process for setting up the VPN tunnel:



IKE Phase 1

In this phase, the firewalls use the parameters defined in the IKE Gateway configuration and the IKE Crypto profile to authenticate each other and set up a secure control channel. IKE Phase supports the use of preshared keys or digital certificates (which use public key infrastructure, PKI) for mutual authentication of the VPN peers. Preshared keys are a simple solution for securing smaller networks because they do not require the support of a PKI infrastructure. Digital certificates can be more convenient for larger networks or implementations that require stronger authentication security.

When using certificates, make sure that the CA issuing the certificate is trusted by both gateway peers and that the maximum length of certificates in the certificate chain is 5 or less. With IKE fragmentation enabled, the firewall can reassemble IKE messages with up to 5 certificates in the certificate chain and successfully establish a VPN tunnel.

The IKE-crypto profile defines the following options that are used in the IKE SA negotiation:

- Diffie-Hellman (DH) Group for generating symmetrical keys for IKE. The Diffie Hellman algorithm uses the private key of one party and the public key of the other to create a shared secret, which is an encrypted key that is shared by both VPN tunnel peers. The DH groups supported on the firewall are: Group 1—768 bits; Group 2—1024 bits (the default); Group 5—1536 bits; Group 14—2048 bits.
- Authentication options—sha1; sha 256; sha 384; sha 512; md5
- Encryption algorithms—3des; aes128; aes192; aes256

IKE Phase 2

After the tunnel is secured and authenticated, in Phase 2 the channel is further secured for the transfer of data between the networks. IKE Phase 2 uses the keys that were established in Phase 1 of the process and the IPsec Crypto profile, which defines the IPsec protocols and keys used for the SA in IKE Phase 2.

The IPSEC uses the following protocols to enable secure communication:

- Encapsulating Security Payload (ESP)—Allows you to encrypt the entire IP packet, and authenticate the source and verify integrity of the data. While ESP requires that you encrypt and authenticate the packet, you can choose to only encrypt or only authenticate by setting the encryption option to Null; using encryption without authentication is discouraged.
- Authentication Header (AH)—Authenticates the source of the packet and verifies data integrity. AH does not encrypt the data payload and is unsuited for deployments where data privacy is important. AH is commonly used when the main concern is to verify the legitimacy of the peer, and data privacy is not required.

Table: Algorithms Supported for IPSEC Authentication and Encryption

ESP	AH
Diffie Hellman Exchange options supported	
<ul style="list-style-type: none"> • Group 1—768 bits • Group 2—1024 bits (the default) • Group 5—1536 bits • Group 14—2048 bits. • no-pfs—By default, Perfect Forward Secrecy (pfs) is enabled. With PFS is enabled, a new DH key is generated in IKE phase 2 using one of the groups listed above; this key is independent of the keys exchanged in IKE phase1, and therefore allows for more secure transfer of data. <p>No-pfs implies that the DH key created at phase 1 is not renewed and a single key is used for the IPSEC SA negotiations. Both VPN peers must be enabled or disabled for Perfect Forward Secrecy.</p>	
Encryption algorithms supported	
• 3des	
• aes128	
• aes192	
• aes256	
• aes128ccm16	
• null	
Authentication algorithms supported	
• md5	• md5
• sha 1	• sha 1
• sha 256	• sha 256
• sha 384	• sha 384
• sha512	• sha 512
• none	

Methods of Securing IPSec VPN Tunnels (IKE Phase 2)

IPSec VPN tunnels can be secured using manual keys or auto keys. In addition, IPSec configuration options include Diffie-Hellman Group for key agreement, and/or an encryption algorithm and a hash for message authentication.

- **Manual Key**—Manual key is typically used if the Palo Alto Networks firewall is establishing a VPN tunnel with a legacy device, or if you want to reduce the overhead of generating session keys. If using manual keys, the same key must be configured on both peers.

Manual keys are not recommended for establishing a VPN tunnel because the session keys can be compromised when relaying the key information between the peers; if the keys are compromised, the data transfer is no longer secure.

- **Auto Key**— Auto Key allows you to automatically generate keys for setting up and maintaining the IPSec tunnel based on the algorithms defined in the IPSec Crypto profile.

Set Up Site-to-Site VPN

To set up site-to-site VPN:

- ❑ Make sure that your Ethernet interfaces, virtual routers, and zones are configured properly. For more information, see [Set Up Interfaces and Zones](#).
- ❑ Create your tunnel interfaces. Ideally, put the tunnel interfaces in a separate zone, so that tunneled traffic can use different policies.
- ❑ Set up static routes or assign routing protocols to redirect traffic to the VPN tunnels. To support dynamic routing (OSPF, BGP, RIP are supported), you must assign an IP address to the tunnel interface.
- ❑ Define IKE gateways for establishing communication between the peers across each end of the VPN tunnel; also define the cryptographic profile that specifies the protocols and algorithms for identification, authentication, and encryption to be used for setting up VPN tunnels in IKEv1 Phase 1. See [Set up an IKE Gateway](#) and [Define IKE Crypto Profiles](#).
- ❑ Configure the parameters that are needed to establish the IPSec connection for transfer of data across the VPN tunnel; See [Set up an IPSec Tunnel](#). For IKEv1 Phase-2, see [Define IPSec Crypto Profiles](#).
- ❑ (Optional) Specify how the firewall will monitor the IPSec tunnels. See [Set up Tunnel Monitoring](#).
- ❑ Define security policies to filter and inspect the traffic.



If there is a deny rule at the end of the security rulebase, intra-zone traffic is blocked unless otherwise allowed. Rules to allow IKE and IPSec applications must be explicitly included above the deny rule.



If your VPN traffic is passing through (not originating or terminating on) a PA-7000 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](#).


Set up an IKE Gateway


To set up a VPN tunnel, the VPN peers or gateways must authenticate each other using preshared keys or digital certificates and establish a secure channel in which to negotiate the IPSec security association (SA) that will be used to secure traffic between the hosts on each side.

Set up an IKE Gateway

Step 1 Define the new [IKE Gateway](#).

1. Select **Network > Network Profiles > IKE Gateway** and enter a **Name** for the new gateway configuration.
2. Select the outgoing **Interface** on the firewall.
3. From the **Local IP address** drop-down list, select the IP address that will be used to as the endpoint for the VPN connection. This is the external-facing interface with a publicly routable IP address on the firewall.

Set up an IKE Gateway	
<p>Step 2 Define the settings for the peer at the far end of the tunnel.</p>	<ol style="list-style-type: none"> 1. Select whether the peer uses a Static or Dynamic IP address in Peer IP Type. 2. If the Peer IP Address is static, enter the IP address of the peer.
<p>Step 3 Select the peer authentication method. This is required for static and dynamic peers.</p>	<ul style="list-style-type: none"> • For configuring a Pre-Shared Key, see Step 4. • For configuring digital Certificates, see Step 5.
<p>Step 4 Configure a pre-shared key.</p>	<ol style="list-style-type: none"> 1. Enter a security key to use for authentication across the tunnel. This key must be the same on both peers. Generate a key that is hard to crack with dictionary attacks; use a pre-shared key generator, if necessary. Use a maximum of 255 ASCII or non-ASCII characters. 2. Continue with Step 6.
<p>Step 5 Configure certificate-based authentication.</p> <p> The pre-requisites for certificate-based authentication are as follows:</p> <ul style="list-style-type: none"> – Obtain a signed certificate: See Generate a Certificate on the Firewall or Obtain a Certificate from an External CA. – Configure the certificate profile: The certificate profile provides the settings that the IKE gateway uses for negotiating and validating certificate authentication with its peer. See Configure a Certificate Profile. 	<ol style="list-style-type: none"> 1. Select Certificate for the Authentication method and select the signed certificate from the Local Certificate drop-down. If your device is enabled for multi-virtual systems, if the certificate belongs to a virtual system, it must be in the same virtual system as the interface used for the IKE gateway. 2. From the Local Identification drop-down list, choose one of the following types and enter the value: IP address, FQDN (hostname), User FQDN (email address), Distinguished Name (subject). 3. From the Peer Identification drop-down, choose one of the following types and enter the value: IP address, FQDN (hostname), User FQDN (email address), Distinguished Name (subject). 4. Select a Certificate Profile to use. 5. Continue with Step 6.

Set up an IKE Gateway	
<p>Step 6 Configure the additional parameters for IKE phase 1 negotiations—Exchange mode, Crypto profile, IKE fragmentation, Dead Peer Detection.</p>	<ol style="list-style-type: none"> 1. Select Network > Network Profiles > IKE Gateways and select the Advanced Phase 1 Options tab. 2. Choose auto, aggressive, or main for the Exchange Mode. When a device is set to use the auto exchange mode, it can accept both main mode and aggressive mode negotiation requests; however, whenever possible, it initiates negotiation and allows exchanges in main mode.  If the exchange mode is not set to auto, you must configure both VPN peers with the same exchange mode to allow each peer to accept negotiation requests. 3. Select an existing profile or keep the default profile from IKE Crypto Profile drop-down. For details on defining an IKE Crypto profile, see Define IKE Crypto Profiles. 4. Select Passive Mode if you want the firewall to only respond to IKE connections and never initiate them. 5. Select NAT Traversal Select to have UDP encapsulation used on IKE and UDP protocols, enabling them to pass through intermediate NAT devices. 6. (Only if using certificate-based authentication and the exchange mode is not set as aggressive mode) Select Enable Fragmentation to enable the firewall to operate with IKE Fragmentation. 7. Select the Dead Peer Detection check box and enter an Interval (2 - 100 seconds); For Retry, define the time to delay (2 - 100 seconds) before attempting to re-check availability. 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.
<p>Step 7 Save the changes.</p>	<p>Click OK and Commit.</p>

Define Cryptographic Profiles

A cryptographic profile specifies the ciphers used for authentication and/or encryption between two IKE peers, and the lifetime of the key. The time period between each renegotiation is known as the lifetime; when the specified time expires, the firewall renegotiates a new set of keys.

For securing communication across the VPN tunnel, the firewall requires IKE and IPSec cryptographic profiles for completing IKE phase 1 and phase 2 negotiations, respectively. The firewall includes a *default* IKE crypto profile and a *default* IPSec crypto profile that is ready for use.

▲ [Define IKE Crypto Profiles](#)

▲ [Define IPSec Crypto Profiles](#)

Define IKE Crypto Profiles

The IKE crypto profile is used to set up the encryption and authentication algorithms used for the key exchange process in [IKE Phase 1](#), and lifetime of the keys, which specifies how long the keys are valid. To invoke the profile, you must attach it to the IKE Gateway configuration.



All IKE gateways configured on the same interface or local IP address must use the same crypto profile.

Define an IKE Crypto Profile	
Step 1 Create a new IKE profile.	<ol style="list-style-type: none"> 1. Select Network > Network Profiles > IKE Crypto and select Add. 2. Enter a Name for the new profile.
Step 2 Select the DH Group to use for setting up the key exchange.	<p>Click Add and select the key strength that you want to use for the DH Group.</p> <p>Select more than one DH group if you are not certain of what is supported by the VPN peer. Prioritize the list of ciphers by strength so that the strongest cipher is used for setting up the tunnel.</p>
Step 3 Select the authentication and encryption algorithm.	<p>Click Add and select the Authentication and Encryption algorithms that you want to use for communication between the IKE peers.</p> <p>Selecting multiple algorithms allows the peers to use the strongest cipher/algorithm that is supported on both IKE peers.</p>
Step 4 Specify the duration for which the key is valid.	<p>Select the Lifetime for which the key is valid. The time period between each renegotiation is known as the lifetime; when the specified time expires, the firewall will renegotiate a new set of keys.</p>
Step 5 Save your IKE Crypto profile.	<p>Click OK and click Commit.</p>
Step 6 Attach the IKE Crypto profile to the IKE Gateway configuration.	<p>See Step 6 in Set up an IKE Gateway.</p>

Define IPSec Crypto Profiles

The IPSec crypto profile is invoked in [IKE Phase 2](#). It specifies how the data is secured within the tunnel when Auto Key IKE is used to automatically generate keys for the IKE SAs.

Define the IPSec Crypto Profile	
<p>Step 1 Create a new IPSec profile.</p>	<ol style="list-style-type: none"> 1. Select Network > Network Profiles > IPSec Crypto and select Add. 2. Enter a Name for the new profile. 3. Select the IPSec Protocol—ESP or AH—that you want to apply to secure the data as it traverses across the tunnel. 4. Click Add and select the Authentication and Encryption algorithms for ESP, and Authentication algorithms for AH, so that the IKE peers can negotiate the keys for the secure transfer of data across the tunnel. Selecting multiple algorithms allows the peers to use the strongest cipher/algorithm that is supported on both IKE peers.
<p>Step 2 Select the DH Group to use for the IPSec SA negotiations in IKE phase 2.</p>	<ol style="list-style-type: none"> 1. Select the key strength that you want to use from the DH Group drop-down. Select more than one DH group if you are not certain of what key strength is supported by the peer at the other end. The strongest cipher will be used for setting up the tunnel. 2. Select no-pfs, if you do not want to renew the key that was created at phase 1, the current key is reused for the IPSEC SA negotiations.
<p>Step 3 Specify the duration of the key—time and volume of traffic.</p>	<p>Using a combination of time and traffic volume allows you to ensure safety of data.</p> <p>Select the Lifetime or time period for which the key is valid. When the specified time expires, the firewall will renegotiate a new set of keys.</p> <p>Select the Lifesize or volume of data after which the keys must be renegotiated.</p>
<p>Step 4 Save your IPSec profile.</p>	<p>Click OK and click Commit.</p>
<p>Step 5 Attach the IPSec Profile to an IPSec tunnel configuration.</p>	<p>See Step 4 in</p>

Set up an IPSec Tunnel

The IPSec tunnel configuration allows you to authenticate and/or encrypt the data (IP packet) as it traverses across the tunnel.

If you are setting up the Palo Alto Networks firewall to work with a peer that supports policy-based VPN, you must define Proxy IDs. Devices that support policy-based VPN use specific security rules/policies or access-lists (source addresses, destination addresses and ports) for permitting interesting traffic through an IPSec tunnel. These rules are referenced during quick mode/IKE phase 2 negotiation, and are exchanged as Proxy-IDs in the first or the second message of the process. So, if you are configuring the Palo Alto Networks firewall to work with a policy-based VPN peer, for a successful phase 2 negotiation you must define the

Proxy-ID so that the setting on both peers is identical. If the Proxy-ID is not configured, because the Palo Alto Networks firewall supports route-based VPN, the default values used as Proxy-ID are source ip: 0.0.0.0/0, destination ip: 0.0.0.0/0 and application: any; and when these values are exchanged with the peer, it results in a failure to set up the VPN connection.

Set up an IPSec Tunnel

Step 1 Select **Network > IPSec Tunnels > General** and enter a **Name** for the new tunnel.


Step 2 Select the **Tunnel interface** that will be used to set up the IPSec tunnel.

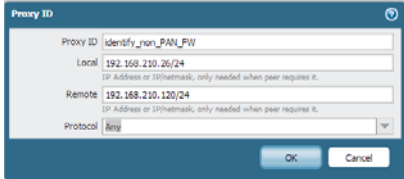
– To create a new tunnel interface:

1. Select **Network > Interfaces > Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
 - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down. Associating the tunnel interface with the same zone (and virtual router) as the external-facing interface on which the packets enter the firewall, mitigates the need to create inter-zone routing.
 - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example vpn-corp), and click **OK**.
4. In the **Virtual Router** drop-down, select **default**.
5. (Optional) If you want to assign an IPv4 address to the tunnel interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.31.32.1/32.
6. If you want to assign an IPv6 address to the tunnel interface, see [Step 3](#).
7. To save the interface configuration, click **OK**.

Step 3 (Optional) Enable IPv6 on the tunnel interface.

1. Select the IPv6 tab on **Network > Interfaces > Tunnel > IPv6**.
2. Select the check box to **Enable IPv6 on the interface**.
This option allows you to route IPv6 traffic over an IPv4 IPSec tunnel and will provide confidentiality between IPv6 networks. The IPv6 traffic is encapsulated by IPv4 and then ESP. To route IPv6 traffic to the tunnel, you can use a static route to the tunnel, or use OSPFv3, or use a Policy-Based Forwarding (PBF) rule to direct traffic to the tunnel.
3. Enter the 64-bit extended unique **Interface ID** in hexadecimal format, for example, 00:26:08:FF:FE:DE:4E:29. By default, the firewall will use the EUI-64 generated from the physical interface's MAC address.
4. To enter an IPv6 **Address**, click **Add** and enter an IPv6 address and prefix length, for example 2001:400:f00::1/64. If Prefix is not selected, the IPv6 address assigned to the interface will be wholly specified in the address text box.
 - a. Select **Use interface ID as host portion** to assign an IPv6 address to the interface that will use the interface ID as the host portion of the address.
 - b. Select **Anycast** to include routing through the nearest node.

Set up an IPSec Tunnel	
<p>Step 4 Select the type of key that will be used to secure the IPSec tunnel.</p>	<p>Continue to one of the following steps, depending on why type of key exchange you are using:</p> <ul style="list-style-type: none"> • Set up Auto Key exchange. • Set up Auto Key exchange.
<ul style="list-style-type: none"> • Set up Auto Key exchange. 	<ol style="list-style-type: none"> 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.
<ul style="list-style-type: none"> • Set up a Manual Key exchange. 	<ol style="list-style-type: none"> 1. Set up the parameters for the local firewall: <ol style="list-style-type: none"> a. Specify the SPI for the local firewall. SPI is a 32-bit hexadecimal index that is added to the header for IPSec tunneling to assist in differentiating between IPSec traffic flows; it is used to create the SA required for establishing a VPN tunnel. b. Select the Interface that will be the tunnel endpoint, and optionally select the IP address for the local interface that is the endpoint of the tunnel. c. Select the protocol to be used—AH or ESP. d. For AH, select the Authentication method from the drop-down and enter a Key and then Confirm Key. e. For ESP, select the Authentication method from the drop-down and enter a Key and then Confirm Key. Then, select the Encryption method and enter a Key and then Confirm Key, if needed. 2. Set up the parameters that pertain to the remote VPN peer. <ol style="list-style-type: none"> a. Specify the SPI for the remote peer. b. Enter the Remote Address, the IP address of the remote peer.
<p>Step 5 Protect against a replay attack.</p> <p>A replay attack occurs when a packet is maliciously intercepted and retransmitted by the interceptor.</p>	<p>Select the Show Advanced Options check box, select Enable Replay Protection to detect and neutralize against replay attacks.</p>
<p>Step 6 Preserve the Type of Service header for the priority or treatment of IP packets.</p>	<p>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.</p>
<p>Step 7 Enable Tunnel Monitoring.</p> <p> You need to assign an IP address to the tunnel interface for monitoring.</p>	<p>To alert the device administrator to tunnel failures and to provide automatic failover to another tunnel interface:</p> <ol style="list-style-type: none"> 1. Specify a Destination IP address on the other side of the tunnel to determine if the tunnel is working properly. 2. Select a Profile to determine the action on tunnel failure. To create a new profile, see Define a Tunnel Monitoring Profile.

Set up an IPsec Tunnel	
<p>Step 8 (Required only if the VPN peer uses policy-based VPN). Create a Proxy ID to identify the VPN peers.</p>	<ol style="list-style-type: none"> 1. Select Network > IPsec Tunnels > ProxyID. 2. Click Add and enter the IP address for the VPN gateway peers. 
<p>Step 9 Save your changes</p>	<p>Click OK and Commit.</p>

Set up Tunnel Monitoring

To provide uninterrupted VPN service, you can use the Dead Peer Detection capability along with the tunnel monitoring capability on the firewall. You can also monitor the status of the tunnel. These monitoring tasks are described in the following sections:

- ▲ [Define a Tunnel Monitoring Profile](#)
- ▲ [View the Status of the Tunnels](#)

Define a Tunnel Monitoring Profile

A tunnel monitoring profile allows you to verify connectivity between the VPN peers; you can configure the tunnel interface to ping a destination IP address at a specified interval and specify the action if the communication across the tunnel is broken.

Define a Tunnel Monitoring Profile	
Step 1	Select Network > Network Profiles > Monitor . A default tunnel monitoring profile is available for use.
Step 2	Click Add , and enter a Name for the profile.
Step 3	<p>Select the Action if the destination IP address is unreachable.</p> <ul style="list-style-type: none"> • 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. <p>In either case, the firewall attempts to accelerate the recovery by negotiating new IPsec keys.</p>
Step 4	<p>Specify the Interval and Threshold to trigger the specified action.</p> <p>The threshold specifies the number of heartbeats to wait before taking the specified action. The range is 2-100 and the default is 5.</p> <p>The Interval measures the time between heartbeats. The range is 2-10 and the default is 3 seconds.</p>
Step 5	Attach the monitoring profile to the IPsec Tunnel configuration. See Enable Tunnel Monitoring .

View the Status of the Tunnels

The status of the tunnel informs you about whether or not valid IKE phase-1 and phase-2 SAs have been established, and whether the tunnel interface is up and available for passing traffic.

Because the tunnel interface is a logical interface, it cannot indicate a physical link status. Therefore, you must enable tunnel monitoring so that the tunnel interface can verify connectivity to an IP address and determine if the path is still usable. If the IP address is unreachable, the firewall will either wait for the tunnel to recover or failover. When a failover occurs, the existing tunnel is torn down and routing changes are triggered to set up a new tunnel and redirect traffic.

View Tunnel Status

1. Select **Network > IPSec Tunnels**.
 2. View the **Tunnel Status**.
 - Green indicates a valid IPSec SA tunnel.
 - Red indicates that IPSec SA is not available or has expired.
 3. View the **IKE Gateway Status**.
 - Green indicates a valid IKE phase-1 SA.
 - Red indicates that IKE phase-1 SA is not available or has expired.
 4. View the **Tunnel Interface Status**.
 - Green indicates that the tunnel interface is up.
 - Red indicates that the tunnel interface is down, because tunnel monitoring is enabled and the status is down.
-

To troubleshoot a VPN tunnel that is not yet up, see [Interpret VPN Error Messages](#).

Test VPN Connectivity

Test Connectivity

- Initiate IKE phase 1 by either pinging a host across the tunnel or using the following CLI command:

```
test vpn ike-sa gateway gateway_name
```

- Then enter the following command to test if IKE phase 1 is set up:

```
show vpn ike-sa gateway gateway_name
```

In the output, check if the Security Association displays. If it does not, review the system log messages to interpret the reason for failure.

- Initiate IKE phase 2 by either pinging a host from across the tunnel or using the following CLI command:

```
test vpn ipsec-sa tunnel tunnel_name
```

- Then enter the following command to test if IKE phase 1 is set up:

```
show vpn ipsec-sa tunnel tunnel_name
```

In the output, check if the Security Association displays. If it does not, review the system log messages to interpret the reason for failure.

- To view the VPN traffic flow information, use the following command:

```
show vpn-flow
```

```
admin@PA-500> show vpn flow
```

```
total tunnels configured:      1
filter - type IPSec, state any
```

```
total IPSec tunnel configured:  1
total IPSec tunnel shown:      1
```

name	id	state	local-ip	peer-ip	tunnel-i/f
vpn-to-siteB	5	active	100.1.1.1	200.1.1.1	tunnel.41

Interpret VPN Error Messages

The following table lists some of the common VPN error messages that are logged in the system log.

Table: Syslog Error Messages for VPN Issues

If error is this:	Try this:
<pre>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]</pre>	<ul style="list-style-type: none"> • 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.
<pre>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.</pre>	<p>Check the IKE Crypto profile configuration to verify that the proposals on both sides have a common encryption, authentication, and DH Group proposal.</p>
<pre>pfs group mismatched:my: 2peer: 0 or IKE phase-2 negotiation failed when processing SA payload. No suitable proposal found in peer's SA payload.</pre>	<p>Check the IPSec Crypto profile configuration to verify that:</p> <ul style="list-style-type: none"> • 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
<pre>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.</pre>	<p>The VPN peer on one end is using policy-based VPN. You must configure a Proxy ID on the Palo Alto Networks firewall. See Step 8.</p>

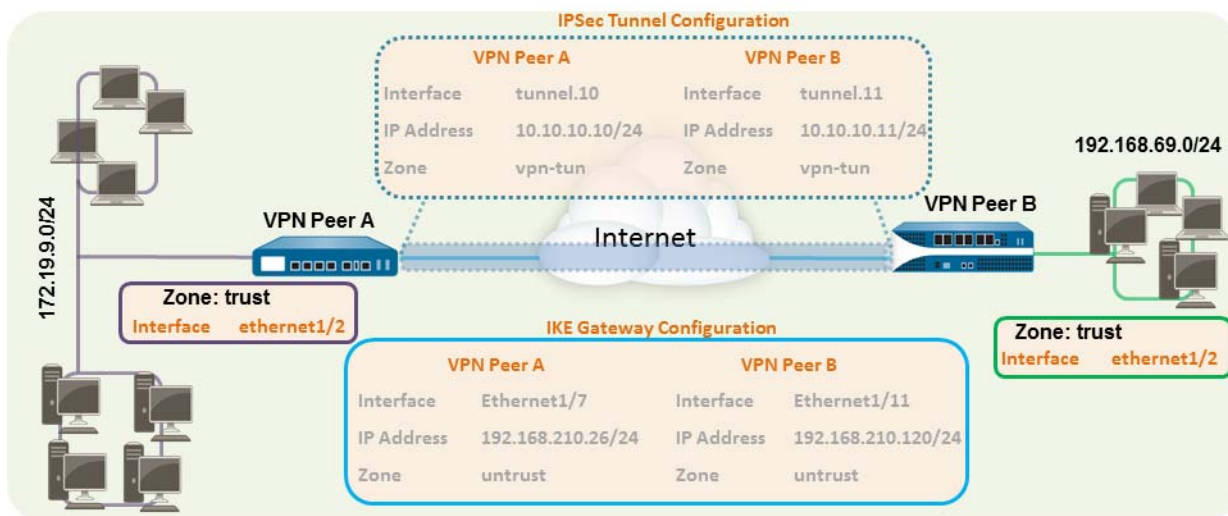
Site-to-Site VPN Quick Configs

The following sections provide instructions for configuring some common VPN deployments:

- ▲ Site-to-Site VPN with Static Routing
- ▲ Site-to-Site VPN with OSPF
- ▲ Site-to-Site VPN with Static and Dynamic Routing

Site-to-Site VPN with Static Routing

The following example shows a VPN connection between two sites that use static routes. Without dynamic routing, the tunnel interfaces on VPN Peer A and VPN Peer B do not require an IP address because the firewall automatically uses the tunnel interface as the next hop for routing traffic across the sites. However, to enable tunnel monitoring, a static IP address has been assigned to each tunnel interface.



Quick Config: Site-to-Site VPN with Static Routing

Step 1 Configure a Layer 3 interface.
This interface is used for the IKE phase-1 tunnel.

1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for VPN.
2. Select **Layer3** from the **Interface Type** drop-down.
3. On the **Config** tab, select the **Security Zone** to which the interface belongs:
 - The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.
 - If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.
4. Select the **Virtual Router** to use.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.
6. To save the interface configuration, click **OK**.
In this example, the configuration for VPN Peer A is:
 - **Interface**—ethernet1/7
 - **Security Zone**—untrust
 - **Virtual Router**—default
 - **IPv4**—192.168.210.26/24
 The configuration for VPN Peer B is:
 - **Interface**—ethernet1/11
 - **Security Zone**—untrust
 - **Virtual Router**—default
 - **IPv4**—192.168.210.120/24

Quick Config: Site-to-Site VPN with Static Routing

Step 2 Create a tunnel interface and attach it to a virtual router and security zone.

Interface	Management Profile	IP Address	Virtual Router	Security Zone
tunnel	none	none	none	none
tunnel.11		172.19.9.2	default	vpn_tun

1. Select **Network > Interfaces > Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.1**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
 - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.
 - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the **Zone** dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.
4. Select the **Virtual Router**.
5. (Optional) Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface.

With static routes, the tunnel interface does not require an IP address. For traffic that is destined to a specified subnet/IP address, the tunnel interface will automatically become the next hop. Consider adding an IP address if you want to enable tunnel monitoring.

6. To save the interface configuration, click **OK**.

In this example, the configuration for VPN Peer A is:

- **Interface**—tunnel.11
- **Security Zone**—vpn_tun
- **Virtual Router**—default
- **IPv4**—172.19.9.2/24

The configuration for VPN Peer B is:

- **Interface**—tunnel.12
- **Security Zone**—vpn_tun
- **Virtual Router**—default
- **IPv4**—192.168.69.2/24

Step 3 Configure a static route, on the virtual router, to the destination subnet.

Virtual Router: default					
Static Routes					
Name	Destination	Interface	Type	Value	Admin Distance
traffic to 192.168.69.0	192.168.69.0	tunnel.11			default

1. Select **Network > Virtual Router** and click the router you defined in [step 4](#) above.
2. Select **Static Route**, click **Add**, and enter a new route to access the subnet that is at the other end of the tunnel.

In this example, the configuration for VPN Peer A is:

- **Destination**—192.168.69.0/24
- **Interface**—tunnel.11

The configuration for VPN Peer B is:

- **Destination**—172.19.9.0/24
- **Interface**—tunnel.12

Quick Config: Site-to-Site VPN with Static Routing

Step 4 Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).

Complete this task on both peers and make sure to set identical values.

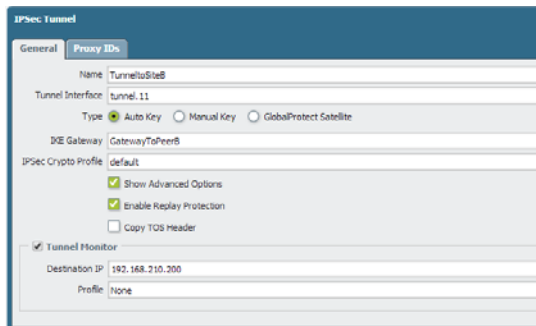
1. Select **Network > Network Profiles > IKE Crypto**. In this example, we use the default profile.

Name	Encryption	Authentication	DH Group	Lifetime
default	aes128, 3des	sha1	group2	8 hours

2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.

Name	ESP/AH	Encryption	Authentication	DH Group	Lifetime	Lifetsize
default	ESP	aes128, 3des	sha1	group2	1 hours	

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.

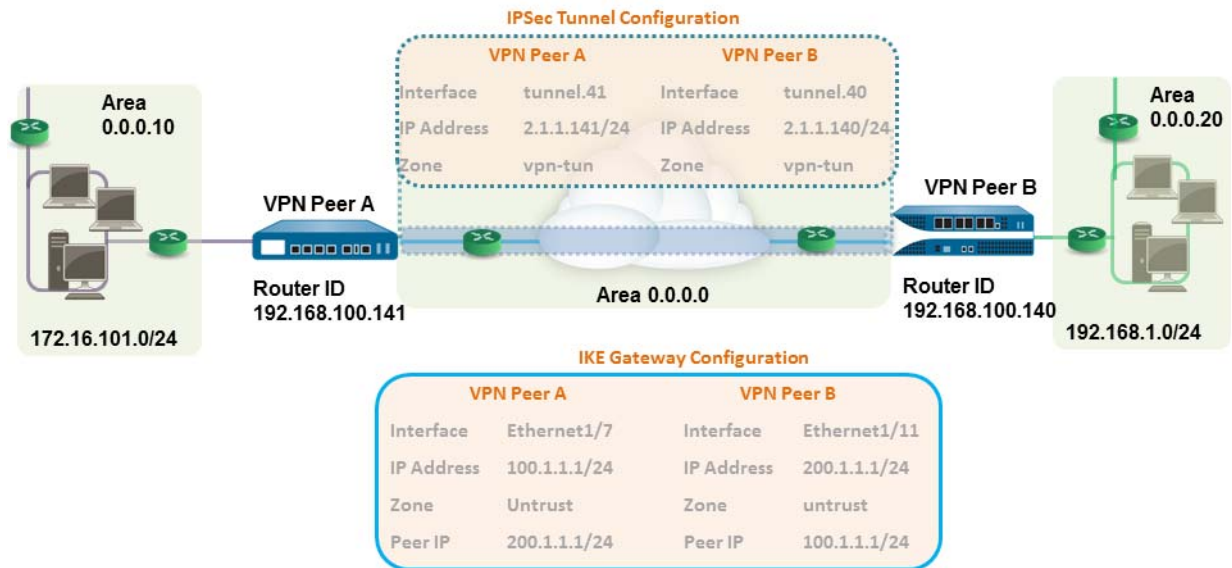
The configuration for VPN Peer B is:

- **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.

Quick Config: Site-to-Site VPN with Static Routing	
<p>Step 6 Set up the IPSec Tunnel.</p>	<ol style="list-style-type: none"> 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: <ul style="list-style-type: none"> • Tunnel Interface—tunnel.11 • Type—Auto Key • IKE Gateway—Select the IKE Gateway defined above. • IPSec Crypto Profile—Select the IPSec Crypto profile defined in Step 4. The configuration for VPN Peer B is: <ul style="list-style-type: none"> • Tunnel Interface—tunnel.12 • Type—Auto Key • IKE Gateway—Select the IKE Gateway defined above. • IPSec Crypto Profile—Select the IPSec Crypto defined in Step 4. 3. (Optional) Select Show Advanced Options, select Tunnel Monitor, and specify a Destination IP address to ping for verifying connectivity. Typically, the tunnel interface IP address for the VPN Peer is used. 4. (Optional) To define the action on failure to establish connectivity, see Define a Tunnel Monitoring Profile.
<p>Step 7 Create policies to allow traffic between the sites (subnets).</p>	<ol style="list-style-type: none"> 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.
<p>Step 8 Save any pending configuration changes.</p>	Click Commit .
<p>Step 9 Test VPN connectivity.</p>	See View the Status of the Tunnels .

Site-to-Site VPN with OSPF

In this example, each site uses OSPF for dynamic routing of traffic. The tunnel IP address on each VPN peer is statically assigned and serves as the next hop for routing traffic between the two sites.



Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF

Step 1 Configure the Layer 3 interfaces on each firewall.

1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for VPN.
2. Select **Layer3** from the **Interface Type** drop-down.
3. On the **Config** tab, select the **Security Zone** to which the interface belongs:
 - The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.
 - If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.
4. Select the **Virtual Router** to use.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.
6. To save the interface configuration, click **OK**.

In this example, the configuration for VPN Peer A is:

- **Interface**—ethernet1/7
- **Security Zone**—untrust
- **Virtual Router**—default
- **IPv4**—100.1.1.1/24

The configuration for VPN Peer B is:

- **Interface**—ethernet1/11
- **Security Zone**—untrust
- **Virtual Router**—default
- **IPv4**—200.1.1.1/24

Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF

Step 2 Create a tunnel interface and attach it to a virtual router and security zone.

Interface	Management Profile	IP Address	Virtual Router	Security Zone
tunnel	none	2.1.1.141/24	default	vpn_tun
tunnel.11				


1. Select **Network > Interfaces > Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, say, **.11**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
 - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.
 - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the **Zone** dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.
4. Select the **Virtual Router**.
5. Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask/prefix to assign to the interface, for example, 172.19.9.2/24.

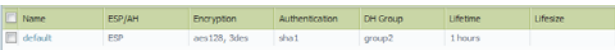
This IP address will be used as the next hop IP address to route traffic to the tunnel and can also be used to monitor the status of the tunnel.
6. To save the interface configuration, click **OK**.

In this example, the configuration for VPN Peer A is:
 - **Interface**—tunnel.41
 - **Security Zone**—vpn_tun
 - **Virtual Router**—default
 - **IPv4**—2.1.1.141/24
 The configuration for VPN Peer B is:
 - **Interface**—tunnel.40
 - **Security Zone**—vpn_tun
 - **Virtual Router**—default
 - **IPv4**—2.1.1.140/24

Step 3 Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).

Complete this task on both peers and make sure to set identical values.

1. Select **Network > Network Profiles > IKE Crypto**. In this example, we use the default profile.
 

Name	Encryption	Authentication	DH Group	Lifetime
default	aes128, 3des	sha1	group2	8 hours
2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.
 

Name	ESP/AH	Encryption	Authentication	DH Group	Lifetime	Lifesize
default	ESP	aes128, 3des	sha1	group2	1 hours	

Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF

<p>Step 4 Set up the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall.</p> <p>For more information on the OSPF options that are available on the firewall, see Configure OSPF.</p> <p>Use Broadcast as the link type when there are more than two OSPF routers that need to exchange routing information.</p>	<ol style="list-style-type: none"> 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: <ul style="list-style-type: none"> – 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 <p>The OSPF configuration for VPN Peer B is:</p> <ul style="list-style-type: none"> – 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
<p>Step 5 Set up the IKE Gateway.</p> <p>This examples uses static IP addresses for both VPN peers. Typically, the corporate office uses a statically configured IP address, and the branch side can be a dynamic IP address; dynamic IP addresses are not best suited for configuring stable services such as VPN.</p>	<ol style="list-style-type: none"> 1. Select Network > Network Profiles > IKE Gateway. 2. Click Add and configure the options in the General tab. <p>In this example, the configuration for VPN Peer A is:</p> <ul style="list-style-type: none"> • Interface—ethernet1/7 • Local IP address—100.1.1.1/24 • Peer IP address—200.1.1.1/24 • Preshared keys—enter a value <p>The configuration for VPN Peer B is:</p> <ul style="list-style-type: none"> • 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.

Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF

<p>Step 6 Set up the IPSec Tunnel.</p>	<ol style="list-style-type: none"> 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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.
<p>Step 7 Create policies to allow traffic between the sites (subnets).</p>	<ol style="list-style-type: none"> 1. Select Policies > Security. 2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses.

Quick Config: Site-to-Site VPN with Dynamic Routing using OSPF

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: 0x0:reserved, 0:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
M/P:NSSA option, MC:multicast, EA: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: 0x0:reserved, 0:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
M/P:NSSA option, MC:multicast, EA: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
Oi:ospf intra-area, Oo:ospf inter-area, O1:ospf ext-type-1, O2:ospf ext-type-2
```

```
VIRTUAL ROUTER: vr1 (id 1)
=====
destination      nexthop          metric flags   age   interface
2.1.1.0/24       0.0.0.0          10   Oi         6760  tunnel.41
172.16.101.0/24 0.0.0.0          10   Oi         6854  ethernet1/1
192.168.1.0/24   2.1.1.140       20   A Co       6754  tunnel.40
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, O1:ospf ext-type-1, O2:ospf ext-type-2
```

```
VIRTUAL ROUTER: vr1 (id 1)
=====
destination      nexthop          metric flags   age   interface
2.1.1.0/24       0.0.0.0          10   Oi         20033  tunnel.40
172.16.101.0/24 2.1.1.141       20   A Oo       6896  tunnel.40
192.168.1.0/24   0.0.0.0          10   Oi         8058  ethernet1/15
total routes shown: 3
```

Step 9 Test VPN connectivity.

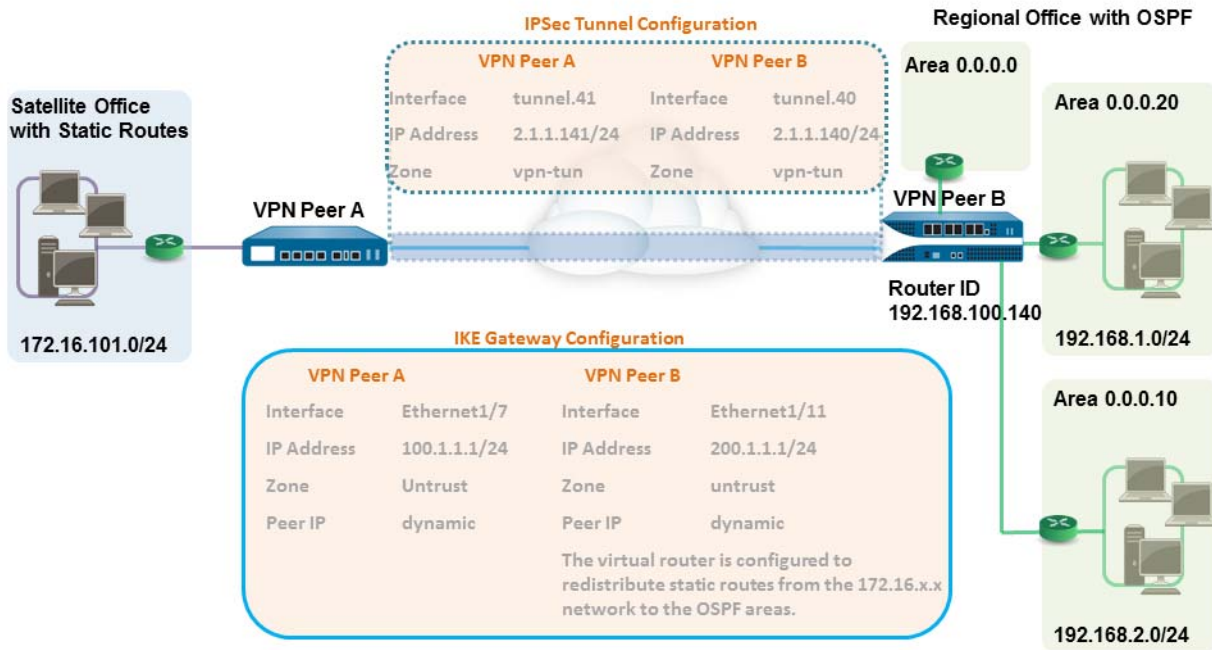
See [Set up Tunnel Monitoring](#) and [View the Status of the Tunnels](#).

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.



Quick Config: Site-to-Site VPN with Static and Dynamic Routing

Step 1 Configure the Layer 3 interfaces on each firewall.

1. Select **Network > Interfaces > Ethernet** and then select the interface you want to configure for VPN.
2. Select **Layer3** from the **Interface Type** drop-down.
3. On the **Config** tab, select the **Security Zone** to which the interface belongs:
 - The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic.
 - If you have not yet created the zone, select **New Zone** from the **Security Zone** drop-down, define a **Name** for the new zone and then click **OK**.
4. Select the **Virtual Router** to use.
5. To assign an IP address to the interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.210.26/24.
6. To save the interface configuration, click **OK**.

In this example, the configuration for VPN Peer A is:

- **Interface**—ethernet1/7
- **Security Zone**—untrust
- **Virtual Router**—default
- **IPv4**—100.1.1.1/24

The configuration for VPN Peer B is:

- **Interface**—ethernet1/11
- **Security Zone**—untrust
- **Virtual Router**—default
- **IPv4**—200.1.1.1/24

Step 2 Set up the Crypto profiles (IKE Crypto profile for phase 1 and IPSec Crypto profile for phase 2).

Complete this task on both peers and make sure to set identical values.

1. Select **Network > Network Profiles > IKE Crypto**. In this example, we use the default profile.

Name	Encryption	Authentication	DH Group	Lifetime
default	aes128, 3des	sha1	group2	8 hours

2. Select **Network > Network Profiles > IPSec Crypto**. In this example, we use the default profile.

Name	ESP/AH	Encryption	Authentication	DH Group	Lifetime	Lifesize
default	ESP	aes128, 3des	sha1	group2	1 hours	

Quick Config: Site-to-Site VPN with Static and Dynamic Routing

Step 3 Set up the IKE Gateway.

With pre-shared keys, to add authentication scrutiny when setting up the IKE phase-1 tunnel, you can set up Local and Peer Identification attributes and a corresponding value that is matched in the IKE negotiation process.

1. Select **Network > Network Profiles > IKE Gateway**.
2. Click **Add** and configure the options in the **General** tab.
 In this example, the configuration for VPN Peer A is:
 - **Interface**—ethernet1/7
 - **Local IP address**—100.1.1.1/24
 - **Peer IP type**—dynamic
 - **Preshared keys**—enter a value
 - **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer A.
 - **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer B
 The configuration for VPN Peer B is:
 - **Interface**—ethernet1/11
 - **Local IP address**—200.1.1.1/24
 - **Peer IP address**—dynamic
 - **Preshared keys**—enter same value as on Peer A
 - **Local identification**—select **FQDN(hostname)** and enter the value for VPN Peer B
 - **Peer identification**—select **FQDN(hostname)** and enter the value for VPN Peer A
3. Select the IKE Crypto profile you created earlier to use for IKE phase 1.

Quick Config: Site-to-Site VPN with Static and Dynamic Routing

Step 4 Create a tunnel interface and attach it to a virtual router and security zone.

Interface	Management Profile	IP Address	Virtual Router	Security Zone
tunnel	none	none	none	none
tunnel.11		2.1.1.141/24	default	vpn_tun

1. Select **Network > Interfaces > Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, say, **.41**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
 - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.
 - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the **Zone** dialog, define a **Name** for new zone (for example *vpn-tun*), and then click **OK**.
4. Select the **Virtual Router**.
5. Assign an IP address to the tunnel interface, select the **IPv4** or **IPv6** tab, click **Add** in the IP section, and enter the IP address and network mask/prefix to assign to the interface, for example, 172.19.9.2/24.

This IP address will be used to route traffic to the tunnel and to monitor the status of the tunnel.

6. To save the interface configuration, click **OK**.
In this example, the configuration for VPN Peer A is:

- **Interface**—tunnel.41
- **Security Zone**—vpn_tun
- **Virtual Router**—default
- **IPv4**—2.1.1.141/24

The configuration for VPN Peer B is:

- **Interface**—tunnel.42
- **Security Zone**—vpn_tun
- **Virtual Router**—default
- **IPv4**—2.1.1.140/24

Step 5 Specify the interface to route traffic to a destination on the 192.168.x.x network.

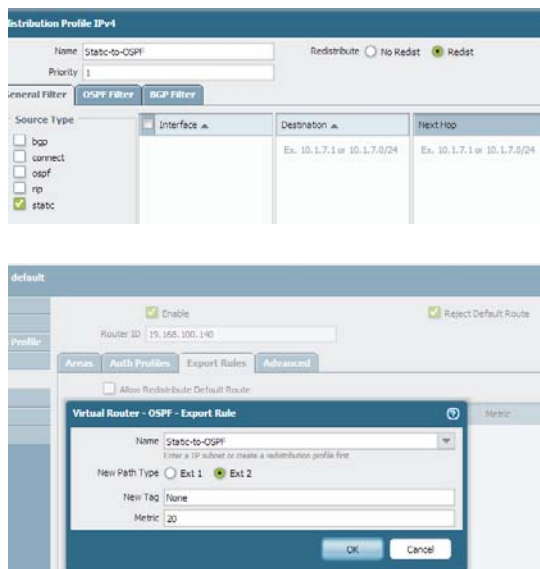
1. On VPN Peer A, select the virtual router.
2. Select **Static Routes**, and **Add** tunnel.41 as the **Interface** for routing traffic with a **Destination** in the 192.168.x.x network.

Quick Config: Site-to-Site VPN with Static and Dynamic Routing

Step 6 Set up the static route and the OSPF configuration on the virtual router and attach the OSPF areas with the appropriate interfaces on the firewall.

1. On VPN Peer B, select **Network > Virtual Routers**, and select the default router or add a new router.
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.
 - a. Select **Network > Virtual Routers**, and select the router you used above.
 - b. Select **Redistribution Profiles**, and click **Add**.
 - c. Enter a Name for the profile and select **Redist** and assign a **Priority** value. If you have configured multiple profiles, the profile with the lowest priority value is matched first.
 - d. Set **Source Type** as **static**, and click **OK**. The static route defined in **Step 6-2** will be used for the redistribution.
2. Inject the static routes in to the OSPF system.
 - a. Select **OSPF> Export Rules** (for IPv4) or **OSPFv3> Export Rules** (for IPv6).
 - b. Click **Add**, and select the redistribution profile that you just created.
 - c. Select how the external routes are brought into the OSPF system. The default option, **Ext2** calculates the total cost of the route using only the external metrics. To use both internal and external OSPF metrics, use **Ext1**.
 - d. Assign a **Metric** (cost value) for the routes injected into the OSPF system. This option allows you to change the metric for the injected route as it comes into the OSPF system.
 - e. Click **OK** to save the changes.

Quick Config: Site-to-Site VPN with Static and Dynamic Routing

Step 8 Set up the IPSec Tunnel.

1. Select **Network > IPSec Tunnels**.
2. Click **Add** and configure the options in the **General** tab.
In this example, the configuration for VPN Peer A is:
 - **Tunnel Interface**—tunnel.41
 - **Type**—Auto Key
 - **IKE Gateway**—Select the IKE Gateway defined above.
 - **IPSec Crypto Profile**—Select the IKE Gateway defined above.
 The configuration for VPN Peer B is:
 - **Tunnel Interface**—tunnel.40
 - **Type**—Auto Key
 - **IKE Gateway**—Select the IKE Gateway defined above.
 - **IPSec Crypto Profile**—Select the IKE Gateway defined above.
3. Select **Show Advanced Options**, select **Tunnel Monitor**, and specify a Destination IP address to ping for verifying connectivity.
4. To define the action on failure to establish connectivity, see [Define a Tunnel Monitoring Profile](#).

Step 9 Create policies to allow traffic between the sites (subnets).

1. Select **Policies > Security**.
 2. Create rules to allow traffic between the untrust and the vpn-tun zone and the vpn-tun and the untrust zone for traffic originating from specified source and destination IP addresses.
-

Quick Config: Site-to-Site VPN with Static and Dynamic Routing

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: 0x0:reserved, 0:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
M/P:NSSA option, MC:multicast, EA: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: 0x0:reserved, 0:Opaq-LSA capability, DC:demand circuits, EA:Ext-Attr LSA capability,
M/P:NSSA option, MC:multicast, EA: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    1    tunnel1.41
192.168.2.0/24  2.1.1.141  20     A S    1    tunnel1.41
172.16.101.0/24  0.0.0.0    1      A H    1    ethernet1/1
2.1.1.140/24    2.1.1.141  20     A S    1    tunnel1.41

VPN PeerB
destination  next hop  metric  flags  age  interface  next-AS
192.168.1.0/24  0.0.0.0    10     A O    1    ethernet1/1
192.168.2.0/24  0.0.0.0    10     A O    1    ethernet1/15
172.16.101.0/24  2.1.1.140  20     A H    1    tunnel1.40
2.1.1.141/24    2.1.1.140  10     A C    1    tunnel1.40
```

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 *satellite* devices. This solution uses certificates for device authentication and IPSec to secure data.



LSVPN enables site-to-site VPNs between Palo Alto Networks firewalls. To set up a site-to-site VPN between a Palo Alto Networks firewall and another device, see [VPNs](#).

The following topics describe the LSVPN components and how to set them up to enable site-to-site VPN services between Palo Alto Networks firewalls:

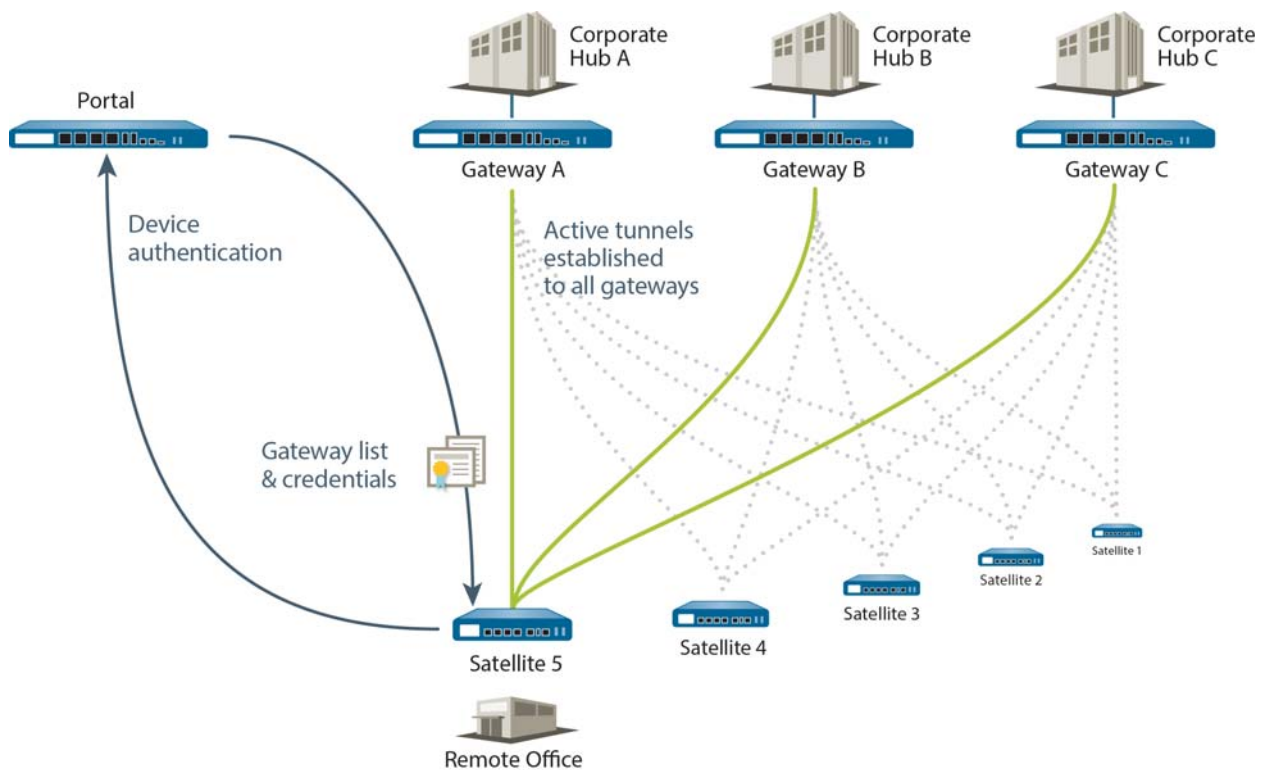
- ▲ [LSVPN Overview](#)
- ▲ [Create Interfaces and Zones for the LSVPN](#)
- ▲ [Enable SSL Between GlobalProtect LSVPN Components](#)
- ▲ [Configure the Portal to Authenticate Satellites](#)
- ▲ [Configure GlobalProtect Gateways for LSVPN](#)
- ▲ [Configure the GlobalProtect Portal for LSVPN](#)
- ▲ [Prepare the Satellite Device to Join the LSVPN](#)
- ▲ [Verify the LSVPN Configuration](#)
- ▲ [LSVPN Quick Configs](#)

LSVPN Overview

GlobalProtect provides a complete infrastructure for managing secure access to corporate resources from your remote sites. This infrastructure includes the following components:

- **GlobalProtect Portal**—Provides the management functions for your GlobalProtect LSVPN infrastructure. Every satellite that participates in the GlobalProtect LSVPN receives configuration information from the portal, including configuration information to enable the satellites (the spokes) to connect to the gateways (the hubs). You configure the portal on an interface on any Palo Alto Networks next-generation firewall.
- **GlobalProtect Gateways**—A Palo Alto Networks firewall that provides the tunnel end point for satellite connections. The resources that the satellites access is protected by security policy on the gateway. It is not required to have a separate portal and gateway; a single firewall can function both as portal and gateway.
- **GlobalProtect Satellite**—A Palo Alto Networks firewall at a remote site that establishes IPSec tunnels with the gateway(s) at your corporate office(s) for secure access to centralized resources. Configuration on the satellite firewall is minimal, enabling you to quickly and easily scale your VPN as you add new sites.

The following diagram illustrates how the GlobalProtect LSVPN components work together.




Create Interfaces and Zones for the LSVPN

You must configure the following interfaces and zones for your LSVPN infrastructure:

- **GlobalProtect portal**—Requires a Layer 3 interface for GlobalProtect satellites to connect to. If the portal and gateway are on the same firewall, they can use the same interface. The portal must be in a zone that is accessible from your branch offices.
- **GlobalProtect gateways**—Requires three interfaces: a Layer 3 interface in the zone that is reachable by the remote satellites, an internal interface in the trust zone that connects to the protected resources, and a logical tunnel interface for terminating the VPN tunnels from the satellites. Unlike other site-to-site VPN solutions, the GlobalProtect gateway only requires a single tunnel interface, which it will use for tunnel connections with all of your remote satellites (point-to-multipoint). If you plan to use dynamic routing, you must assign an IP address to the tunnel interface.
- **GlobalProtect satellites**—Requires a single tunnel interface for establishing a VPN with the remote gateways (up to a maximum of 25 gateways). If you plan to use dynamic routing, you must assign an IP address to the tunnel interface.

For more information about portals, gateways, and satellites see [LSVPN Overview](#).

Set Up Interfaces and Zones for the GlobalProtect LSVPN

<p>Step 1 Configure a Layer 3 interface.</p> <p>The portal and each gateway and satellite all require a Layer 3 interface to enable traffic to be routed between sites.</p> <p>If the gateway and portal are on the same firewall, you can use a single interface for both components.</p> <p> IPv6 addresses are not supported with LSVPN.</p>	<ol style="list-style-type: none"> 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: <ul style="list-style-type: none"> • The interface must be accessible from a zone outside of your trust network. Consider creating a dedicated VPN zone for visibility and control over your VPN traffic. • If you have not yet created the zone, select New Zone from the Security Zone drop-down, define a Name for the new zone and then click OK. 4. Select the Virtual Router to use. 5. To assign an IP address to the interface, select the IPv4 tab, click Add in the IP section, and enter the IP address and network mask to assign to the interface, for example 203.0.11.100/24. 6. To save the interface configuration, click OK.
---	---

Set Up Interfaces and Zones for the GlobalProtect LSVPN (Continued)

Step 2 On the firewall(s) hosting GlobalProtect gateway(s), configure the logical tunnel interface that will terminate VPN tunnels established by the GlobalProtect satellites.



IP addresses are not required on the tunnel interface unless plan to use dynamic routing. However, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues.



Make sure to enable User-ID in the zone where the VPN tunnels terminate.

1. Select **Network > Interfaces > Tunnel** and click **Add**.
2. In the **Interface Name** field, specify a numeric suffix, such as **.2**.
3. On the **Config** tab, expand the **Security Zone** drop-down to define the zone as follows:
 - To use your trust zone as the termination point for the tunnel, select the zone from the drop-down.
 - (Recommended) To create a separate zone for VPN tunnel termination, click **New Zone**. In the Zone dialog, define a **Name** for new zone (for example *lsvpn-tun*), select the **Enable User Identification** check box, and then click **OK**.
4. Select the **Virtual Router**.
5. (Optional) If you want to assign an IP address to the tunnel interface, select the **IPv4** tab, click **Add** in the IP section, and enter the IP address and network mask to assign to the interface, for example 203.0.11.33/24.
6. To save the interface configuration, click **OK**.

Step 3 If you created a separate zone for tunnel termination of VPN connections, create a security policy to enable traffic flow between the VPN zone and your trust zone. For example, the following policy rule enables traffic between the *lsvpn-tun* zone and the *L3-Trust* zone.

Name	Zone	Address	User	Zone	Address	Application	Service	Action
LSVPN Access	lsvpn-tun	any	any	L3-Trust	any	adobe-creati... ms-exchange ms-office365	application-d...	

Step 4 Save the configuration.

Click **Commit**.

Enable SSL Between GlobalProtect LSVPN Components

All interaction between the GlobalProtect components occurs over an SSL connection. Therefore, you must generate and/or install the required certificates before configuring each component so that you can reference the appropriate certificate(s) and/or certificate profiles in the configurations for each component. The following sections describe the supported methods of certificate deployment, descriptions and best practice guidelines for the various GlobalProtect certificates, and provide instructions for generating and deploying the required certificates:

- ▲ [About Certificate Deployment](#)
- ▲ [Deploy Server Certificates to the GlobalProtect LSVPN Components](#)

About Certificate Deployment

There are two basic approaches to deploying certificates for GlobalProtect LSVPN:

- **Enterprise Certificate Authority**—If you already have your own enterprise certificate authority, you can use this internal CA to issue an intermediate CA certificate for the GlobalProtect portal to enable it to issue certificates to the GlobalProtect gateways and satellites.
- **Self-Signed Certificates**—You can generate a self-signed root CA certificate on the firewall and use it to issue server certificates for the portal, gateway(s), and satellite(s). As a best practice, create a self-signed root CA certificate on the portal and use it to issue server certificates for the gateways and satellites. This way, the private key used for certificate signing stays on the portal.

Deploy Server Certificates to the GlobalProtect LSVPN Components


The GlobalProtect LSVPN components use SSL/TLS to mutually authenticate. Before deploying the LSVPN, you must issue server certificates to the portal and gateways. You do not need to create server certificates for the satellite devices because the portal will issue a server certificate for each satellite during the first connection.

In addition, you must import the root CA certificate used to issue the server certificates onto each firewall that you plan to host as a gateway or satellite. Finally, on each gateway and satellite participating in the LSVPN, you must configure a certificate profile that will enable them to establish an SSL/TLS connection using mutual authentication.

The following workflow shows the best practice steps for deploying SSL certificates to the GlobalProtect LSVPN components:



You do not need to issue server certificates for the satellite devices because the portal will issue them as part of the satellite registration process.

Deploy SSL Server Certificates to the GlobalProtect Components	
<p>Step 1 On the firewall hosting the portal, create the root CA certificate for issuing self-signed certificates for the GlobalProtect components.</p>	<p>To use self-signed certificates, you must first create the root CA certificate that will be used to sign the certificates for the GlobalProtect components as follows:</p> <ol style="list-style-type: none"> 1. To create a root CA certificate, select Device > Certificate Management > Certificates > Device Certificates and then click Generate. 2. Enter a Certificate Name, such as <i>LSVPN_CA</i>. The certificate name cannot contain any spaces. 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.
<p>Step 2 Generate server certificates for the GlobalProtect portal and gateway(s).</p> <p>You must issue a unique self-signed server certificate for the portal and for each GlobalProtect gateway. The best practice is to issue all of the required certificates on the portal, so that the signing certificate (with the private key) does not have to be exported.</p> <p> If the GlobalProtect portal and gateway are on the same firewall interface, you can use the same server certificate for both components.</p>	<p>Use the root CA on the portal to generate server certificates for each gateway you plan to deploy:</p> <ol style="list-style-type: none"> 1. Select Device > Certificate Management > Certificates > Device Certificates and then click Generate. 2. Enter a Certificate Name. The Certificate Name cannot contain any spaces. 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 <i>LSVPN_CA</i> you created previously. 5. In the Certificate Attributes section, click Add and define the attributes to uniquely identify the gateway. Keep in mind that if you add a Host Name attribute (which populates the SAN field of the certificate), it must exactly match the value you defined for the Common Name. 6. Generate the certificate.

Deploy SSL Server Certificates to the GlobalProtect Components (Continued)

<p>Step 3 Deploy the self-signed server certificates to the gateways.</p> <p>Best Practices:</p> <ul style="list-style-type: none"> • 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. 	<ol style="list-style-type: none"> 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.
<p>Step 4 Import the root CA certificate used to issue server certificates for the LSVPN components.</p> <p>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.</p>	<ol style="list-style-type: none"> 1. Download the root CA certificate from the portal. <ol style="list-style-type: none"> a. Select Device > Certificate Management > Certificates > Device Certificates. b. Select the root CA certificate used to issue certificates for the LSVPN components and click Export. c. Select Base64 Encoded Certificate (PEM) from the File Format drop-down and click OK to download the certificate. (Do not export the private key.) 2. On the firewalls hosting the gateways and satellites, import the root CA certificate. <ol style="list-style-type: none"> a. Select Device > Certificate Management > Certificates > Device Certificates and click Import. b. Enter a Certificate Name that identifies the certificate as your client CA certificate. c. Browse to the Certificate File you downloaded from the CA. d. Select Base64 Encoded Certificate (PEM) as the File Format and then click OK. e. Select the certificate you just imported on the Device Certificates tab to open it. f. Select Trusted Root CA and then click OK. g. Commit the changes.

Deploy SSL Server Certificates to the GlobalProtect Components (Continued)

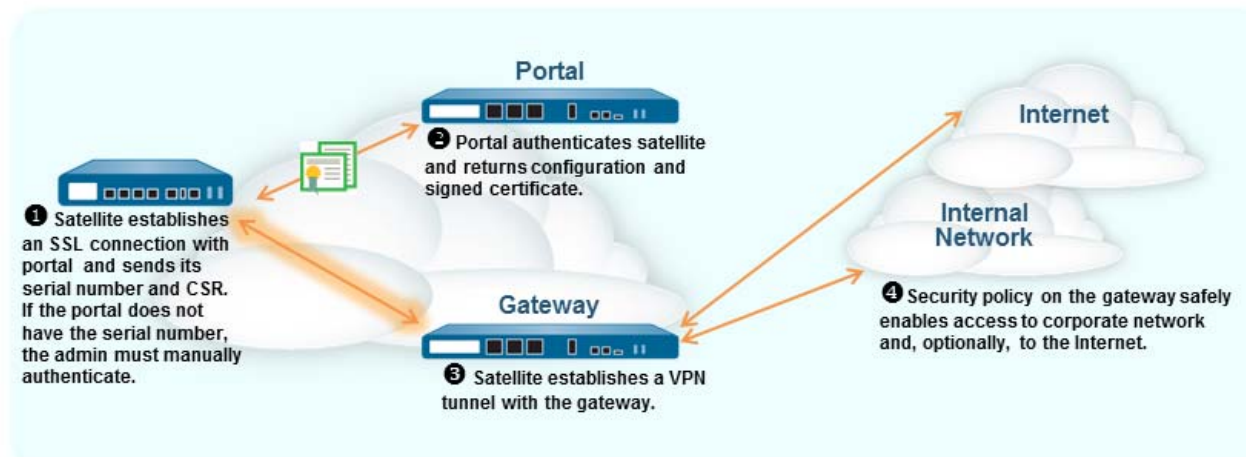
<p>Step 5 Create a certificate profile.</p> <p>The GlobalProtect LSVPN portal and each gateway require a certificate profile that specifies which certificate to use to authenticate the satellites.</p>	<ol style="list-style-type: none"> 1. Select Device > Certificate Management > Certificate Profile and 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 Step 4. 4. (Optional, but recommended) Enable use of CRL and/or OCSP to enable certificate status verification. 5. Click OK to save the profile.
<p>Step 6 Save the configuration.</p>	<p>Click Commit.</p>

Configure the Portal to Authenticate Satellites


In order to register with the LSVPN, each satellite must establish an SSL/TLS connection with the portal. After establishing the connection, the portal authenticates the satellite device to ensure that 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 satellite devices into the portal configuration, you can instead require the satellite administrator to authenticate when establishing the initial connection to the portal. Although the portal will always look for the serial number in the initial request from the satellite, if it cannot identify the serial number, the satellite administrator must provide a username and password to authenticate to the portal. Because the portal will always fall back to this form of authentication, you must create an authentication profile in order to commit the portal configuration. This requires that you set up an authentication profile for the portal LSVPN configuration even if you plan to authenticate satellites using the serial number.



The following workflow describes how to set up the portal to authenticate satellites against an existing authentication service. GlobalProtect LSVPN supports external authentication using a local database, LDAP (including Active Directory), Kerberos, or RADIUS.

Set Up Satellite Authentication	
<p>Step 1 Create a server profile on the portal.</p> <p>The server profile instructs the firewall how to connect to an external authentication service to validate the authentication credentials provided by the satellite administrator.</p> <p>If you are using local authentication, you can skip this step and instead add a local user configuration for authenticating the satellite administrator.</p> <p> If you are using LDAP to connect to Active Directory (AD), you must create a separate LDAP server profile for every AD domain.</p>	<ol style="list-style-type: none"> 1. Select Device > Server Profiles and select type of profile (LDAP, Kerberos, or RADIUS). 2. Click Add and enter a Name for the profile, such as <i>LSVPN-Auth</i>. 3. (LDAP only) Select the Type of LDAP server you are connecting to. 4. Click Add in the Servers section and then enter information required to connect to the authentication service, including the server Name, IP Address (or FQDN), and Port. 5. (RADIUS and LDAP only) Specify settings to enable the firewall to authenticate to the authentication service as follows: <ul style="list-style-type: none"> • RADIUS—Enter the shared Secret when adding the server entry. • LDAP—Enter the Bind DN and Bind Password. 6. (LDAP and Kerberos only) Specify where to search for credentials in the directory service: <ul style="list-style-type: none"> • LDAP—The Base DN specifies where in the LDAP tree to begin searching for users and groups. This field should populate automatically when you enter the server address and port. If it doesn't, check the service route to the LDAP server. • Kerberos—Enter the Kerberos Realm name. 7. Specify the Domain name (without dots, for example acme not acme.com). 8. Click OK to save the server profile.
<p>Step 2 Create an authentication profile.</p> <p>The authentication profile specifies which server profile to use to authenticate satellites.</p>	<ol style="list-style-type: none"> 1. Select Device > Authentication Profile and click Add. 2. Enter a Name for the profile and then select the Authentication type (Local Database, LDAP, Kerberos, or RADIUS). 3. Select the Server Profile you created in Step 1. 4. (LDAP AD) Enter sAMAccountName as the Login Attribute. 5. Click OK.
<p>Step 3 Save the configuration.</p>	<p>Click Commit.</p>

Configure GlobalProtect Gateways for LSVPN

Because the GlobalProtect configuration that the portal delivers to the satellites includes the list of gateways the satellite can connect to, it is a good idea to configure the gateways before configuring the portal.

- ▲ [Prerequisite Tasks](#)
- ▲ [Configure the Gateway](#)

Prerequisite Tasks

Before you can configure the GlobalProtect gateway, you must have completed the following tasks:


- ❑ Created the interfaces (and zones) for the interface where you plan to configure each gateway. You must configure both the physical interface and the virtual tunnel interface. See [Create Interfaces and Zones for the LSVPN](#).
- ❑ Set up the gateway server certificates and certificate profile required for enable GlobalProtect satellite and gateway to establish a mutual SSL/TLS connection. See [Enable SSL Between GlobalProtect LSVPN Components](#).

Configure the Gateway

After you have completed the [Prerequisite Tasks](#), configure each GlobalProtect gateway to participate in the LSVPN as follows:

Configure the Gateway for LSVPN	
<p>Step 1 Add a gateway.</p>	<ol style="list-style-type: none"> 1. Select Network > GlobalProtect > Gateways and click Add. 2. On the General tab, enter a Name for the gateway. The gateway name should not contain any spaces and as a best practice it should include the location or other descriptive information that will help identify the gateway. 3. (Optional) Select the virtual system to which this gateway belongs from the Location field.
<p>Step 2 Specify the network information to enable satellites to connect to the gateway.</p> <p>If you have not yet created the network interface for the gateway, see Create Interfaces and Zones for the LSVPN for instructions. If you haven't yet created a server certificate for the gateway, see Deploy Server Certificates to the GlobalProtect LSVPN Components.</p>	<ol style="list-style-type: none"> 1. Select the Interface that satellites will use for ingress access to the gateway. 2. Select the IP Address for gateway access. 3. Select the Server Certificate for the gateway from the drop-down.

Configure the Gateway for LSVPN (Continued)	
<p>Step 3 Select the certificate profile for the gateway to use to authenticate satellites attempting to establish tunnels.</p> <p>If you have not yet set up the certificate profile, see Enable SSL Between GlobalProtect LSVPN Components for instructions.</p>	<p>Select the Certificate Profile to you created for SSL communication between the LSVPN components.</p>
<p>Step 4 Configure the tunnel parameters and enable tunneling.</p>	<ol style="list-style-type: none"> 1. On the GlobalProtect Gateway dialog, select Satellite Configuration > Tunnel Settings. 2. Select the Tunnel Configuration check box to enable tunneling. 3. Select the Tunnel Interface you defined in Step 2 in 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 the Copy TOS check box.
<p>Step 5 (Optional) Enable tunnel monitoring.</p> <p>Tunnel monitoring enables satellite devices to monitor its gateway tunnel connection, allowing it to failover to a backup gateway if the connection fails. Failover to another gateway is the only type of tunnel monitoring profile supported with LSVPN.</p>	<ol style="list-style-type: none"> 1. Select the Tunnel Monitoring check box. 2. Specify the Destination IP address the satellite devices should use to determine if the gateway is active. Alternatively, if you configured an IP address for the tunnel interface, you can leave this field blank and the tunnel monitor will instead use the tunnel interface to determine if the connection is active. 3. Select Failover from the Tunnel Monitor Profile drop-down (this is the only supported tunnel monitor profile for LSVPN).
<p>Step 6 Select the Crypto Profile to use when establishing tunnel connections.</p> <p>The crypto profile specifies the type IPsec encryption and/or the authentication method for securing the data that will traverse the tunnel. Because both tunnel endpoints in an LSVPN are trusted firewalls within your organization, you can typically use the default profile, which uses ESP-DH group2-AES 128 with SHA-1 encryption.</p> <p>However, if you require a different mix of encryption and authentication mechanisms, you can optionally create a custom IPsec crypto profile.</p>	<p>Select default from the IPsec Crypto Profile drop-down or, optionally, select New IPsec Crypto Profile to define a new profile. For details on the authentication and encryption options in the crypto profile, refer to the online help.</p>

Configure the Gateway for LSVPN (Continued)	
<p>Step 7 Configure the network settings to assign the satellites during establishment of the IPsec tunnel.</p> <p> You can also configure the satellite device to push the DNS settings to its local clients by configuring a DHCP server on the firewall hosting the satellite. In this configuration, the satellite will push DNS settings it learns from the gateway to the DHCP clients.</p>	<ol style="list-style-type: none"> 1. On the GlobalProtect Gateway dialog, select Satellite Configuration > Network Settings. 2. (Optional) If clients local to the satellite device need to resolve FQDNs on the corporate network, configure the gateway to push DNS settings to the satellites in one of the following ways: <ul style="list-style-type: none"> • Manually define the Primary DNS, Secondary DNS, and DNS Suffix settings to push to the satellites. • If the gateway has an interface that is configured as a DHCP client, you can set the Inheritance Source to that interface and the GlobalProtect satellites will be assigned the same settings received by the DHCP client. 3. To specify the IP Pool of addresses to assign the tunnel interface on the satellite devices when the VPN is established, click Add and then specify the IP address range(s) to use. 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: <ul style="list-style-type: none"> • If you want to route all traffic from the satellites through the tunnel, leave this field blank. Note that in this case, all traffic except traffic destined for the local subnet will be tunneled to the gateway. • To route only some traffic through the gateway (called <i>split tunneling</i>), 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.
<p>Step 8 (Optional) Define what routes, if any, the gateway will accept from satellites.</p> <p>By default, the gateway will not add any routes satellites advertise to its routing table. If you do not want the gateway to accept routes from gateways, you do not need to complete this step.</p>	<ol style="list-style-type: none"> 1. To enable the gateway to accept routes advertised by satellites, select Satellite Configuration > 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.
<p>Step 9 Save the gateway configuration.</p>	<ol style="list-style-type: none"> 1. Click OK to save the settings and close the GlobalProtect Gateway dialog. 2. Commit the configuration.

Configure the GlobalProtect Portal for LSVPN

The GlobalProtect portal provides the management functions for your GlobalProtect LSVPN. Every satellite system that participates in the LSVPN receives configuration information from the portal, including information about available gateways as well as the certificate it needs in order to connect to the gateways.

The following sections provide procedures for setting up the portal:

- ▲ [Prerequisite Tasks](#)
- ▲ [Configure the Portal](#)
- ▲ [Define the Satellite Configurations](#)


Prerequisite Tasks

Before you can configure the GlobalProtect portal, you must have completed the following tasks:

- Created the interfaces (and zones) for the firewall interface where you plan to configure the portal. See [Create Interfaces and Zones for the LSVPN](#).
- Issued the portal server certificate, gateway server certificates, and set up the portal to issue server certificates for the satellites. See [Enable SSL Between GlobalProtect LSVPN Components](#).
- Defined the authentication profile that will be used to authenticate GlobalProtect satellites in the event that the serial number is not available. See [Configure the Portal to Authenticate Satellites](#).
- Configured the global protect gateways. See [Configure GlobalProtect Gateways for LSVPN](#).

Configure the Portal

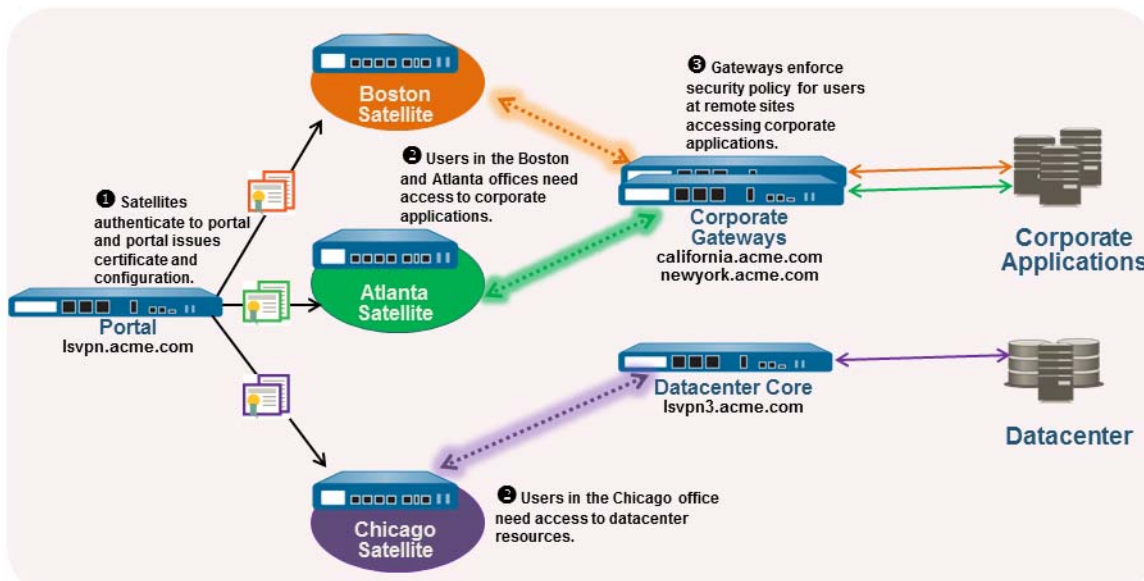
After you have completed the [Prerequisite Tasks](#), configure the GlobalProtect portal as follows:

Configure the Portal for LSVPN	
<p>Step 1 Add the portal.</p>	<ol style="list-style-type: none"> 1. Select Network > GlobalProtect > Portals and click Add. 2. On the Portal Configuration 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.
<p>Step 2 Specify the network information to enable satellites to connect to the portal.</p> <p>If you have not yet created the network interface for the portal, see Create Interfaces and Zones for the LSVPN for instructions. If you haven't yet created a server certificate for the portal and issued gateway certificates, see Deploy Server Certificates to the GlobalProtect LSVPN Components.</p>	<ol style="list-style-type: none"> 1. Select the Interface that satellites will use for ingress access to the portal. 2. Select the IP Address for satellite access to the portal. 3. Select the Server Certificate you generated to enable the satellite to establish SSL connection with the portal.
<p>Step 3 Specify an authentication profile for authenticating satellite devices.</p> <p> Even if you plan to manually configure the portal with the serial numbers of the satellites, you must define an authentication profile or you will not be able to save the configuration.</p>	<ul style="list-style-type: none"> • Select the Authentication Profile you defined for authenticating satellites. • If you have not yet set up the authentication profile, select New Authentication Profile to create one now. See Configure the Portal to Authenticate Satellites for instructions. If the portal is unable to validate the serial number of a connecting satellite, it will fall back to the authentication profile and therefore you must configure an authentication profile in order to save the portal configuration.
<p>Step 4 Continue with defining the configurations to push to the satellites, or, if you have already created the satellite configurations, save the portal configuration.</p>	<p>Click OK to save the portal configuration or continue to Define the Satellite Configurations.</p>



Define the Satellite Configurations


When a GlobalProtect satellite connects and successfully authenticates to the GlobalProtect portal, the portal delivers a satellite configuration, which specifies what gateways the satellite can connect to. If all your satellites will use the same gateway and certificate configurations, you can create a single satellite configuration to deliver to all satellites upon successful authentication. However, if you require different satellite configurations—for example if you want one group of satellites to connect to one gateway and another group of satellites to connect to a different gateway—you can create a separate satellite configuration for each. The portal will then use the enrollment username/group name or the serial number of the satellite device to determine which satellite configuration to deploy. As with security rule evaluation, the portal looks for a match starting from the top of the list. When it finds a match, it delivers the corresponding configuration to the satellite.

For example, the following figure shows a network in which some branch offices require VPN access to the corporate applications protected by your perimeter firewalls and another site needs VPN access to the datacenter.



Use the following procedure to create one or more satellite configurations.


Create a GlobalProtect Satellite Configuration	
<p>Step 1 Specify the certificates required to enable satellites to participate in the LSVPN.</p>	<ol style="list-style-type: none"> 1. Select Network > GlobalProtect > Portals and select the portal configuration for which you want to add a satellite configuration and then select the Satellite Configuration tab. 2. In the Trusted Root CA field, click Add and then select the CA certificate used to issue the gateway server certificates. The portal will deploy the root CA certificate you add here to all satellites as part of the configuration to enable the satellite to establish an SSL connection with the gateways. As a best practice, all of your gateways should use the same issuer. <p> 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.</p> 3. Select the Root CA certificate that the portal will use to issue certificates to satellites upon successfully authenticating them from the Issuing Certificate drop-down.
<p>Step 2 Add a satellite configuration.</p> <p>The satellite configuration specifies the GlobalProtect LSVPN configuration settings to deploy to the connecting satellites. You must define at least one satellite configuration.</p>	<p>In the Satellite Configuration section, click Add and enter a Name for the configuration.</p> <p>If you plan to create multiple configurations, make sure the name you define for each is descriptive enough to allow you to distinguish them.</p>
<p>Step 3 Specify which satellites to deploy this configuration to. There are two ways to specify which satellites will get the configuration: by enrollment user/group name and/or using the serial number of the satellite devices.</p> <p>The portal uses the Enrollment User/User Group settings and/or Devices serial numbers to match a satellite to a configuration. Therefore, if you have multiple configurations, be sure to order them properly. As soon as the portal finds a match, it will deliver the configuration. Therefore, more specific configurations must precede more general ones. See Step 6 for instructions on ordering the list of satellite configurations.</p>	<p>Specify the match criteria for the satellite configuration as follows:</p> <ul style="list-style-type: none"> • To restrict this configuration to satellite devices with specific serial numbers, select the Devices tab, click Add, and enter serial number (you do not need to enter the satellite hostname; it will be automatically added when the satellite connects). Repeat this step for each satellite you want to receive this configuration. • 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). <p> Before you can restrict the configuration to specific groups, you must Map Users to Groups.</p>

Create a GlobalProtect Satellite Configuration (Continued)	
<p>Step 4 Specify the gateways that satellites with this configuration can establish VPN tunnels with.</p> <p> 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.</p>	<ol style="list-style-type: none"> 1. On the Gateways tab, click Add. 2. Enter a descriptive Name for the gateway. The name you enter here should match the name you defined when you configured the gateway and should be descriptive enough identify the location of the gateway. 3. Enter the FQDN or IP address of the interface where the gateway is configured in the Gateways field. The address you specify must exactly match the Common Name (CN) in the gateway server certificate. 4. (Optional) If you are adding two or more gateways to the configuration, the Routing Priority helps the satellite pick the preferred gateway. Enter a value in the range of 1-25, with lower numbers having the higher priority (that is, the gateway the satellite will connect to if all gateways are available). The satellite will multiply the routing priority by 10 to determine the routing metric.
<p>Step 5 Save the satellite configuration.</p>	<ol style="list-style-type: none"> 1. Click OK to save the satellite configuration. 2. If you want to add another satellite configuration, repeat Step 2 through Step 5.
<p>Step 6 Arrange the satellite configurations so that the proper configuration is deployed to each satellite.</p>	<ul style="list-style-type: none"> • 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.
<p>Step 7 Save the portal configuration.</p>	<ol style="list-style-type: none"> 1. Click OK to save the settings and close the GlobalProtect Portal dialog. 2. Commit your changes.

Prepare the Satellite Device to Join the LSVPN

In order to participate in the LSVPN, the satellite devices require a minimal amount of configuration. Because the required configuration is minimal, you can pre-configure the devices before shipping them to your branch offices for installation.

Prepare the Satellite Device to Join the GlobalProtect LSVPN

<p>Step 1 Configure a Layer 3 interface.</p>	<p>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.</p>
<p>Step 2 Configure the logical tunnel interface for the tunnel to use to establish VPN tunnels with the GlobalProtect gateways.</p> <p> IP addresses are not required on the tunnel interface unless plan to use dynamic routing. However, assigning an IP address to the tunnel interface can be useful for troubleshooting connectivity issues.</p>	<ol style="list-style-type: none"> 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 <i>lsvpsat</i>). 4. In the Virtual Router drop-down, select default. 5. (Optional) If you want to assign an IP address to the tunnel interface, select the IPv4 tab, click Add in the IP section, and enter the IP address and network mask to assign to the interface, for example 2.2.2.11/24. 6. To save the interface configuration, click OK.

Prepare the Satellite Device to Join the GlobalProtect LSVPN (Continued)	
<p>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.</p> <p>The root CA certificate is required to enable the satellite device to establish the initial connection with the portal to obtain the LSVPN configuration.</p>	<ol style="list-style-type: none"> 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: <ol style="list-style-type: none"> a. Select Device > Certificate Management > Certificates > Device Certificates. b. Select the CA certificate, and click Export. c. Select Base64 Encoded Certificate (PEM) from the File Format drop-down and click OK to download the certificate. (You do not need to export the private key.) 2. Import the root CA certificate you just exported onto each satellite as follows. <ol style="list-style-type: none"> a. Select Device > Certificate Management > Certificates > Device Certificates and click Import. b. Enter a Certificate Name that identifies the certificate as your client CA certificate. c. Browse to the Certificate File you downloaded from the CA. d. Select Base64 Encoded Certificate (PEM) as the File Format and then click OK. e. Select the certificate you just imported on the Device Certificates tab to open it. f. Select Trusted Root CA and then click OK.
<p>Step 4 Configure the IPSec tunnel configuration.</p>	<ol style="list-style-type: none"> 1. Select Network > IPSec Tunnels and click Add. 2. On the General tab, enter a descriptive Name for the IPSec configuration. 3. Select the Tunnel Interface you created for the satellite. 4. Select GlobalProtect Satellite as the Type. 5. Enter the IP address or FQDN of the portal as the Portal Address. 6. Select the Layer 3 Interface you configured for the satellite. 7. Select the Local IP Address to use on the selected interface.
<p>Step 5 (Optional) Configure the satellite to publish local routes to the gateway.</p> <p>Pushing routes to the gateway enables traffic to the subnets local to the satellite via the gateway. However, the gateway must also be configured to accept the routes.</p>	<ol style="list-style-type: none"> 1. To enable the satellite to push routes to the gateway, on the Advanced tab select Publish all static and connected routes to Gateway. 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.
<p>Step 6 Save the satellite configuration.</p>	<ol style="list-style-type: none"> 1. Click OK to save the IPSec tunnel settings. 2. Click Commit.

Prepare the Satellite Device to Join the GlobalProtect LSVPN (Continued)

Step 7 If required, provide the credentials to allow the satellite to authenticate to the portal.

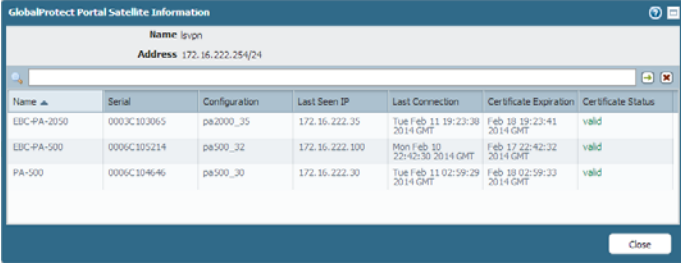
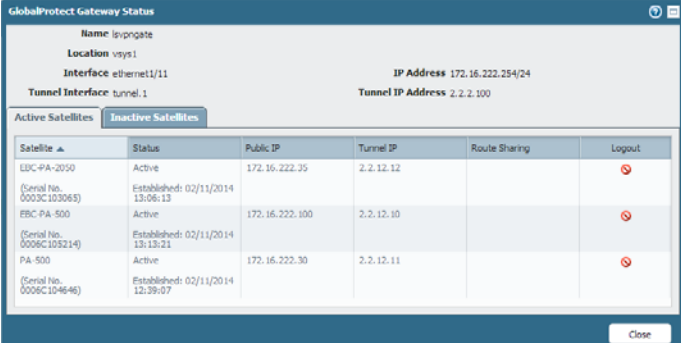
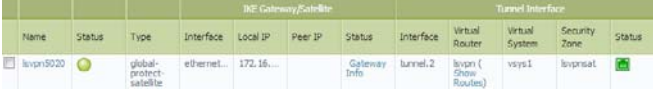
This step is only required if the portal was unable to find a serial number match in its configuration or if the serial number didn't work. In this case, the satellite will not be able to establish the tunnel with the gateway(s).

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 satellite devices, verify that the satellites are able to connect to the portal and gateway and establish VPN tunnels with the gateway(s).

Verify the LSVPN Configuration																																
<p>Step 1 Verify satellite connectivity with portal.</p>	<p>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.</p>  <table border="1"> <caption>GlobalProtect Portal Satellite Information</caption> <thead> <tr> <th>Name</th> <th>Serial</th> <th>Configuration</th> <th>Last Seen IP</th> <th>Last Connection</th> <th>Certificate Expiration</th> <th>Certificate Status</th> </tr> </thead> <tbody> <tr> <td>EBC-PA-2050</td> <td>000C103065</td> <td>pa2000_35</td> <td>172.16.222.35</td> <td>Tue Feb 11 19:23:38 2014 GMT</td> <td>Feb 18 19:23:41 2014 GMT</td> <td>valid</td> </tr> <tr> <td>EBC-PA-500</td> <td>0006C105214</td> <td>pa500_32</td> <td>172.16.222.100</td> <td>Mon Feb 10 22:42:30 2014 GMT</td> <td>Feb 17 22:42:32 2014 GMT</td> <td>valid</td> </tr> <tr> <td>PA-500</td> <td>0006C104646</td> <td>pa500_30</td> <td>172.16.222.30</td> <td>Tue Feb 11 02:59:29 2014 GMT</td> <td>Feb 18 02:59:33 2014 GMT</td> <td>valid</td> </tr> </tbody> </table>	Name	Serial	Configuration	Last Seen IP	Last Connection	Certificate Expiration	Certificate Status	EBC-PA-2050	000C103065	pa2000_35	172.16.222.35	Tue Feb 11 19:23:38 2014 GMT	Feb 18 19:23:41 2014 GMT	valid	EBC-PA-500	0006C105214	pa500_32	172.16.222.100	Mon Feb 10 22:42:30 2014 GMT	Feb 17 22:42:32 2014 GMT	valid	PA-500	0006C104646	pa500_30	172.16.222.30	Tue Feb 11 02:59:29 2014 GMT	Feb 18 02:59:33 2014 GMT	valid			
Name	Serial	Configuration	Last Seen IP	Last Connection	Certificate Expiration	Certificate Status																										
EBC-PA-2050	000C103065	pa2000_35	172.16.222.35	Tue Feb 11 19:23:38 2014 GMT	Feb 18 19:23:41 2014 GMT	valid																										
EBC-PA-500	0006C105214	pa500_32	172.16.222.100	Mon Feb 10 22:42:30 2014 GMT	Feb 17 22:42:32 2014 GMT	valid																										
PA-500	0006C104646	pa500_30	172.16.222.30	Tue Feb 11 02:59:29 2014 GMT	Feb 18 02:59:33 2014 GMT	valid																										
<p>Step 1 Verify satellite connectivity with the gateway(s).</p>	<p>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.</p>  <table border="1"> <caption>GlobalProtect Gateway Status</caption> <thead> <tr> <th>Satellite</th> <th>Status</th> <th>Public IP</th> <th>Tunnel IP</th> <th>Route Sharing</th> <th>Logout</th> </tr> </thead> <tbody> <tr> <td>EBC-PA-2050 (Serial No. 0000C103065)</td> <td>Active</td> <td>172.16.222.35</td> <td>2.2.12.12</td> <td></td> <td>Logout</td> </tr> <tr> <td>EBC-PA-500 (Serial No. 0006C105214)</td> <td>Active</td> <td>172.16.222.100</td> <td>2.2.12.10</td> <td></td> <td>Logout</td> </tr> <tr> <td>PA-500 (Serial No. 0006C104646)</td> <td>Active</td> <td>172.16.222.30</td> <td>2.2.12.11</td> <td></td> <td>Logout</td> </tr> </tbody> </table>	Satellite	Status	Public IP	Tunnel IP	Route Sharing	Logout	EBC-PA-2050 (Serial No. 0000C103065)	Active	172.16.222.35	2.2.12.12		Logout	EBC-PA-500 (Serial No. 0006C105214)	Active	172.16.222.100	2.2.12.10		Logout	PA-500 (Serial No. 0006C104646)	Active	172.16.222.30	2.2.12.11		Logout							
Satellite	Status	Public IP	Tunnel IP	Route Sharing	Logout																											
EBC-PA-2050 (Serial No. 0000C103065)	Active	172.16.222.35	2.2.12.12		Logout																											
EBC-PA-500 (Serial No. 0006C105214)	Active	172.16.222.100	2.2.12.10		Logout																											
PA-500 (Serial No. 0006C104646)	Active	172.16.222.30	2.2.12.11		Logout																											
<p>Step 1 Verify LSVPN tunnel status on the satellite.</p>	<p>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.</p>  <table border="1"> <thead> <tr> <th rowspan="2">Name</th> <th rowspan="2">Status</th> <th rowspan="2">Type</th> <th colspan="3">IKE Gateway/Satellite</th> <th colspan="4">Tunnel Interface</th> </tr> <tr> <th>Interface</th> <th>Local IP</th> <th>Peer IP</th> <th>Status</th> <th>Interface</th> <th>Virtual Router</th> <th>Virtual System</th> <th>Security Zone</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td>lsvpn5020</td> <td>Active</td> <td>global-protect-satellite</td> <td>ethernet...</td> <td>172.16...</td> <td></td> <td>Gateway Info</td> <td>tunnel.2</td> <td>lsvpn (Show Routes)</td> <td>vsys1</td> <td>lsvpnrat</td> <td>Active</td> </tr> </tbody> </table>	Name	Status	Type	IKE Gateway/Satellite			Tunnel Interface				Interface	Local IP	Peer IP	Status	Interface	Virtual Router	Virtual System	Security Zone	Status	lsvpn5020	Active	global-protect-satellite	ethernet...	172.16...		Gateway Info	tunnel.2	lsvpn (Show Routes)	vsys1	lsvpnrat	Active
Name	Status				Type	IKE Gateway/Satellite			Tunnel Interface																							
		Interface	Local IP	Peer IP		Status	Interface	Virtual Router	Virtual System	Security Zone	Status																					
lsvpn5020	Active	global-protect-satellite	ethernet...	172.16...		Gateway Info	tunnel.2	lsvpn (Show Routes)	vsys1	lsvpnrat	Active																					

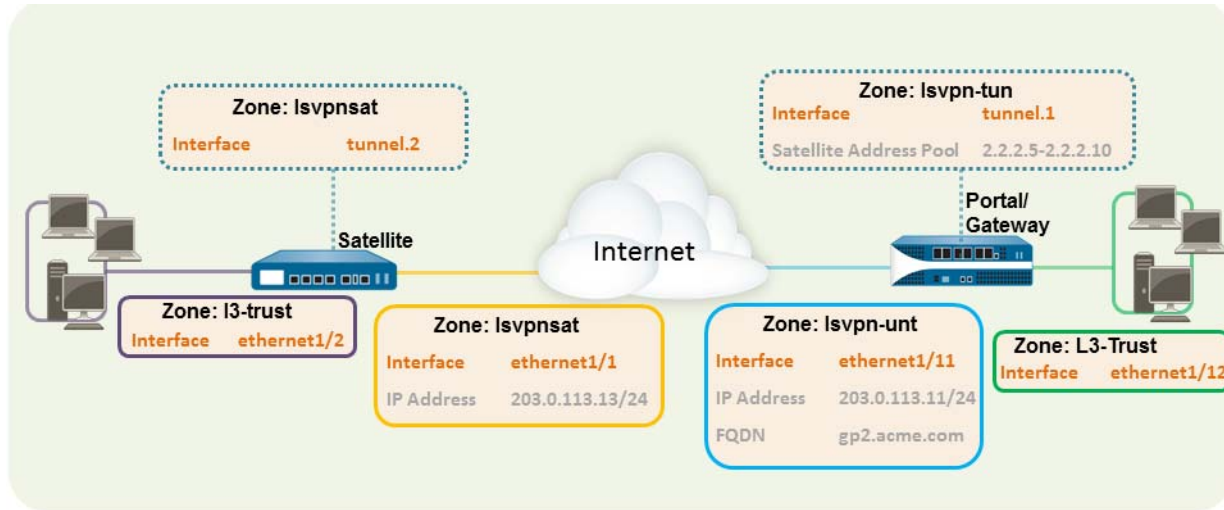
LSVPN Quick Configs

The following sections provide step-by-step instructions for configuring some common GlobalProtect LSVPN deployments:

- ▲ [Basic LSVPN Configuration with Static Routing](#)
- ▲ [Advanced LSVPN Configuration with Dynamic Routing](#)


Basic LSVPN Configuration with Static Routing

This quick config shows the fastest way to get up and running with LSVPN. In this example, a single firewall at the corporate headquarters site is configured as both a portal and a gateway. Satellite devices can be quickly and easily deployed with minimal configuration for optimized scalability.





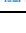




The following workflow shows the steps for setting up this basic configuration:

Quick Config: Basic LSVPN with Static Routing

<p>Step 1 Configure a Layer 3 interface.</p>	<p>In this example, the Layer 3 interface on the portal/gateway requires the following configuration:</p> <ul style="list-style-type: none"> • Interface—ethernet1/11 • Security Zone—lsvpn-unt • IPv4—203.0.113.11/24
<p>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.</p> <p> To enable visibility into users and groups connecting over the VPN, enable User-ID in the zone where the VPN tunnels terminate.</p>	<p>In this example, the Tunnel interface on the portal/gateway requires the following configuration:</p> <ul style="list-style-type: none"> • 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).

Name	Zone	Address	User	Zone	Address	Application	Service	Action
LSVPN Access	 lsvpn-tun	any	any	 L3-Trust	any	 adobe-creati...  ms-exchange  ms-office365	 application-d...	

Quick Config: Basic LSVPN with Static Routing (Continued)	
<p>Step 4 Issue a self-signed server certificate for the portal/gateway.</p> <p>The certificate subject name must match the FQDN or IP address of the Layer 3 interface you create for the portal/gateway.</p>	<ol style="list-style-type: none"> 1. On the firewall hosting the portal, create the root CA certificate for issuing self-signed certificates for the GlobalProtect components. In this example, the root CA certificate, <i>lsvpn-CA</i>, will be used to issue the server certificate for the portal/gateway. In addition, the portal will use this root CA certificate to sign the CSRs from the satellite devices. 2. Generate server certificates for the GlobalProtect portal and gateway(s). Because the portal and gateway will be on the same interface in this example, they can share a server certificate. In this example, the server certificate is named <i>lsvpnserver</i>.
<p>Step 5 Create a certificate profile.</p>	<p>In this example, the certificate profile, <i>lsvpn-profile</i>, references the root CA certificate <i>lsvpn-CA</i>. The gateway will use this certificate profile to authenticate satellites attempting to establish VPN tunnels.</p>
<p>Step 6 Configure an authentication profile for the portal to use if the satellite serial number is not available.</p>	<ol style="list-style-type: none"> 1. Create a server profile on the portal. 2. Create an authentication profile. In this example, the profile <i>lsvpn-sat</i> is used to authenticate satellites.
<p>Step 7 Configure the Gateway for LSVPN.</p>	<p>Select Network > GlobalProtect > Gateways and Add a configuration. This example requires the following gateway configuration:</p> <ul style="list-style-type: none"> • Interface—ethernet1/11 • IP Address—203.0.113.11/24 • Server Certificate—lsvpnserver • Certificate Profile—lsvpn-profile • Tunnel Interface—tunnel.1 • Primary DNS/Secondary DNS—4.2.2.1/4.2.2.2 • IP Pool—2.2.2.111-2.2.2.120 • Access Route—10.2.10.0/24
<p>Step 8 Configure the Portal for LSVPN.</p>	<p>Select Network > GlobalProtect > Portal and Add a configuration. This example requires the following portal configuration:</p> <ul style="list-style-type: none"> • Interface—ethernet1/11 • IP Address—203.0.113.11/24 • Server Certificate—lsvpnserver • Authentication Profile—lsvpn-sat

Quick Config: Basic LSVPN with Static Routing (Continued)	
<p>Step 9 Create a GlobalProtect Satellite Configuration.</p>	<p>On the Satellite Configuration tab in the portal configuration, Add a Satellite Configuration and a Trusted Root CA and specify the CA the portal will use to issue certificates for the satellites. In this example the required settings are as following:</p> <ul style="list-style-type: none"> • Gateway—203.0.113.11 • Issuing Certificate—lsvpn-CA • Trusted Root CA—lsvpn-CA
<p>Step 10 Prepare the Satellite Device to Join the LSVPN.</p>	<p>The satellite configuration in this example requires the following settings:</p> <p>Interface Configuration</p> <ul style="list-style-type: none"> • Layer 3 interface—ethernet1/1, 203.0.113.13/24 • Tunnel interface—tunnel.2 • Zone—lsvpnsat <p>Root CA Certificate from Portal</p> <ul style="list-style-type: none"> • lsvpn-CA <p>IPSec Tunnel Configuration</p> <ul style="list-style-type: none"> • 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

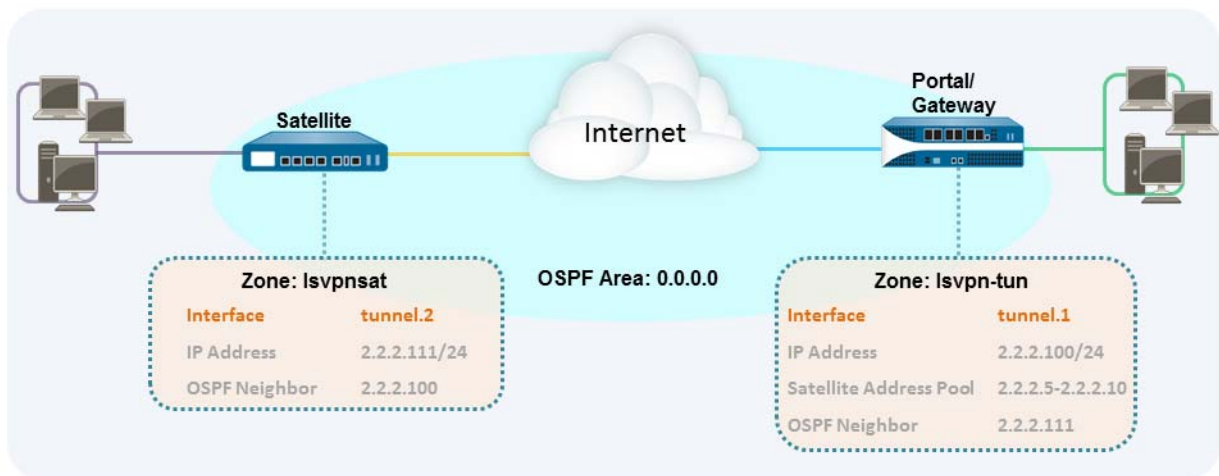
In larger LSVPN deployments with multiple gateways and many satellites, investing a little more time in the initial configuration to set up dynamic routing will simplify the maintenance of gateway configurations because access routes will update dynamically. The following example configuration shows 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:

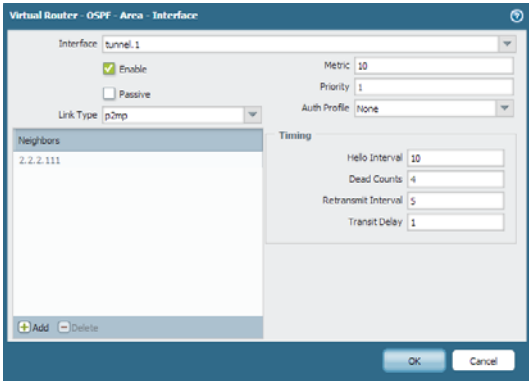
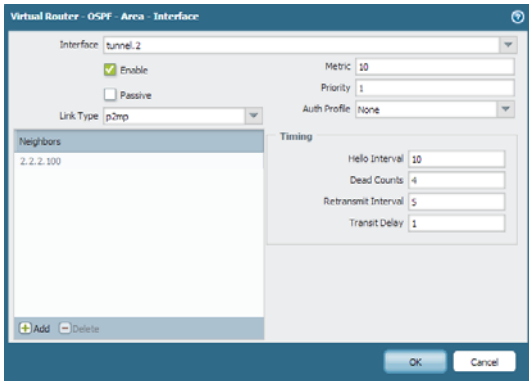
- Manual assignment of IP addresses to tunnel interfaces on all gateways and satellites.
- Configuration of OSPF point-to-multipoint (P2MP) on the virtual router on all gateways and satellites. In addition, as part of the OSPF configuration on each gateway, you must manually define the tunnel IP address 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.



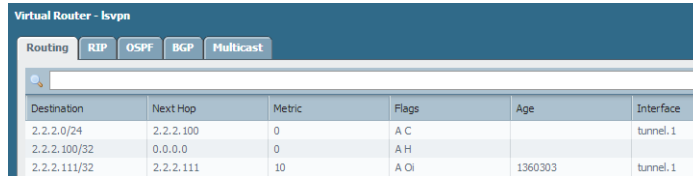
For a basic setup of a LSVPN, follow the steps in [Basic LSVPN Configuration with Static Routing](#). You can then complete the steps in the following workflow to extend the configuration to use dynamic routing rather than static routing.

Quick Config: LSVPN with Dynamic Routing	
<p>Step 1 Add an IP address to the tunnel interface configuration on each gateway and each satellite.</p>	<p>Complete the following steps on each gateway and each satellite:</p> <ol style="list-style-type: none"> 1. Select Network > Interfaces > Tunnel and select the tunnel configuration you created for the LSVPN to open the Tunnel Interface dialog. If you have not yet created the tunnel interface, see Step 2 in Quick Config: Basic LSVPN with Static Routing. 2. On the IPv4 tab, click Add and then enter an IP address and subnet mask. For example, to add an IP address for the gateway tunnel interface you would enter 2.2.2.100/24. 3. Click OK to save the configuration.
<p>Step 2 Configure the dynamic routing protocol on the gateway.</p> 	<p>To configure OSPF on the gateway:</p> <ol style="list-style-type: none"> 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 device, 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.
<p>Step 3 Configure the dynamic routing protocol on the satellite.</p> 	<p>To configure OSPF on the satellite:</p> <ol style="list-style-type: none"> 1. Select Network > Virtual Routers and select the virtual router associated with your VPN interfaces. 2. On the Areas tab, click Add to create the backbone area, or, if it is already configured, click on the area ID to edit it. 3. If you are creating a new area, enter an Area ID on the Type tab. 4. On the Interface tab, click Add and select the tunnel Interface you created for the LSVPN. 5. Select p2mp as the Link Type. 6. Click Add in the Neighbors section and enter the IP address of the tunnel interface of each GlobalProtect gateway, for example 2.2.2.100. 7. Click OK twice to save the virtual router configuration and then Commit the changes on the gateway. 8. Repeat this step each time you add a new gateway.

Quick Config: LSVPN with Dynamic Routing (Continued)

Step 4 Verify that the gateways and satellites are able to form router adjacencies.

- On each satellite and each gateway, confirm that peer adjacencies have formed and that routing table entries have been created for the peers (that is, the satellites have routes to the gateways and the gateways have routes to the satellites). Select **Network > Virtual Router** and click the **More Runtime Stats** link for the virtual router you are using for the LSVPN. On the Routing tab, verify that the LSVPN peer has a route.

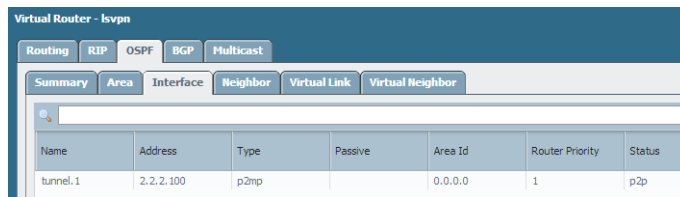


Virtual Router - Isvpn

Routing RIP OSPF BGP Multicast

Destination	Next Hop	Metric	Flags	Age	Interface
2.2.2.0/24	2.2.2.100	0	A C		tunnel.1
2.2.2.100/32	0.0.0.0	0	A H		tunnel.1
2.2.2.111/32	2.2.2.111	10	A OI	1360303	tunnel.1

- On the **OSPF > Interface** tab, verify that the **Type** is **p2mp**.



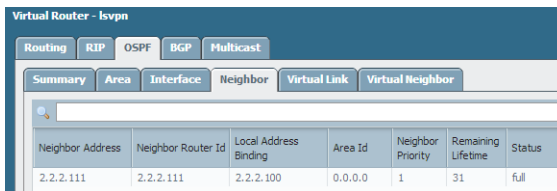
Virtual Router - Isvpn

Routing RIP OSPF BGP Multicast

Summary Area Interface Neighbor Virtual Link Virtual Neighbor

Name	Address	Type	Passive	Area Id	Router Priority	Status
tunnel.1	2.2.2.100	p2mp		0.0.0.0	1	p2p

- 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.



Virtual Router - Isvpn

Routing RIP OSPF BGP Multicast

Summary Area Interface Neighbor Virtual Link Virtual Neighbor

Neighbor Address	Neighbor Router Id	Local Address Binding	Area Id	Neighbor Priority	Remaining Lifetime	Status
2.2.2.111	2.2.2.111	2.2.2.100	0.0.0.0	1	31	full



Networking

All Palo Alto Networks next-generation firewalls provide a flexible networking architecture that includes support for dynamic routing, switching, and VPN connectivity, enabling you to deploy the firewall into nearly any networking environment. When configuring the Ethernet ports on your firewall, you can choose from virtual wire, Layer 2, or Layer 3 interface deployments. In addition, to allow you to integrate into a variety of network segments, you can configure different types of interfaces on different ports. The [Interface Deployments](#) section provides basic information on each type of deployment. For more detailed deployment information, refer to [Designing Networks with Palo Alto Networks Firewalls](#).

The following topics describe networking concepts and how to integrate Palo Alto Networks next-generation firewalls into your network.

- ▲ [Interface Deployments](#)
- ▲ [Virtual Routers](#)
- ▲ [Static Routes](#)
- ▲ [RIP](#)
- ▲ [OSPF](#)
- ▲ [BGP](#)
- ▲ [Session Settings and Timeouts](#)
- ▲ [DHCP](#)
- ▲ [NAT](#)
- ▲ [NAT64](#)
- ▲ [LACP](#)

For information on route distribution, refer to [Understanding Route Redistribution and Filtering](#).

Interface Deployments

A Palo Alto Networks firewall can operate in multiple deployments at once because the deployments occur at the interface level. The following sections describe the deployments.

- ▲ [Virtual Wire Deployments](#)
- ▲ [Layer 2 Deployments](#)
- ▲ [Layer 3 Deployments](#)
- ▲ [Tap Mode Deployments](#)

Virtual Wire Deployments


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

A virtual wire deployment allows the following conveniences:

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

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

Figure: Virtual Wire Deployment



Virtual Wire Subinterfaces

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

- **VLAN tags** —The example in [Figure: Virtual Wire Deployment with Subinterfaces \(VLAN Tags only\)](#), shows an Internet Service Provider (ISP) using virtual wire subinterfaces with VLAN tags to separate traffic for two different customers.

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

Virtual Wire Subinterface Workflow

1. Configure two Ethernet interfaces as type virtual wire, and assign these interfaces to a virtual wire.
2. Create subinterfaces on the parent Virtual Wire to separate CustomerA and CustomerB traffic. Make sure that the VLAN tags defined on each pair of subinterfaces that are configured as virtual wire(s) are identical. This is essential because a virtual wire does not switch VLAN tags.
3. Create new subinterfaces and define IP classifiers. This task is optional and only required if you wish to add additional subinterfaces with IP classifiers for further managing traffic from a customer based on the combination of VLAN tags and a specific source IP address, range or subnet.
You can also use IP classifiers for managing untagged traffic. To do so, you must create a sub-interface with the vlan tag “0”, and define sub-interface(s) with IP classifiers for managing untagged traffic using IP classifiers



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

Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags only)




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

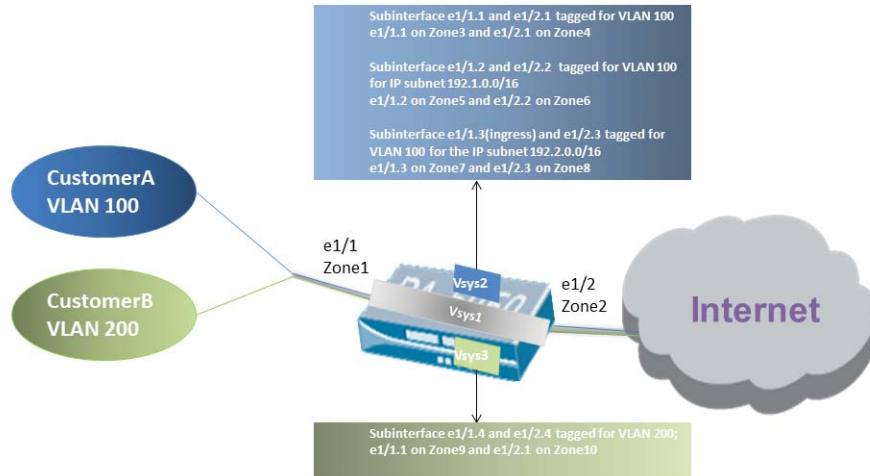
When traffic enters the firewall from CustomerA or CustomerB, the VLAN tag on the incoming packet is first matched against the VLAN tag defined on the ingress subinterfaces. In this example, a single subinterface matches the VLAN tag on the incoming packet, hence that subinterface is selected. The policies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



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

Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags and IP Classifiers) depicts CustomerA and CustomerB connected to one physical firewall that has two virtual systems (vsys), in addition to the default virtual system (vsys1). Each virtual system is an independent virtual firewall that is managed separately for each customer. Each vsys has attached interfaces/subinterfaces and security zones that are managed independently.

Figure: Virtual Wire Deployment with Subinterfaces (VLAN Tags and IP Classifiers)



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

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

Customer	Vsys	Vwire Subinterfaces	Zone	VLAN Tag	IP Classifier
A	2	e1/1.1 (ingress) e1/2.1 (egress)	Zone3 Zone4	100 100	None
	2	e1/1.2 (ingress) e1/2.2 (egress)	Zone5 Zone6	100 100	IP subnet 192.1.0.0/16
	2	e1/1.3 (ingress) e1/2.3 (egress)	Zone7 Zone8	100 100	IP subnet 192.2.0.0/16
B	3	e1/1.4 (ingress)	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.

Layer 2 Deployments

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


Figure: Layer 2 Deployment



Layer 3 Deployments

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

Figure: Layer 3 Deployment



In addition, because the firewall must route traffic in a Layer 3 deployment, you must configure a virtual router. See [Virtual Routers](#).

Point-to-Point Protocol over Ethernet Support

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

You can choose the PPPoE option and configure the associated settings when an interface is defined as a Layer 3 interface.



PPPoE is not supported in HA active/active mode.

DHCP Client

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



DHCP client is not supported in HA active/active mode.

For more information, see [DHCP](#).

Tap Mode Deployments

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

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



When deployed in tap mode, the firewall is not able to take action, such as block traffic or apply QoS traffic control.

Virtual Routers

The firewall uses virtual routers to obtain routes to other subnets by manually defining a route (static routes) or through participation in Layer 3 routing protocols (dynamic routes). The best routes obtained through these methods are used to populate the firewall's IP route table. When a packet is destined for a different subnet, the Virtual Router obtains the best route from this IP route table and forwards the packet to the next hop router defined in the table.

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

You can configure the virtual router to participate with dynamic routing protocols (BGP, OSPF, or RIP) as well as adding static routes. You can also create multiple virtual routers, each maintaining a separate set of routes that are not shared between virtual routers, enabling you to configure different routing behaviors for different interfaces.

Each Layer 3 interface, loopback interface, and VLAN interface defined on the firewall must be associated with a virtual router. While each interface can belong to only one virtual router, multiple routing protocols and static routes can be configured for a virtual router. Regardless of the static routes and dynamic routing protocols configured for a virtual router, a common general configuration is required. The firewall uses Ethernet switching to reach other devices on the same IP subnet.

The following Layer 3 routing protocols are supported from Virtual Routers:

- RIP
- OSPF
- OSPFv3
- BGP

Define a Virtual Router General Configuration	
Step 1 Gather the required information from your network administrator.	<ul style="list-style-type: none"> • Interfaces that you want to route • Administrative distances for static, OSFP internal, OSPF external, IBGP, EBGP and RIP
Step 2 Create the virtual router and name it.	<ol style="list-style-type: none"> 1. Select Network > Virtual Routers. 2. Click Add and enter a name for the virtual router. 3. Select interfaces to apply to the virtual router. 4. Click OK.
Step 3 Select interfaces to apply to the virtual router.	<ol style="list-style-type: none"> 1. Click Add in the Interfaces box. 2. Select an already defined interface from the drop-down. 3. Repeat Step 2 for all interfaces that you want to add to the virtual router.



Define a Virtual Router General Configuration (Continued)

<p>Step 4 Set Administrative Distances for static and dynamic routing.</p>	<p>1. Set Administrative Distances as required.</p> <ul style="list-style-type: none"> • Static—Range: 10-240, Default: 10 • OSPF Internal—Range: 10-240, Default: 30 • OSPF External—Range: 10-240, Default: 110 • IBGP—Range: 10-240, Default: 200 • EBGP—Range: 10-240, Default: 20 • RIP—Range: 10-240, Default: 120
<p>Step 5 Save virtual router general settings.</p>	<p>Click OK to save your settings.</p>
<p>Step 6 Commit your changes.</p>	<p>Click Commit. The device may take up to 90 seconds to save your changes.</p>

Static Routes

The following procedure shows how to integrate the firewall into the network using static routing.

Set Up Interfaces and Zones	
<p>Step 1 Configure a default route to your Internet router.</p>	<ol style="list-style-type: none"> 1. Select Network > Virtual Router and then select the default link to open the Virtual Router dialog. 2. Select the Static Routes tab and click Add. Enter a Name for the route and enter the route in the Destination field (for example, 0.0.0.0/0). 3. Select the IP Address radio button in the Next Hop field and then enter the IP address and netmask for your Internet gateway (for example, 208.80.56.1). 4. Click OK twice to save the virtual router configuration.
<p>Step 2 Configure the external interface (the interface that connects to the Internet).</p>	<ol style="list-style-type: none"> 1. Select Network > Interfaces and then select the interface you want to configure. In this example, we are configuring Ethernet1/3 as the external interface. 2. Select the Interface Type. Although your choice here depends on your network topology, this example shows the steps for Layer3. 3. In the Virtual Router drop-down, select default. 4. On the Config tab, select New Zone from the Security Zone drop-down. In the Zone dialog, define a Name for new zone, for example Untrust, and then click OK. 5. To assign an IP address to the interface, select the IPv4 tab and Static radio button. Click Add in the IP section, and enter the IP address and network mask to assign to the interface, for example 208.80.56.100/24. 6. To enable you to ping the interface, select Advanced > Other Info, expand the Management Profile drop-down, and select New Management Profile. Enter a Name for the profile, select Ping and then click OK. 7. To save the interface configuration, click OK.

Set Up Interfaces and Zones (Continued)	
<p>Step 3 Configure the interface that connects to your internal network.</p> <p> In this example, the interface connects to a network segment that uses private IP addresses. Because private IP addresses cannot be routed externally, you will have to configure NAT. See Configure NAT for details.</p>	<ol style="list-style-type: none"> 1. Select Network > Interfaces and select the interface you want to configure. In this example, we are configuring Ethernet1/4 as the internal interface. 2. Select Layer3 from the Interface Type drop-down. 3. On the Config tab, expand the Security Zone drop-down and select New Zone. In the Zone dialog, define a Name for new zone, for example Trust, and then click OK. 4. Select the same Virtual Router you used in Step 2, default in this example. 5. To assign an IP address to the interface, select the IPv4 tab and the Static radio button, click Add in the IP section, and enter the IP address and network mask to assign to the interface, for example 192.168.1.4/24. 6. To enable you to ping the interface, select the management profile that you created in Step 2-6. 7. To save the interface configuration, click OK.
<p>Step 4 Configure the interface that connects to the DMZ.</p>	<ol style="list-style-type: none"> 1. Select the interface you want to configure. 2. Select Layer3 from the Interface Type drop-down. In this example, we are configuring Ethernet1/13 as the DMZ interface. 3. On the Config tab, expand the Security Zone drop-down and select New Zone. In the Zone dialog, define a Name for new zone, for example DMZ, and then click OK. 4. Select the Virtual Router you used in Step 2, default in this example. 5. To assign an IP address to the interface, select the IPv4 tab and the Static radio button, click Add in the IP section, and enter the IP address and network mask to assign to the interface, for example 10.1.1.1/24. 6. To enable you to ping the interface, select the management profile that you created in Step 2-6. 7. To save the interface configuration, click OK.
<p>Step 5 Save the interface configuration.</p>	<p>Click Commit.</p>
<p>Step 6 Cable the firewall.</p>	<p>Attach straight through cables from the interfaces you configured to the corresponding switch or router on each network segment.</p>
<p>Step 7 Verify that the interfaces are active.</p>	<p>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.</p> 

RIP

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

Perform the following procedure to configure RIP.

Configure RIP	
Step 1 Configure general virtual router configuration settings.	See Virtual Routers for details.
Step 2 Configure general RIP configuration settings.	<ol style="list-style-type: none"> 1. Select the RIP tab. 2. Select the Enable check box to enable the RIP protocol. 3. Select the Reject Default Route check box if you do not want to learn any default routes through RIP. This is the recommended default setting. 4. De-select the Reject Default Route check box if you want to permit redistribution of default routes through RIP.
Step 3 Configure interfaces for the RIP protocol.	<ol style="list-style-type: none"> 1. On Interfaces tab, select an interface from the drop-down in the Interface configuration section. 2. Select an already defined interface from the drop-down. 3. Select the Enable check box. 4. Select the Advertise check box to advertise a default route to RIP peers with the specified metric value. 5. You can optionally select a profile from the Auth Profile drop-down. See Step 5 for details. 6. Select normal, passive or send-only from the Mode drop-down. 7. Click OK.
Step 4 Configure RIP timers.	<ol style="list-style-type: none"> 1. On the Timers tab, enter a value in the Interval Seconds (sec) box. This setting defines the length of the timer interval in seconds. This duration is used for the remaining RIP timing fields. Default: 1. Range: 1 - 60. 2. Enter an Update Intervals to define the number of intervals between route update announcements Default: 30. Range: 1 - 3600. 3. Specify the Delete Intervals to define the number of intervals between the time that the route expires to its deletion Default: 180. Range: 1 - 3600. 4. Specify the Expire Intervals to define the number of intervals between the time that the route was last updated to its expiration Default: 120. Range: 1 - 3600.

Configure RIP**Step 5** (Optional) Configure Auth Profiles.

By default, the firewall does not use RIP authentication for the exchange between RIP neighbors. Optionally, you can configure RIP authentication between RIP neighbors by either a simple password or using MD5 authentication.

Simple Password RIP authentication

1. Select **Auth Profiles** and click **Add**.
2. Enter a name for the authentication profile to authenticate RIP messages.
3. Select **Simple Password** as the **Password Type**.
4. Enter a simple password and then confirm.

MD5 RIP authentication

1. Select **Auth Profiles** and click **Add**.
2. Enter a name for the authentication profile to authenticate RIP messages.
3. Select **MD5** as the **Password Type**.
4. Click **Add**.
5. Enter one or more password entries, including:
 - Key-ID Range 0-255
 - Key
6. You can optionally select **Preferred** status.
7. Click **OK** to specify the key to be used to authenticate outgoing message.
8. Click **OK** again in the Virtual Router - RIP Auth Profile dialog box.

OSPF

Open Shortest Path First (OSPF) is an interior gateway protocol (IGP) that is most often used to dynamically manage network routes in large enterprise network. It determines routes dynamically by obtaining information from other routers and advertising routes to other routers by way of Link State Advertisements (LSAs). The information gathered from the LSAs is used to construct a topology map of the network. This topology map is shared across routers in the network and used to populate the IP routing table with available routes.

Changes in the network topology are detected dynamically and used to generate a new topology map within seconds. A shortest path tree is computed of each route. Metrics associated with each routing interface are used to calculate the best route. These can include distance, network throughput, link availability etc. Additionally, these metrics can be configured statically to direct the outcome of the OSPF topology map.

Palo Alto networks implementation of OSPF fully supports the following RFCs:

- RFC 2328 (for IPv4)
- RFC 5340 (for IPv6)

The following topics provide more information about the OSPF and procedures for configuring OSPF on the firewall:

- ▲ [OSPF Concepts](#)
- ▲ [Configure OSPF](#)
- ▲ [Configure OSPFv3](#)
- ▲ [Configure OSPF Graceful Restart](#)
- ▲ [Confirm OSPF Operation](#)

OSPF Concepts

The following topics introduce the OSPF concepts you will need to understand in order to configure the firewall to participate in an OSPF network:

- ▲ [OSPFv3](#)
- ▲ [OSPF Neighbors](#)
- ▲ [OSPF Areas](#)
- ▲ [OSPF Router Types](#)

OSPFv3

OSPFv3 provides support for the OSPF routing protocol within an IPv6 network. As such, it provides support for IPv6 addresses and prefixes. It retains most of the structure and functions in OSPFv2 (for IPv4) with some minor changes. The following are some of the additions and changes to OSPFv3:

- **Support for multiple instances per link**—With OSPFv3, you can run multiple instances of the OSPF protocol over a single link. This is accomplished by assigning an OSPFv3 instance ID number. An interface that is assigned to an instance ID drops packets that contain a different ID.
- **Protocol Processing Per-link**—OSPFv3 operates per-link instead of per-IP-subnet as on OSPFv2.
- **Changes to Addressing**—IPv6 addresses are not present in OSPFv3 packets, except for LSA payloads within link state update packets. Neighboring routers are identified by the Router ID.
- **Authentication Changes**—OSPFv3 doesn't include any authentication capabilities. Configuring OSPFv3 on a firewall requires an authentication profile that specifies Encapsulating Security Payload (ESP) or IPv6 Authentication Header (AH). The re-keying procedure specified in RFC 4552 is not supported in this release.
- **Support for multiple instances per-link**—Each instance corresponds to an instance ID contained in the OSPFv3 packet header.
- **New LSA Types**—OSPFv3 supports two new LSA types: Link LSA and Intra Area Prefix LSA.

All additional changes are described in detail in RFC 5340.

OSPF Neighbors

Two OSPF-enabled routers connected by a common network and in the same OSPF area that form a relationship are OSPF neighbors. The connection between these routers can be through a common broadcast domain or by a point-to-point connection. This connection is made through the exchange of hello OSPF protocol packets. These neighbor relationships are used to exchange routing updates between routers.

OSPF Areas

OSPF operates within a single autonomous system (AS). Networks within this single AS, however, can be divided into a number of areas. By default, Area 0 is created. Area 0 can either function alone or act as the OSPF backbone for a larger number of areas. Each OSPF area is named using a 32-bit identifier which in most cases is written in the same dotted-decimal notation as an IP4 address. For example, Area 0 is usually written as 0.0.0.0.

The topology of an area is maintained in its own link state database and is hidden from other areas, which reduces the amount of traffic routing required by OSPF. The topology is then shared in a summarized form between areas by a connecting router.

Table: OSPF Area Types

Area Type	Description
Backbone Area	The backbone area (Area 0) is the core of an OSPF network. All other areas are connected to it and all traffic between areas must traverse it. All routing between areas is distributed through the backbone area. While all other OSPF areas must connect to the backbone area, this connection doesn't need to be direct and can be made through a virtual link.
Normal OSPF Area	In a normal OSPF area there are no restrictions; the area can carry all types of routes.
Stub OSPF Area	A stub area does not receive routes from other autonomous systems. Routing from the stub area is performed through the default route to the backbone area.
NSSA Area	The Not So Stubby Area (NSSA) is a type of stub area that can import external routes, with some limited exceptions.

OSPF Router Types

Within an OSPF area, routers are divided into the following categories.

- **Internal Router**—A router with that has OSPF neighbor relationships only with devices in the same area.
- **Area Border Router (ABR)**—A router that has OSPF neighbor relationships with devices in multiple areas. ABRs gather topology information from their attached areas and distribute it to the backbone area.
- **Backbone Router**—A backbone router is any OSPF router that is attached to the OSPF backbone. Since ABRs are always connected to the backbone, they are always classified as backbone routers.
- **Autonomous System Boundary Router (ASBR)**—An ASBR is a router that attaches to more than one routing protocol and exchanges routing information between them.

Configure OSPF

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

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

Configure OSPF	
Step 1	Configure general virtual router configuration settings.
	See Virtual Routers for details.
Step 2	Configure general OSPF configuration settings.
	<ol style="list-style-type: none"> 1. Select the OSPF tab. 2. Select the Enable check box to enable the OSPF protocol. 3. Optionally enter the Router ID. 4. Select the Reject Default Route check box if you do not want to learn any default routes through OSPF. This is the recommended default setting. <p>De-select the Reject Default Route check box if you want to permit redistribution of default routes through OSPF.</p>

Configure OSPF (Continued)

Step 3 Configure Areas - Type for the OSPF protocol.

1. On the **Areas** tab, click **Add**.
 2. Enter an Area ID for the area in *x.x.x.x* format. This is the identifier that each neighbor must accept to be part of the same area.
 3. On the **Type** tab, select one of the following from the area **Type** drop-down:
 - **Normal**—There are no restrictions; the area can carry all types of routes.
 - **Stub**—There is no outlet from the area. To reach a destination outside of the area, it is necessary to go through the border, which connects to other areas. If you select this option, configure the following:
 - **Accept Summary**—Link state advertisements (LSA) are accepted from other areas. If this option on a stub area Area Border Router (ABR) interface is disabled, the OSPF area will behave as a Totally Stubby Area (TSA) and the ABR will not propagate any summary LSAs.
 - **Advertise Default Route**—Default route LSAs will be included in advertisements to the stub area along with a configured metric value in the configured range: 1-255.
 - **NSSA (Not-So-Stubby Area)**—The firewall can only leave the area by routes other than OSPF routes. If selected, configure **Accept Summary** and **Advertise Default Route** as described for **Stub**. If you select this option, configure the following:
 - **Type**—Select either **Ext 1** or **Ext 2** route type to advertise the default LSA.
 - **Ext Ranges**—Click **Add** in the section to enter ranges of external routes that you want to enable or suppress advertising for.
 4. **Priority**—Enter the OSPF priority for this interface (0-255). This is the priority for the router to be elected as a designated router (DR) or as a backup DR (BDR) according to the OSPF protocol. When the value is zero, the router will not be elected as a DR or BDR.
 - **Auth Profile**—Select a previously-defined authentication profile.
 - **Timing**—It is recommended that you keep the default timing settings.
 - **Neighbors**—For p2pmp interfaces, enter the neighbor IP address for all neighbors that are reachable through this interface.
 5. Select **normal**, **passive** or **send-only** as the **Mode**.
 6. Click **OK**.
-

Configure OSPF (Continued)	
<p>Step 4 Configure Areas - Range for the OSPF protocol</p>	<ol style="list-style-type: none"> 1. On the Range tab, click Add to aggregate LSA destination addresses in the area into subnets. 2. Advertise or Suppress advertising LSAs that match the subnet, and click OK. Repeat to add additional ranges.
<p>Step 5 Configure Areas - Interfaces for the OSPF protocol</p>	<ol style="list-style-type: none"> 1. On the Interface tab, click Add and enter the following information for each interface to be included in the area: <ul style="list-style-type: none"> • Interface—Select an interface from the drop-down box. • Enable—Selecting this option causes the OSPF interface settings to take effect. • Passive—Select the check box to if you do not want the OSPF interface to send or receive OSPF packets. Although OSPF packets are not sent or received if you choose this option, the interface is included in the LSA database. • Link type—Choose Broadcast if you want all neighbors that are accessible through the interface to be discovered automatically by multicasting OSPF hello messages, such as an Ethernet interface. Choose p2p (point-to-point) to automatically discover the neighbor. Choose p2mp (point-to-multipoint) when neighbors must be defined manually. Defining neighbors manually is allowed only for p2mp mode. • Metric—Enter an OSPF metric for this interface. Default: 10. Range: 0-65535. • Priority—Enter an OSPF priority for this interface. This is the priority for the router to be elected as a designated router (DR) or as a backup DR (BDR). Default: 1. Range: 0 - 255. If zero is configured, the router will not be elected as a DR or BDR. • Auth Profile—Select a previously-defined authentication profile. • Timing—The following OSPF timing settings can be set. Palo Alto Networks recommends that you retain the default timing settings. <ul style="list-style-type: none"> – Hello Interval (sec)—Interval at which the OSPF process sends hello packets to its directly connected neighbors. Default: 10 seconds. Range: 0-3600 seconds. – Dead Counts—Number of times the hello interval can occur for a neighbor without OSPF receiving a hello packet from the neighbor, before OSPF considers that neighbor down. The Hello Interval multiplied by the Dead Counts equals the value of the dead timer. Default: 4. Range: 3-20.

Configure OSPF (Continued)	
	<ul style="list-style-type: none"> – Retransmit Interval (sec)—Length of time that OSPF waits to receive a link state advertisement (LSA) from a neighbor before OSPF retransmits the LSA. Default: 10 seconds. Range: 0-3600 seconds. – Transit Delay (sec)—Length of time that an LSA is delayed before it is sent out of an interface. Default: 1 second. Range: 0-3600 seconds. – Graceful Restart Hello Delay (sec)—Applies to an OSPF interface when Active/Passive High Availability is configured. Graceful Restart Hello Delay is the length of time during which the firewall sends Grace LSA packets at 1-second intervals. During this time, no hello packets are sent from the restarting firewall. During the restart, the dead timer (which is the Hello Interval multiplied by the Dead Counts) is also counting down. If the dead timer is too short, the adjacency will go down during the graceful restart because of the hello delay. Therefore, it is recommended that the dead timer be at least four times the value of the Graceful Restart Hello Delay. For example, a Hello Interval of 10 seconds and a Dead Counts of 4 yield a dead timer of 40 seconds. If the Graceful Restart Hello Delay is set to 10 seconds, that 10-second delay of hello packets is comfortably within the 40-second dead timer, so the adjacency will not time out during a graceful restart. Default: 10 seconds. Range: 1-10 seconds. • If p2mp is selected for Link Type interfaces, enter the neighbor IP addresses for all neighbors that are reachable through this interface. <p>2. Click OK</p>
<p>Step 6 Configure Areas - Virtual Links.</p>	<p>1. On the Virtual Link tab, click Add and enter the following information for each virtual link to be included in the backbone area:</p> <ul style="list-style-type: none"> • Name—Enter a name for the virtual link. • Neighbor ID—Enter the router ID of the router (neighbor) on the other side of the virtual link. • Transit Area—Enter the area ID of the transit area that physically contains the virtual link. • Enable—Select to enable the virtual link. • Timing—It is recommended that you keep the default timing settings. • Auth Profile—Select a previously-defined authentication profile. <p>2. Click OK.</p>

Configure OSPF (Continued)	
<p>Step 7 (Optional) Configure Auth Profiles.</p>	<p>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.</p> <p>Simple Password OSPF authentication</p> <ol style="list-style-type: none"> 1. On the Auth Profiles tab, click Add. 2. Enter a name for the authentication profile to authenticate OSPF messages. 3. Select Simple Password as the Password Type. 4. Enter a simple password and then confirm. <p>MD5 OSPF authentication</p> <ol style="list-style-type: none"> 1. On the Auth Profiles tab, click Add. 2. Enter a name for the authentication profile to authenticate OSPF messages. 3. Select MD5 as the Password Type. 4. Click Add. 5. Enter one or more password entries, including: <ul style="list-style-type: none"> • Key-ID Range 0-255 • Key • Select the Preferred option to specify that the key be used to authenticate outgoing messages. 6. Click OK. 7. Click OK again in the Virtual Router - OSPF Auth Profile dialog box.
<p>Step 8 Configure Advanced OSPF options.</p>	<ol style="list-style-type: none"> 1. On the Advanced tab, select the RFC 1583 Compatibility check box to ensure compatibility with RFC 1583. 2. Configure a value for the SPF Calculation Delay (sec) timer. This timer allows you to tune the delay time between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should be tuned in a similar manner to optimize convergence times. 3. Configure a value for the LSA Interval (sec) time. This timer specifies the minimum time between transmissions of two instances of the same LSA (same router, same type, same LSA ID). This is equivalent to MinLSInterval in RFC 2328. Lower values can be used to reduce re-convergence times when topology changes occur.

Configure OSPFv3

Configure OSPFv3	
Step 1 Configure general virtual router configuration settings.	See Virtual Routers for details.
Step 2 Configure general OSPF configuration settings.	<ol style="list-style-type: none"> 1. Select the OSPF tab. 2. Select the Enable check box to enable the OSPF protocol. 3. Select the Reject Default Route check box if you do not want to learn any default routes through OSPF. This is the recommended default setting. 4. Clear the Reject Default Route check box if you want to permit redistribution of default routes through OSPF.
Step 3 Configure general OSPFv3 configuration settings.	<ol style="list-style-type: none"> 1. Select the OSPFv3 tab. 2. Select the Enable check box to enable the OSPF protocol. 3. Select the Reject Default Route check box if you do not want to learn any default routes through OSPFv3. This is the recommended default setting. De-select the Reject Default Route check box if you want to permit redistribution of default routes through OSPFv3.
Step 4 Configure Auth Profile for the OSPFv3 protocol. While OSPFv3 doesn't include any authentication capabilities of its own, it relies entirely on IPsec to secure communications between neighbors.	<p>When configuring an authentication profile, you must use Encapsulating Security Payload (ESP) or IPv6 Authentication Header (AH).</p> <p>ESP OSPFv3 authentication</p> <ol style="list-style-type: none"> 1. On the Auth Profiles tab, click Add. 2. Enter a name for the authentication profile to authenticate OSPFv3 messages. 3. Specify a Security Policy Index (SPI). The SPI must match between both ends of the OSPFv3 adjacency. The SPI number must be a hexadecimal value between 00000000 and FFFFFFFF. 4. Select ESP for Protocol. 5. Select a Crypto Algorithm from the drop-down. You can enter none or one of the following algorithms: SHA1, SHA256, SHA384, SHA512 or MD5. 6. If a Crypto Algorithm other than none was selected, enter a value for Key and then confirm.

Configure OSPFv3 (Continued)	
	<p>AH OSPFv3 authentication</p> <ol style="list-style-type: none"> 1. On the Auth Profiles tab, click Add. 2. Enter a name for the authentication profile to authenticate OSPFv3 messages. 3. Specify a Security Policy Index (SPI). The SPI must match between both ends of the OSPFv3 adjacency. The SPI number must be a hexadecimal value between 00000000 and FFFFFFFF. 4. Select AH for Protocol. 5. Select a Crypto Algorithm from the drop-down. You must enter one of the following algorithms: SHA1, SHA256, SHA384, SHA512 or MD5. 6. Enter a value for Key and then confirm. 7. Click OK. 8. Click OK again in the Virtual Router - OSPF Auth Profile dialog.
<p>Step 5 Configure Areas - Type for the OSPF protocol.</p>	<ol style="list-style-type: none"> 1. On the Areas tab, click Add. 2. Enter an Area ID. This is the identifier that each neighbor must accept to be part of the same area. 3. On the General tab, select one of the following from the area Type drop-down: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> – 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.

Configure OSPFv3 (Continued)	
	<ul style="list-style-type: none"> • NSSA (Not-So-Stubby Area)—The firewall can only leave the area by routes other than OSPF routes. If selected, configure Accept Summary and Advertise Default Route as described for Stub. If you select this option, configure the following: <ul style="list-style-type: none"> – Type—Select either Ext 1 or Ext 2 route type to advertise the default LSA. – Ext Ranges—Click Add in the section to enter ranges of external routes that you want to enable or suppress advertising for.
<p>Step 6 Associate an OSPFv3 authentication profile to an area or an interface.</p>	<p>To an Area</p> <ol style="list-style-type: none"> 1. On the Areas tab, select an existing area from the table. 2. On the General tab, select a previously defined Authentication Profile from the Authentication drop-down. 3. Click OK. <p>To an Interface</p> <ol style="list-style-type: none"> 1. On the Areas tab, select an existing area from the table. 2. Select the Interface tab and click Add. 3. Select the authentication profile you want to associate with the OSPF interface from the Auth Profile drop-down.
<p>Step 7 (Optional) Configure Export Rules</p>	<ol style="list-style-type: none"> 1. On the Export tab, click Add. 2. Select the Allow Redistribute Default Route check box to permit redistribution of default routes through OSPFv3. 3. Select the name of a redistribution profile. The value must be an IP subnet or valid redistribution profile name. 4. Select a metric to apply for New Path Type. 5. Specify a New Tag for the matched route that has a 32-bit value. 6. Assign a metric for the new rule. Range: 1 - 65535. 7. Click OK.

Configure OSPFv3 (Continued)	
<p>Step 8 Configure Advanced OSPFv3 options.</p>	<ol style="list-style-type: none"> 1. On the Advanced tab, select the Disable Transit Routing for SPF Calculation check box if you want the firewall to participate in OSPF topology distribution without being used to forward transit traffic. 2. Configure a value for the SPF Calculation Delay (sec) timer. This timer allows you to tune the delay time between receiving new topology information and performing an SPF calculation. Lower values enable faster OSPF re-convergence. Routers peering with the firewall should be tuned in a similar manner to optimize convergence times. 3. Configure a value for the LSA Interval (sec) time. This timer specifies the minimum time between transmissions of two instances of the same LSA (same router, same type, same LSA ID). This is equivalent to MinLSInterval in RFC 2328. Lower values can be used to reduce re-convergence times when topology changes occur. 4. Optionally, Configure OSPF Graceful Restart.

Configure OSPF Graceful Restart

OSPF Graceful Restart directs OSPF neighbors to continue using routes through a device during a short transition when it is out of service. This behavior increases network stability by reducing the frequency of routing table reconfiguration and the related route flapping that can occur during short periodic down times.

For a Palo Alto Networks firewall, OSPF Graceful Restart involves the following operations:

- **Firewall as a restarting device**—In a situation where the firewall will be down for a short period of time or is unavailable for short intervals, it sends Grace LSAs to its OSPF neighbors. The neighbors must be configured to run in Graceful Restart Helper mode. In Helper Mode, the neighbors receive the Grace LSAs that inform it that the firewall will perform a graceful restart within a specified period of time defined as the **Grace Period**. During the grace period, the neighbor continues to forward routes through the firewall and to send LSAs that announce routes through the firewall. If the firewall resumes operation before expiration of the grace period, traffic forwarding will continue as before without network disruption. If the firewall does not resume operation after the grace period has expired, the neighbors will exit helper mode and resume normal operation, which will involve reconfiguring the routing table to bypass the firewall.
- **Firewall as a Graceful Restart Helper**—In a situation where neighboring routers may be down for a short periods of time, the firewall can be configured to operate in Graceful Restart Helper mode. If configured in this mode, the firewall will be configured with a **Max Neighbor Restart Time**. When the firewall receives the Grace LSAs from its OSPF neighbor, it will continue to route traffic to the neighbor and advertise routes through the neighbor until either the grace period or max neighbor restart time expires. If neither expires before the neighbor returns to service, traffic forwarding continues as before without network disruption. If either period expires before the neighbor returns to service, the firewall will exit helper mode and resume normal operation, which will involve reconfiguring the routing table to bypass the neighbor.

Configure OSPF Graceful Restart

1. Select **Network > Virtual Routers** and select the virtual router you want to configure.
2. Select **OSPF > Advanced**.
3. Verify that the following check boxes are selected (they are enabled by default):
 - **Enable Graceful Restart**
 - **Enable Helper Mode**
 - **Enable Strict LSA checking**
 These check boxes should remain selected unless required by your topology.
4. Configure a **Grace Period** in seconds.
5. Configure a **Max Neighbor Restart Time** in seconds.

Confirm OSPF Operation

Once an OSPF configuration has been committed, you can use any of the following operations to confirm that OSPF is operating:

- ▲ [View the Routing Table](#)
- ▲ [Confirm OSPF Adjacencies](#)
- ▲ [Confirm that OSPF Connections are Established](#)

View the Routing Table

By viewing the routing table, you can see whether OSPF routes have been established. The routing table is accessible from either the web interface or the CLI. If you are using the CLI, use the following commands:

- `show routing route`
- `show routing fib`

The following procedure describes how to use the web interface to view the routing table.

View the Routing Table

1. Select **Network > Virtual Routers**.



Name	Interfaces	Configuration	RIP	OSPF	BGP	Multicast	Runtime Stats
<input checked="" type="checkbox"/> default	ethernet1/1 ethernet1/2 loopback.3			Enabled Area Count: 1 Subnet Count: 1 Neighbor Count: 0 Virtual Link Count: 0 Virtual Neighbor Count: 0			More Runtime Stats

View the Routing Table

- Select the **Routing** tab and examine the **Flags** column of the routing table for routes that were learned by OSPF.

Destination	Next Hop	Metric	Flags	Age	Interface
192.0.2.0/30	0.0.0.0	10	O	7515	ethernet1/1
192.0.2.0/30	192.0.2.2	0	A C		ethernet1/1
192.0.2.2/32	0.0.0.0	0	A H		
192.0.2.4/30	192.0.2.13	20	A O	7364	ethernet1/3
192.0.2.8/30	192.0.2.1	20	A O	7399	ethernet1/1
192.0.2.12/30	0.0.0.0	10	O	7515	ethernet1/3
192.0.2.12/30	192.0.2.14	0	A C		ethernet1/3
192.0.2.14/32	0.0.0.0	0	A H		
192.0.2.16/30	0.0.0.0	10	O	7515	ethernet1/2
192.0.2.16/30	192.0.2.17	0	A C		ethernet1/2
192.0.2.17/32	0.0.0.0	0	A H		
192.0.2.20/30	0.0.0.0	10	O	7515	ethernet1/4
192.0.2.20/30	192.0.2.21	0	A C		ethernet1/4
192.0.2.21/32	0.0.0.0	0	A H		

Confirm OSPF Adjacencies

By viewing the **Neighbor** tab as described in the following procedure, you can confirm that OSPF adjacencies have been established.

View the Neighbor Tab to Confirm OSPF Adjacencies

- Select **Network > Virtual Routers**.

Name	Interfaces	Configuration	RIP	OSPF	BGP	Multicast	Runtime Stats
default	ethernet1/1 ethernet1/2 loopback.3			Enabled Area Count: 1 Subnet Count: 1 Neighbor Count: 0 Virtual Link Count: 0 Virtual Neighbor Count: 0			More Runtime Stats

- Select **OSPF > Neighbor** and examine the **Status** column to determine if OSPF adjacencies have been established.

Neighbor Address	Neighbor Router Id	Local Address Binding	Area Id	Neighbor Priority	Remaining Lifetime	Status
192.0.2.1	192.0.2.1	0.0.0.0	0.0.0.0	1	30	full
192.0.2.13	192.0.2.13	0.0.0.0	0.0.0.0	1	30	full
192.0.2.18	192.0.2.18	0.0.0.0	0.0.0.0	1	34	full
192.0.2.22	192.0.2.22	0.0.0.0	0.0.0.0	1	33	full

Confirm that OSPF Connections are Established

By viewing the system log, you can confirm that OSPF connections have been established, as described in the following procedure:

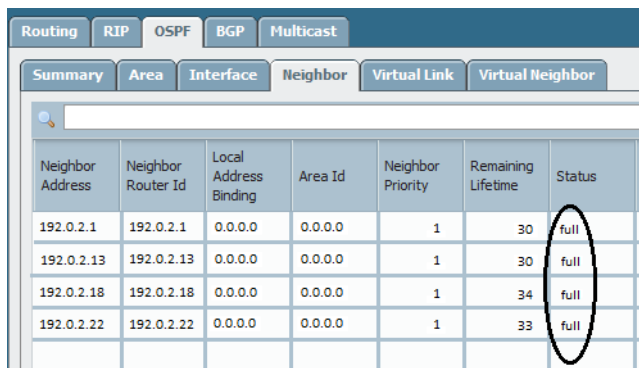
Examine the System Log

1. Select **Monitor > System** and look for messages to confirm that OSPF adjacencies have been established.



Name	Interfaces	Configuration	RIP	OSPF	BGP	Multicast	Runtime Stats
<input checked="" type="checkbox"/> default	ethernet1/1 ethernet1/2 loopback.3			Enabled Area Count: 1 Subnet Count: 1 Neighbor Count: 0 Virtual Link Count: 0 Virtual Neighbor Count: 0			More Runtime Stats

2. Select **OSPF > Neighbor** and examine the **Status** column to determine if OSPF adjacencies have been established.



Neighbor Address	Neighbor Router Id	Local Address Binding	Area Id	Neighbor Priority	Remaining Lifetime	Status
192.0.2.1	192.0.2.1	0.0.0.0	0.0.0.0	1	30	full
192.0.2.13	192.0.2.13	0.0.0.0	0.0.0.0	1	30	full
192.0.2.18	192.0.2.18	0.0.0.0	0.0.0.0	1	34	full
192.0.2.22	192.0.2.22	0.0.0.0	0.0.0.0	1	33	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.

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

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

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

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

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

BGP configuration consists of the following elements:

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

Perform the following procedure to configure BGP.

Configure BGP	
Step 1	Configure general virtual router configuration settings.
	See Virtual Routers for details.

Configure BGP (Continued)	
<p>Step 2 Configure standard BGP configuration settings.</p>	<ol style="list-style-type: none"> 1. Select the BGP tab. 2. Select the Enable check box to enable the BGP protocol. 3. Assign an IP address to the virtual router in the Router ID box. 4. Enter the number of the AS to which the virtual router belongs in the AS Number box, based on the router ID. Range: 1-4294967295
<p>Step 3 Configure general BGP configuration settings.</p>	<ol style="list-style-type: none"> 1. Select BGP > General. 2. Select the Reject Default Route check box to ignore any default routes that are advertised by BGP peers. 3. Select the Install Route check box to install BGP routes in the global routing table. 4. Select the Aggregate MED check box to enable route aggregation even when routes have different Multi-Exit Discriminator (MED) values. 5. Enter a value for the Default Local Preference that specifies a value that can be used to determine preferences among different paths. 6. Select one of the following values for the AS format for interoperability purposes: <ul style="list-style-type: none"> • 2 Byte (default value) • 4 Byte 7. Enable or disable each of the following values for Path Selection: <ul style="list-style-type: none"> • Always Compare MED—Enable this comparison to choose paths from neighbors in different autonomous systems. • Deterministic MED Comparison—Enable this comparison to choose between routes that are advertised by IBGP peers (BGP peers in the same autonomous system). 8. Click Add to include a new authentication profile and configure the following settings: <ul style="list-style-type: none"> • Profile Name—Enter a name to identify the profile. • Secret/Confirm Secret—Enter and confirm a passphrase for BGP peer communications.

Configure BGP (Continued)

Step 4 Configure BGP Advanced settings (Optional)

1. On the **Advanced** tab, select **Graceful Restart** and configure the following timers:
 - **Stale Route Time (sec)**—Specifies the length of time in seconds that a route can stay in the stale state. Range: 1 - 3600 seconds. Default: 120 seconds.
 - **Local Restart Time (sec)**—Specifies the length of time in seconds that the local device waits to restart. This value is advertised to peers. Range: 1 - 3600 seconds. Default: 120 seconds.
 - **Max Peer Restart Time (sec)**—Specifies the maximum length of time in seconds that the local device accepts as a grave period restart time for peer devices. Range: 1 - 3600 seconds. Default: 120 seconds.
2. Specify an IPv4 identifier to represent the reflector cluster in the **Reflector Cluster ID** box.
3. Specify the identifier for the AS confederation to be presented as a single AS to external BGP peers in the **Confederation Member AS** box.
4. Click **Add** and enter the following information for each Dampening Profile that you want to configure, select **Enable**, and click **OK**:
 - **Profile Name**—Enter a name to identify the profile.
 - **Cutoff**—Specify a route withdrawal threshold above which a route advertisement is suppressed. Range: 0.0-1000.0. Default: 1.25.
 - **Reuse**—Specify a route withdrawal threshold below which a suppressed route is used again. Range: 0.0-1000.0. Default: 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: 0-3600 seconds. Default: 900 seconds.
 - **Decay Half Life Reachable (sec)**—Specify the length of time in seconds after which a route's stability metric is halved if the route is considered reachable. Range: 0-3600 seconds. Default: 300 seconds.
 - **Decay Half Life Unreachable (sec)**—Specify the length of time in seconds after which a route's stability metric is halved if the route is considered unreachable. Range: 0 - 3600 seconds. Default: 300 seconds.
5. Click **OK**.

Configure BGP (Continued)

<p>Step 5 Configure the BGP peer group.</p>	<ol style="list-style-type: none"> 1. Select the Peer Group tab and click Add. 2. Enter a Name for the peer group and select Enable. 3. Select the Aggregated Confed AS Path check box to include a path to the configured aggregated confederation AS. 4. Select the Soft Reset with Stored Info check box to perform a soft reset of the firewall after updating the peer settings. 5. Specify the type of peer or group from the Type drop-down and configure the associated settings (see below in this table for descriptions of Import Next Hop and Export Next Hop). <ul style="list-style-type: none"> • IBGP—Export Next Hop: Specify Original or Use self • EBGP Confed—Export Next Hop: Specify Original or Use self • EBGP Confed—Export Next Hop: Specify Original or Use self • EBGP—Import Next Hop: Specify Original or Use self, Export Next Hop: Specify Resolve or Use self. Select Remove Private AS if you want to force BGP to remove private AS numbers. 6. Click OK to save.
<p>Step 6 Configure Import and Export rules.</p> <p>The import/export rules are used to import/export routes from/to other routers. For example, importing the default route from your Internet Service Provider.</p>	<ol style="list-style-type: none"> 1. Select the Import tab and then click Add and enter a name in the Rules field and select the Enable check box. 2. Click Add and select the Peer Group to which the routes will be imported from. 3. Click the Match tab and define the options used to filter routing information. You can also define the Multi-Exit Discriminator (MED) value and a next hop value to routers or subnets for route filtering. The MED option is an external metric that lets neighbors know about the preferred path into an AS. A lower value is preferred over a higher value. 4. Click the Action tab and define the action that should occur (allow/deny) based on the filtering options defined in the Match tab. If Deny is selected, no further options need to be defined. If the Allow action is selected, define the other attributes. 5. Click the Export tab and define export attributes, which are similar to the Import settings, but are used to control route information that is exported from the firewall to neighbors. 6. Click OK to save.

Configure BGP (Continued)	
<p>Step 7 Configure conditional advertising, which allows you to control what route to advertise in the event that a different route is not available in the local BGP routing table (LocRIB), indicating a peering or reachability failure. This is useful in cases where you want to try to force routes to one AS over another, for example if you have links to the Internet through multiple ISPs and you want traffic to be routed to one provider instead of the other unless there is a loss of connectivity to the preferred provider.</p>	<ol style="list-style-type: none"> 1. Select the Conditional Adv tab, click Add and enter a name in the Policy field. 2. Select the Enable check box. 3. Click Add and in the Used By section enter the peer group(s) that will use the conditional advertisement policy. 4. Select the Non Exist Filter tab and define the network prefix(es) of the preferred route. This specifies the route that you want to advertise, if it is available in the local BGP routing table. If a prefix is going to be advertised and matches a Non Exist filter, the advertisement will be suppressed. 5. Select the Advertise Filters tab and define the prefix(es) of the route in the Local-RIB routing table that should be advertised in the event that the route in the non-exist filter is not available in the local routing table. If a prefix is going to be advertised and does not match a Non Exist filter, the advertisement will occur.
<p>Step 8 Configure aggregate options to summaries routes in the BGP configuration.</p> <p>BGP route aggregation is used to control how BGP aggregates addresses. Each entry in the table results in one aggregate address being created. This will result in an aggregate entry in the routing table when at least one or more specific route matching the address specified is learned.</p>	<ol style="list-style-type: none"> 1. Select the Aggregate tab, click Add and enter a name for the aggregate address. 2. In the Prefix field, enter the network prefix that will be the primary prefix for the aggregated prefixes. 3. Select the Suppress Filters tab and define the attributes that will cause the matched routes to be suppressed. 4. Select the Advertise Filters tab and define the attributes that will cause the matched routes to always be advertised to peers.
<p>Step 9 Configure redistribution rules.</p> <p>This rule is used to redistribute host routes and unknown routes that are not on the local RIB to the peers routers.</p>	<ol style="list-style-type: none"> 1. Select the Redist Rules tab and click Add. 2. In the Name field, enter an IP subnet or select a redistribution profile. You can also configure a new redistribution profile from the drop-down if needed. 3. Click the Enable check box to enable the rule. 4. In the Metric field, enter the route metric that will be used for the rule. 5. In the Set Origin drop-down, select incomplete, igp, or egp. 6. Optionally set MED, local preference, AS path limit and community values.

Session Settings and Timeouts

This section describes the global settings that affect TCP, UDP, and ICMPv6 sessions, in addition to IPv6, NAT64, NAT oversubscription, jumbo frame size, MTU, accelerated aging, and captive portal authentication. There is also a setting (Rematch Sessions) that allows you to apply newly configured security policies to sessions that are already in progress.

The **Device > Setup > Session** tab is where these session settings and timeouts are configured.

The screenshot shows the Palo Alto Networks configuration interface. The top navigation bar includes tabs for Dashboard, ACC, Monitor, Policies, Objects, Network, and Device. The left sidebar shows the Setup menu with various configuration options. The main content area is divided into two panels: Session Settings and Session Timeouts.

Session Settings	
Rematch Sessions	<input checked="" type="checkbox"/>
ICMPv6 Token Bucket Size	100
ICMPv6 Error Packet Rate (per sec)	100
IPv6 Firewalling	<input checked="" type="checkbox"/>
Enable Jumbo Frame	<input type="checkbox"/>
Global MTU	1500
NAT64 IPv6 Minimum Network MTU	1280
NAT Oversubscription Rate	1x
Accelerated Aging	<input checked="" type="checkbox"/>
Accelerated Aging Threshold	80
Accelerated Aging Scaling Factor	2

Session Timeouts	
Default (sec)	30
Discard Default (sec)	60
Discard TCP (sec)	90
Discard UDP (sec)	60
ICMP (sec)	6
Scan (sec)	10
TCP (sec)	3600
TCP handshake (sec)	10
TCP init (sec)	5
TCP Half Closed (sec)	120
TCP Time Wait (sec)	15
Unverified RST (sec)	30
UDP (sec)	30
Captive Portal (sec)	30

Below the Session Settings panel, there is a Decryption Settings section with the following options:

- Certificate Revocation Checking
- Forward Proxy Server Certificate Settings

The first few topics below provide brief summaries of the Transport Layer of the OSI model, TCP, UDP, and ICMP. For more information about the protocols, refer to their respective RFCs. The remaining topics describe the session timeouts and settings.

- ▲ [Transport Layer Sessions](#)
- ▲ [TCP](#)
- ▲ [UDP](#)
- ▲ [ICMP](#)
- ▲ [Configure Session Timeouts](#)
- ▲ [Configure Session Settings](#)

Transport Layer Sessions

A network session is a conversation that occurs between two or more communication devices, lasting for some period of time. A session is established and is torn down when the session ends. Different types of sessions occur at three layers of the OSI model: the Transport layer, the Session layer, and the Application layer.

The Transport Layer operates at Layer 4 of the OSI model, providing reliable or unreliable, end-to-end delivery and flow control of data. Internet protocols that implement sessions at the Transport layer include Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).

TCP

Transmission Control Protocol (TCP) ([RFC 793](#)) is one of the main protocols in the Internet Protocol (IP) suite, and is so prevalent that it is frequently referenced together with IP as *TCP/IP*. TCP is considered a reliable transport protocol because it provides error-checking while transmitting and receiving segments, it acknowledges segments received, and it reorders segments that arrive out of order from the device that sent them. TCP also requests and provides retransmission of segments that were dropped. TCP is stateful and connection-oriented, in that a connection between the sender and receiver is established for the duration of the session. TCP provides flow control of packets, so it can handle congestion over networks.

TCP performs a three-way handshake during session setup to initiate and acknowledge a session. After the data is transferred, the session is closed in an orderly manner.

Applications that use TCP as their transport protocol include Hypertext Transfer Protocol (HTTP), Hypertext Transfer Protocol Secure (HTTPS), File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), Telnet, Post Office Protocol version 3 (POP3), Internet Message Access Protocol (IMAP), and Secure Shell (SSH).

The following topics describe details of the PAN-OS implementation of TCP.

▲ [TCP Half Closed and TCP Time Wait Timers](#)

▲ [Unverified RST Timer](#)

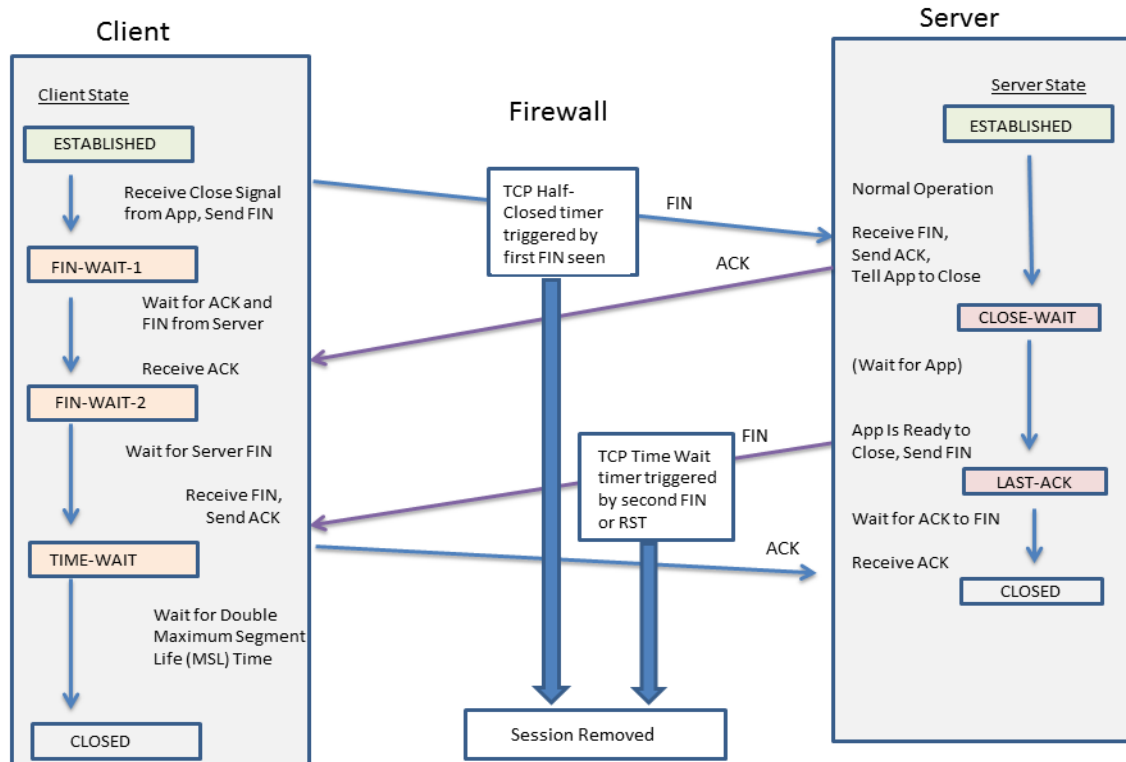
You can use zone protection profiles on the firewall to configure packet-based attack protection and thereby drop IP, TCP, and IPv6 packets with undesirable characteristics or strip undesirable options from packets before allowing them into the zone. You can also configure flood protection, specifying 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.

UDP

User Datagram Protocol (UDP) ([RFC 768](#)) is another main protocol of the IP suite, and is an alternative to TCP. UDP is stateless and connectionless in that there is no handshake to set up a session, and no connection between the sender and receiver; the packets may take different routes to get to a single destination. UDP is considered an unreliable protocol because it does not provide acknowledgments, error-checking, retransmission, or reordering of datagrams. Without the overhead required to provide those features, UDP has reduced latency and is faster than TCP. UDP is referred to as a best-effort protocol because there is no mechanism or guarantee to ensure that the data will arrive at its destination.

Although UDP uses a checksum for data integrity, it performs no error checking at the network interface level. Error checking is assumed to be unnecessary or is performed by the application rather than UDP itself. UDP has no mechanism to handle flow control of packets.

UDP is often used for applications that require faster speeds and time-sensitive, real-time delivery, such as Voice over IP (VoIP), streaming audio and video, and online games. UDP is transaction-oriented, so it is also used for applications that respond to small queries from many clients, such as Domain Name System (DNS) and Trivial File Transfer Protocol (TFTP).

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. (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 do not flood the network segments protected by the firewall.

First the global **ICMPv6 error packet rate (per sec)** controls the rate at which ICMP error packets are allowed through the firewall; the default is 100 packets per second; the range is 10 to 65535 packets per second. If the firewall reaches the ICMP error packet rate, then the token bucket comes into play and throttling occurs, as follows.

The concept of a logical token bucket controls the rate at which ICMP messages can be transmitted. The number of tokens in the bucket is configurable, and each token represents an ICMP message that can be sent. The token count is decremented each time an ICMP message is sent; when the bucket reaches zero tokens, no more ICMP messages can be sent until another token is added to the bucket. The default size of the token bucket is 100 tokens (packets); the range is 10 to 65535 tokens.

To change the default token bucket size or error packet rate, see the section [Configure Session Settings](#).

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.

Control Specific ICMP or ICMPv6 Types and Codes	
Step 1 Create a custom application for ICMP or ICMPv6 message types and codes.	<ol style="list-style-type: none"> 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.	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.

On the firewall, you can define a number of timeouts for TCP, UDP, and ICMP sessions in particular. The Default timeout applies to any other type of session. All of these timeouts are global, meaning they apply to all of the sessions of that type on the firewall.

In addition to the global settings, you have the flexibility to define timeouts for an individual application in the **Objects > Applications** tab. The firewall applies application timeouts to an application that is in established state. When configured, timeouts for an application override the global TCP or UDP session timeouts.

Returning to the global settings, perform the optional tasks below if you need to change default values of the global session timeout settings for TCP, UDP, ICMP, Captive Portal authentication, or other types of sessions. All values are in seconds.




The defaults are optimal values. However, you can modify these according to your network needs. Setting a value too low could cause sensitivity to minor network delays and could result in a failure to establish connections with the firewall. Setting a value too high could delay failure detection.

Change Session Timeouts	
<p>Step 1 Access the Session Settings.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Session. 2. In the Session Timeouts section, click the Edit icon.
<p>Step 2 (Optional) Change miscellaneous timeouts.</p>	<ul style="list-style-type: none"> • Default—Maximum length of time that a non-TCP/UDP or non-ICMP session can be open without a response. Default: 30. Range: 1-15,999,999. • Discard Default—Maximum length of time that a non-TCP/UDP session remains open after PAN-OS denies a session based on security policies configured on the firewall. Default: 60. Range: 1-15,999,999. • Scan—Maximum length of time that any session remains open after it is considered inactive; an application is regarded as inactive when it exceeds the application trickling threshold defined for the application. Default: 10. Range: 5-30. • Captive Portal—Authentication session timeout for the Captive Portal web form. To access the requested content, the user must enter the authentication credentials in this form and be successfully authenticated. Default: 30. Range: 1-15,999,999. <p>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.</p>

Change Session Timeouts (Continued)	
<p>Step 3 (Optional) Change TCP timeouts.</p>	<ul style="list-style-type: none"> • 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: 1-15,999,999. • 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: 3600. Range: 1-15,999,999. • TCP Handshake—Maximum length of time permitted between receiving the SYN-ACK and the subsequent ACK to fully establish the session. Default: 10. Range: 1-60. • TCP init—Maximum length of time permitted between receiving the SYN and SYN-ACK prior to starting the TCP handshake timer. Default: 5. Range: 1-60. • TCP Half Closed—Maximum length of time between receiving the first FIN and receiving the second FIN or a RST. Default: 120. Range: 1-604,800. • TCP Time Wait—Maximum length of time after receiving the second FIN or a RST. Default: 15. Range: 1-600. • Unverified RST—Maximum length of time after receiving a RST that cannot be verified (the RST is within the TCP window but has an unexpected sequence number, or the RST is from an asymmetric path). Default: 30. Range: 1-600. • See also the Scan timeout in the section (Optional) Change miscellaneous timeouts.
<p>Step 4 (Optional) Change UDP timeouts.</p>	<ul style="list-style-type: none"> • 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: 1-15,999,999. • UDP—Maximum length of time that a UDP session remains open without a UDP response. Default: 30. Range: 1-1599999. • See also the Scan timeout in the section (Optional) Change miscellaneous timeouts.
<p>Step 5 (Optional) Change ICMP timeouts.</p>	<ul style="list-style-type: none"> • ICMP—Maximum length of time that an ICMP session can be open without an ICMP response. Default: 6. Range: 1-15,999,999. • See also the Discard Default and Scan timeout in the section (Optional) Change miscellaneous timeouts.
<p>Step 6 Commit the changes.</p>	<p>Click OK and Commit the changes.</p>

Configure Session Settings

This topic describes various settings for sessions other than timeouts values. Perform these optional tasks if you need to change the default settings.

Change Session Settings	
Step 1 Access the Session Settings.	<ol style="list-style-type: none"> 1. Select Device > Setup > Session. 2. In the Session Settings section, click the Edit icon.
Step 2 (Optional) Change session settings.	<ul style="list-style-type: none"> • Rematch Sessions—Causes the firewall to apply newly configured security policies to sessions that are already in progress. This capability is enabled by default. If this setting is disabled, any policy change applies only to sessions initiated after the policy change was committed. For example, if a Telnet session started while an associated policy was configured that allowed Telnet, and you subsequently committed a policy change to deny Telnet, the firewall applies the revised policy to the current session and blocks it. • 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. • Enable Jumbo Frame—Enables jumbo frame support on Ethernet interfaces. Jumbo frames have a maximum transmission unit (MTU) of 9216 bytes and are available on certain platforms. Global MTU—Establishes the maximum transmission unit (MTU) available globally; this field is tied to the Enable Jumbo Frame field. <ul style="list-style-type: none"> • If you do not check Enable Jumbo Frame, the Global MTU defaults to 1500 bytes; the range is 576 to 1500 bytes. • If you check Enable Jumbo Frame, the Global MTU defaults to 9192 bytes; the range is 9192 to 9216 bytes. <p> 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.</p> • NAT64 IPv6 Minimum Network MTU—Sets the global MTU for IPv6 translated traffic. The default of 1280 bytes is based on the standard minimum MTU for IPv6 traffic.

Change Session Settings (Continued)

	<ul style="list-style-type: none"> • NAT Oversubscription Rate—If NAT is configured to be Dynamic IP and Port (DIPP) translation, an oversubscription rate can be configured to multiply the number of times that the same translated IP address and port pair can be used concurrently. The rate is 1, 2, 4, or 8. The default setting is based on the firewall platform. <ul style="list-style-type: none"> • A rate of 1 means no oversubscription; each translated IP address and port pair can be used only once at a time. • If the setting is Platform Default, user configuration of the rate is disabled and the default oversubscription rate for the platform applies. <p>Reducing the oversubscription rate decreases the number of source device translations, but provides higher NAT rule capacities.</p> • Accelerated Aging—Enables faster aging-out of idle sessions. Select the check box to enable accelerated aging, and, if necessary, change the threshold (%) and scaling factor. <ul style="list-style-type: none"> • 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. <p>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.</p>
Step 3 Commit the changes.	Click OK and Commit the changes.
Step 4 Reboot the firewall after changing the jumbo frame configuration.	<ol style="list-style-type: none"> 1. Select Device > Setup > Operation. 2. Click Reboot Device.

DHCP

This section describes Dynamic Host Configuration Protocol (DHCP) and the tasks required to configure an interface on a Palo Alto Networks firewall to act as a DHCP server, client, or relay agent. By assigning these roles to different interfaces, the firewall can perform multiple roles.

- ▲ [DHCP Overview](#)
- ▲ [DHCP Messages](#)
- ▲ [DHCP Addressing](#)
- ▲ [DHCP Options](#)
- ▲ [Restrictions to Palo Alto Networks Implementation of DHCP](#)
- ▲ [Configure an Interface as a DHCP Server](#)
- ▲ [Configure an Interface as a DHCP Client](#)
- ▲ [Configure an Interface as a DHCP Relay Agent](#)
- ▲ [Monitor and Troubleshoot DHCP](#)

DHCP Overview

DHCP is a standardized protocol defined in [RFC 2131, Dynamic Host Configuration Protocol](#). DHCP has two main purposes: to provide TCP/IP and link-layer configuration parameters and to provide network addresses to dynamically configured hosts on a TCP/IP network.

DHCP uses a client-server model of communication. This model consists of three roles that the device can fulfill: DHCP client, DHCP server, and DHCP relay agent.

- A device acting as a DHCP client (host) can request an IP address and other configuration settings from a DHCP server. Users on client devices save configuration time and effort, and need not know the network's addressing plan or other resources and options they are inheriting from the DHCP server.
- A device acting as a DHCP server can service clients. By using any of three [DHCP Addressing](#) mechanisms, the network administrator saves configuration time and has the benefit of reusing a limited number of IP addresses when a client no longer needs network connectivity. The server can deliver IP addressing and many DHCP options to many clients.
- A device acting as a DHCP relay agent transmits DHCP messages between DHCP clients and servers.

DHCP uses [User Datagram Protocol \(UDP\), RFC 768](#), as its transport protocol. DHCP messages that a client sends to a server are sent to well-known port 67 (UDP—Bootstrap Protocol and DHCP). [DHCP Messages](#) that a server sends to a client are sent to port 68.

An interface on a Palo Alto Networks firewall can perform the role of a DHCP server, client, or relay agent. The interface of a DHCP server or relay agent must be a Layer 3 Ethernet, Aggregated Ethernet, or Layer 3 VLAN interface. The network administrator configures the firewall's interfaces with the appropriate settings for any combination of roles. The configuration tasks for each role are provided in this document.

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: <ul style="list-style-type: none"> Request parameters from one server and implicitly decline offers from other servers. Confirm that a previously allocated address is correct after, for example, a system reboot. Extend the lease of a network address.
DHCPACK	Server to client acknowledgment message containing configuration parameters, including a confirmed network address.
DHCPNAK	Server to client negative acknowledgment indicating the client's understanding of the network address is incorrect (for example, if the client has moved to a new subnet), or a client's lease has expired.
DHCPDECLINE	Client to server message indicating the network address is already being used.
DHCPRELEASE	Client to server message giving up the user of the network address and canceling the remaining time on the lease.
DHCPINFORM	Client to server message requesting only local configuration parameters; client has an externally configured network address.

DHCP Addressing

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 do not want to wait for expired leases to be released automatically, you can use the `clear dhcp lease interface value expired-only` command to clear expired leases, making those addresses available in the pool again. You can use the `clear dhcp lease interface value ip ip` command to release a particular IP address. Use the `clear dhcp lease interface value mac mac_address` command to release a particular MAC address.

DHCP Options

The history of DHCP and DHCP options traces back to the Bootstrap Protocol (BOOTP). BOOTP was used by a host to configure itself dynamically during its booting procedure. A host could receive an IP address and a file from which to download a boot program from a server, along with the server's address and the address of an Internet gateway.

Included in the BOOTP packet was a vendor information field, which could contain a number of tagged fields containing various types of information, such as the subnet mask, the BOOTP file size, and many other values. [RFC 1497](#) describes the *BOOTP Vendor Information Extensions*.

These extensions eventually expanded with the use of DHCP and DHCP host configuration parameters, also known as options. Similar to vendor extensions, DHCP options are tagged data items that provide information to a DHCP client. The options are sent in a variable-length field at the end of a DHCP message. For example, the DHCP Message Type is option 53, and a value of 1 indicates the DHCPDISCOVER message. DHCP options are defined in [RFC 2132, DHCP Options and BOOTP Vendor Extensions](#).

A DHCP client can negotiate with the server, limiting the server to send only those options that the client requests.

The Palo Alto Networks firewall supports a subset of the standard DHCP options for its implementation of DHCP server. Such options are configured on the DHCP server and sent to the clients that sent a DHCPREQUEST to the server. The clients are said to *inherit* and implement the options that they are programmed to accept. The firewall supports the following predefined options on its DHCP servers, shown in the order in which they appear on the **DHCP Server** configuration screen:

DHCP Option	DHCP Option Name
51	Lease duration
3	Gateway
1	IP Pool Subnet (mask)
6	Domain Name System (DNS) server address (primary and secondary)
44	Windows Internet Name Service (WINS) server address (primary and secondary)

DHCP Option	DHCP Option Name
41	Network Information Service (NIS) server address (primary and secondary)
42	Network Time Protocol (NTP) server address (primary and secondary)
70	Post Office Protocol Version 3 (POP3) server address
69	Simple Mail Transfer Protocol (SMTP) server address
15	DNS suffix

Restrictions to Palo Alto Networks Implementation of DHCP

- The DHCP standard, [RFC 2131](#), is designed to support IPv4 and IPv6 addresses. The Palo Alto Networks implementations of DHCP server and DHCP client support IPv4 addresses only. Its DHCP relay implementation supports IPv4 and IPv6.
- DHCP client is not supported in High Availability active/active mode.
- The firewall supports DHCPv4 Server and DHCPv6 Relay. However, a single interface cannot support both DHCPv4 Server and DHCPv6 Relay.
- DHCP replaces BOOTP; BOOTP is not supported

Configure an Interface as a DHCP Server

The prerequisites for this task are:

- You have configured a Layer 3 Ethernet or Layer 3 VLAN interface.
- The interface is assigned to a virtual router and a zone.
- You know a valid pool of IP addresses from your network plan that can be designated to be assigned by your DHCP server to clients.

Perform the following task to configure an interface on the firewall to act as a DHCP server. You may configure multiple DHCP servers.

Configure an Interface as a DHCP Server

Step 1 Select an interface to be a DHCP Server.

1. Select **Network > DHCP > DHCP Server** and click **Add**.
2. Enter an **Interface** name or select one from the drop-down.
3. For **Mode**, select **enabled** or **auto** mode. Auto mode enables the server and disables it if another DHCP server is detected on the network. The **disabled** setting disables the server.
4. Optionally click the **Ping IP when allocating new IP** check box if you want the server to ping the IP address before it assigns that address to its client.



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 setting options and returning to the DHCP server tab, the **Probe IP** column for the interface indicates if this check box was checked.



Configure an Interface as a DHCP Server (Continued)

Step 2 Configure the [DHCP Options](#) that the server sends to its clients.

- In the Options section, select a **Lease** type:
 - **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.

Benefits of specifying an inheritance source are that it quickly adds DHCP options from the upstream server of the source DHCP client. It also keeps the client's options updated if the source changes an option. For example, if the source replaces its NTP server (which had been identified as the **Primary NTP** server), the client will automatically inherit the new address as its **Primary NTP** server.
- **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.
- **Ippool Subnet**—Network mask that applies to the addresses in the **IP Pools**.

Configure an Interface as a DHCP Server (Continued)

	<p>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.</p> <p>The DHCP server sends these settings to its clients.</p> <ul style="list-style-type: none"> • Primary DNS, Secondary DNS—IP address of the preferred and alternate Domain Name System (DNS) servers. • Primary WINS, Secondary WINS—IP address of the preferred and alternate Windows Internet Naming Service (WINS) servers. • Primary NIS, Secondary NIS—IP address of the preferred and alternate Network Information Service (NIS) servers. • Primary NTP, Secondary NTP—IP address of the available Network Time Protocol servers. • POP3 Server—IP address of a Post Office Protocol (POP3) server. • SMTP Server—IP address of a Simple Mail Transfer Protocol (SMTP) server. • DNS Suffix—Suffix for the client to use locally when an unqualified hostname is entered that it cannot resolve.
<p>Step 3 Identify the stateful pool of IP addresses from which the DHCP server chooses an address and assigns it to a DHCP client.</p> <p> 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.</p>	<ol style="list-style-type: none"> 1. In the IP Pools field, click Add and enter the range of IP addresses from which this server assigns an address to a client. Enter an IP subnet and subnet mask (for example, 192.168.1.0/24) or a range of IP addresses (for example, 192.168.1.10-192.168.1.20). At least one IP pool is required. 2. Optionally repeat Step 1 to specify another IP address pool.
<p>Step 4 (Optional) Specify an IP address from the IP pools that will not be assigned dynamically. If you also specify a MAC Address, the Reserved Address is assigned to that device when the device requests an IP address through DHCP.</p> <p> See the DHCP Addressing section for an explanation of allocation of a Reserved Address.</p>	<ol style="list-style-type: none"> 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. Optionally specify the MAC Address (format xx:xx:xx:xx:xx:xx) of the device to which you want to permanently assign the IP address specified in Step 2. 4. Optionally repeat Steps 2-3 to reserve another address.
<p>Step 5 Save the configuration.</p>	<p>Click OK and Commit the change.</p>

Configure an Interface as a DHCP Client

Before configuring a firewall interface as a DHCP client, make sure you have configured a Layer 3 Ethernet or Layer 3 VLAN interface, and the interface is assigned to a virtual router and a zone. Perform this task if you need to use DHCP to request an IPv4 address for an interface on your firewall.

Configure an Interface as a DHCP Client	
<p>Step 1 Configure an interface as a DHCP client.</p>	<ol style="list-style-type: none"> Select Network > Interfaces. On the Ethernet tab or the VLAN tab, click Add and enter an interface, or click a configured interface, that you want to be a DHCP client. The following shows the Ethernet Interface screen: <div data-bbox="685 661 1377 1018" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> </div> Click the IPv4 tab; for Type, select DHCP Client. Select the Enable check box. Optionally select the Automatically create default route pointing to default gateway provided by server check box. This causes the firewall to create a static route to a default gateway that will be useful when clients are trying to access many destinations that do not need to have routes maintained in a routing table on the firewall. Optionally enter a Default Route Metric (priority level) for the route between the firewall and the DHCP server. A route with a lower number has higher priority during route selection. For example, a route with a metric of 10 is used before a route with a metric of 100. Range is 1-65535. There is no default metric. Optionally click Show DHCP Client Runtime Info to see all of the settings the client has inherited from its DHCP server.
<p>Step 2 Save the configuration.</p>	<p>Click OK and Commit the change.</p> <p>Now the Ethernet interface indicates Dynamic-DHCP Client in its IP Address field on the Ethernet tab.</p>
<p>Step 3 (Optional) See which interfaces on the firewall are configured as DHCP clients.</p>	<ol style="list-style-type: none"> Select Network > Interfaces > Ethernet and look in the IP Address field to see which interfaces indicate DHCP Client. Select Network > Interfaces > VLAN and look in the IP Address field to see which interfaces indicate DHCP Client.

Configure an Interface as a DHCP Relay Agent

Before configuring a firewall interface as a DHCP relay agent, make sure you have configured a Layer 3 Ethernet or Layer 3 VLAN interface, and the interface is assigned to a virtual router and a zone. You want that interface be able to pass DHCP messages between clients and servers. The interface can forward messages to a maximum of four external DHCP servers. A client DHCPDISCOVER message is sent to all configured servers, and the DHCPOFFER message of the first server that responds is relayed back to the requesting client.

Perform the following task to configure an interface on the firewall to act as a DHCP relay agent:

Configure an Interface as a DHCP Relay Agent	
Step 1 Select DHCP Relay.	Select Network > DHCP > DHCP Relay .
Step 2 Specify the IP address of each DHCP server with which the DHCP relay agent will communicate.	<ol style="list-style-type: none"> 1. In the Interface field, select from the drop-down the interface you want to be the DHCP relay agent. 2. Check either the IPv4 or IPv6 check box, indicating the type of DHCP server address you will specify. 3. If you checked IPv4, in the DHCP Server IP Address field, click Add. Enter the address of the DHCP server to and from which you will relay DHCP messages. 4. If you checked IPv6, in the DHCP Server IPv6 Address field, click Add. Enter the address of the DHCP server to and from which you will relay DHCP messages. If you specify a <i>multicast</i> address, also specify an outgoing Interface. 5. Optionally repeat Steps 2-4 to enter a maximum of four DHCP server addresses.
Step 3 Save the configuration.	Click OK and Commit the change.

Monitor and Troubleshoot DHCP

You can view the status of dynamic address leases that your DHCP server has assigned or that your DHCP client has been assigned by issuing commands from the CLI. You can also clear leases before they time out and are released automatically. For complete command syntax, see the [PAN-OS Command Line Interface Reference Guide](#).

- ▲ [View DHCP Server Information](#)
- ▲ [Clear Leases Before They Expire Automatically](#)
- ▲ [View DHCP Client Information](#)
- ▲ [Gather Debug Output about DHCP](#)

View DHCP Server Information

To view DHCP pool statistics, IP addresses the server has assigned, the corresponding MAC address, state and duration of the lease, and time the lease began, use the following command. If the address was configured as a **Reserved Address**, the `state` column indicates `reserved` and there is no `duration` or `lease_time`. If the lease was configured as **Unlimited**, the `duration` column displays a value of 0.

```
admin@PA-200> show dhcp server lease all
interface: "ethernet1/2"
Allocated IPs: 1, Total number of IPs in pool: 5. 20.0000% used
ip           mac           state      duration    lease_time
192.168.3.11 f0:2f:af:42:70:cf committed  0           Wed Jul  2 08:10:56 2014
admin@PA-200>
```

To view the options that a DHCP server has assigned to clients, use the following command:

```
admin@PA-200> show dhcp server settings all
Interface      GW           DNS1         DNS2         DNS-Suffix   Inherit source
-----
ethernet1/2    192.168.3.1 10.43.2.10  10.44.2.10
admin@PA-200>
```

Clear Leases Before They Expire Automatically

The following example shows how to release expired [DHCP Leases](#) of an interface (server) before the hold timer releases them automatically. Those addresses will be available in the IP pool again.

```
admin@PA-200> clear dhcp lease interface ethernet1/2 expired-only
```

The following example shows how to release the lease of a particular IP address:

```
admin@PA-200> clear dhcp lease interface ethernet1/2 ip 192.168.3.1
```

The following example shows how to release the lease of a particular MAC address:

```
admin@PA-200> clear dhcp lease interface ethernet1/2 mac f0:2c:ae:29:71:34
```

View DHCP Client Information

To view the status of IP address leases sent to the firewall when it is acting as a DHCP client, use the `show dhcp client state interface_name` command or the following command:

```
admin@PA-200> show dhcp client state all
Interface      State      IP           Gateway      Leased-until
-----
ethernet1/1    Bound     10.43.14.80  10.43.14.1  70315
admin@PA-200>
```

Gather Debug Output about DHCP

To gather debug output about DHCP, use one of the following commands:

```
admin@PA-200> debug dhcpd
admin@PA-200> debug management-server dhcpd
```

NAT

This section describes Network Address Translation (NAT) and how to configure NAT rules and features.

- ▲ [Purpose of NAT](#)
- ▲ [NAT Rules and Security Policies](#)
- ▲ [Source NAT and Destination NAT](#)
- ▲ [NAT Rule Capacities](#)
- ▲ [Dynamic IP and Port NAT Oversubscription](#)
- ▲ [Dataplane NAT Memory Statistics](#)
- ▲ [Configure NAT](#)

Purpose of NAT

NAT was introduced to solve the problem of an organization not having enough public, globally-routable IPv4 addresses assigned to it by the Internet Assigned Numbers Authority (IANA) or a Regional Internet Registry. NAT translates private, non-routable IPv4 addresses to one or more globally-routable IPv4 addresses, thereby conserving an organization's routable IP addresses.

Since its origination, the use of NAT has expanded. For example, NAT is used as a way to not disclose the real IP addresses of hosts that need access to public addresses. It is also used to manage traffic by performing port forwarding. The [NAT64](#) option translates between IPv6 and IPv4 addresses, providing connectivity between networks using disparate IP addressing schemes, and therefore a migration path to IPv6 addressing. PAN-OS supports all of these uses.

If you use private IP addresses within your internal networks, you must use NAT to translate the private addresses to public addresses that can be routed on external networks. In PAN-OS, you create NAT policy rules that instruct the firewall which packet addresses need translation and what the translated addresses are.

NAT Rules and Security Policies

You configure a NAT rule to match a packet's source zone and destination zone, at a minimum. In addition to zones, you can configure matching criteria based on the packet's destination interface, source and destination address, and service. You can configure multiple NAT rules. The firewall evaluates the rules in order from the top down. Once a packet matches the criteria of a single NAT rule, the packet is not subjected to additional NAT rules. Therefore, your list of NAT rules should be in order from most specific to least specific so that packets are subjected to the most specific rule you created for them.

It is important to understand the firewall's flow logic when it applies NAT rules and security policies so that you can determine what rules you need, based on the zones you have defined.

Upon ingress, the firewall inspects the packet and does a route lookup to determine the egress interface and zone. Then the firewall determines if the packet matches one of the NAT rules that have been defined, based on source and/or destination zone. It then evaluates and applies any security policies that match the packet based on the original (pre-NAT) source and destination addresses, but the post-NAT zones. Finally, upon egress, for a matching NAT rule, the firewall translates the source and/or destination address and port numbers.

Keep in mind that the translation of the IP address and port do not occur until the packet leaves the firewall. The NAT rules and security policies apply to the original IP address (the pre-NAT address). A NAT rule is configured based on the zone associated with a pre-NAT IP address.

Security policies differ from NAT rules because security policies examine post-NAT zones to determine whether the packet is allowed or not. Because the very nature of NAT is to modify source or destination IP addresses, which can result in modifying the packet's outgoing interface and zone, security policies are enforced on the post-NAT zone.

Address Pools Identified as Address Objects

When configuring a NAT IP address pool, it is typical to identify it as address object. It can be a host IP address, IP address range, or IP subnet. Because both NAT rules and security policies use address objects, it is a best practice to distinguish between them by naming an address object used for NAT with a prefix, such as "NAT-*name*."

Source NAT and Destination NAT

The firewall supports both source address and/or port translation and destination address and/or port translation.

- ▲ Source NAT
- ▲ Destination NAT

Source NAT

Source NAT is typically used by internal users to access the Internet; the source address is translated and thereby kept private. There are three types of source NAT:

- **Dynamic IP and Port (DIPP)**—Allows multiple hosts to have their source IP addresses translated to the same public IP address with different port numbers. The dynamic translation is to the next available address in the NAT address pool, which you configure as a **Translated Address** pool to be an IP address, range of addresses, a subnet, or a combination of these.

As an alternative to the next address in the NAT address pool, DIPP allows you to specify the address of the **Interface** itself. The advantage of specifying the interface in the NAT rule is that the NAT rule will be automatically updated to use address subsequently acquired by the interface.

DIPP has a default NAT oversubscription rate, which is the number of times that the same translated IP address and port pair can be used concurrently. For more information, see [Dynamic IP and Port NAT Oversubscription](#) and [Modify the Oversubscription Rate for DIPP NAT](#).

- **Dynamic IP**—Allows the 1-to-1, dynamic translation of a source IP address only (no port number) to the next available address in the NAT address pool. By default, if the source address pool is larger than the NAT address pool and eventually all of the NAT addresses are allocated, new connections that need address translation are dropped. The default behavior can be changed by clicking **Advanced (Dynamic IP/Port Fallback)**, which causes DIPP addresses to be used when necessary. In either event, as sessions terminate and the addresses in the pool become available, they can be allocated to translate new connections.
- **Static IP**—Allows the 1-to-1, static translation of a source IP address, but leaves the source port unchanged. A common scenario for a static IP translation is an internal server that must be available to the Internet.

Destination NAT

Destination NAT is performed on incoming packets, when the firewall translates a public destination address to a private address. Destination NAT does not use address pools or ranges. It is a 1-to-1, static translation with the option to perform port forwarding or port translation.

- **Static IP**—Allows the 1-to-1, static translation of a destination IP address and optionally the port number.

One common use of destination NAT is to configure several NAT rules that map a single public destination address to several private destination addresses assigned to servers or services. For example:

- **Port Forwarding**—Can translate a public destination address and port number to a private destination address, but keeps the same port number.
- **Port Translation**—Can translate a public destination address and port number to a private destination address and a different port number, thus keeping the real port number private. It is configured by entering a **Translated Port** on the **Translated Packet** tab in the NAT policy rule.

NAT Rule Capacities

The number of NAT rules allowed is based on the firewall platform. Individual rule limits are set for static, Dynamic IP (DIP), and Dynamic IP and Port (DIPP) NAT. The sum of the number of rules used for these NAT types cannot exceed the total NAT rule capacity. For DIPP, the rule limit is based on the device's oversubscription setting (8, 4, 2, or 1) and the assumption of one translated IP address per rule. To see platform-specific NAT rule limits and translated IP address limits, use the [Compare Firewalls](#) tool.

Consider the following when working with NAT rules:

- If you run out of pool resources, you cannot create more NAT rules, even if the platform's maximum rule count has not been reached.
- If you consolidate NAT rules, the logging and reporting will also be consolidated. The statistics are provided per the rule, not per all of the addresses within the rule. If you need granular logging and reporting, do not combine the rules.

Dynamic IP and Port NAT Oversubscription

Dynamic IP and Port (DIPP) NAT allows you to use each translated IP address and port pair multiple times (8, 4, or 2 times) in concurrent sessions. This reusability of an IP address and port (known as oversubscription) provides scalability for customers who have too few public IP addresses. The design is based on the assumption that hosts are connecting to different destinations, therefore sessions can be uniquely identified and collisions are unlikely. The oversubscription rate in effect multiplies the original size of the address/port pool to 8, 4, or 2 times the size. For example, the default limit of 64K concurrent sessions allowed, when multiplied by an oversubscription rate of 8, results in 512K concurrent sessions allowed.

The oversubscription rates that are allowed vary based on the platform. The oversubscription rate is global; it applies to the device. This oversubscription rate is set by default and consumes memory, even if you have enough public IP addresses available to make oversubscription unnecessary. You can reduce the rate from the default setting to a lower setting or even 1 (which means no oversubscription). By configuring a reduced rate, you decrease the number of source device translations possible, but increase the DIP and DIPP NAT rule capacities. To change the default rate, see [Modify the Oversubscription Rate for DIPP NAT](#).

If you select **Platform Default**, your explicit configuration of oversubscription is turned off and the default oversubscription rate for the platform applies, as shown in the table below. The **Platform Default** setting allows for an upgrade or downgrade of a software release.

The following table lists the default (highest) oversubscription rate for each platform.

Platform	Default Oversubscription Rate
PA-200	2
PA-500	2
PA-2020	2
PA-2050	2
PA-3020	2
PA-3050	2
PA-3060	2
PA-4020	4
PA-4050	8
PA-4060	8
PA-5020	4
PA-5050	8
PA-5060	8
PA-7050	8
VM-100	1
VM-200	1

Platform	Default Oversubscription Rate
VM-300	2
VM-1000-HV	2

The firewall supports a maximum of 256 translated IP addresses per NAT rule, and each platform supports a maximum number of translated IP addresses (for all NAT rules combined). If oversubscription causes the maximum translated addresses per rule (256) to be exceeded, the firewall will automatically reduce the oversubscription ratio in an effort to have the commit succeed. However, if your NAT rules result in translations that exceed the maximum translated addresses for the platform, the commit will fail.

Dataplane NAT Memory Statistics

The `show running global-ippool` command displays statistics related to NAT memory consumption for a pool. The Size column displays the number of bytes of memory that the resource pool is using. The Ratio column displays the oversubscription ratio (for DIPP pools only). The lines of pool and memory statistics are explained in the following sample output:

```
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	200.0.3.11	1	1	68720	8

Usable NAT DIP/DIPP shared memory size: 58490064 ← 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-HA-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 IP/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-HA-0 (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 =>  
210.0.0.0-210.0.15.255
```

```
Source      Xlat-Source  Ref.Cnt(F)  TTL(s)
```

```
Total IPs in use: 0
```

```
Total entries in time-reserve cache: 0
```

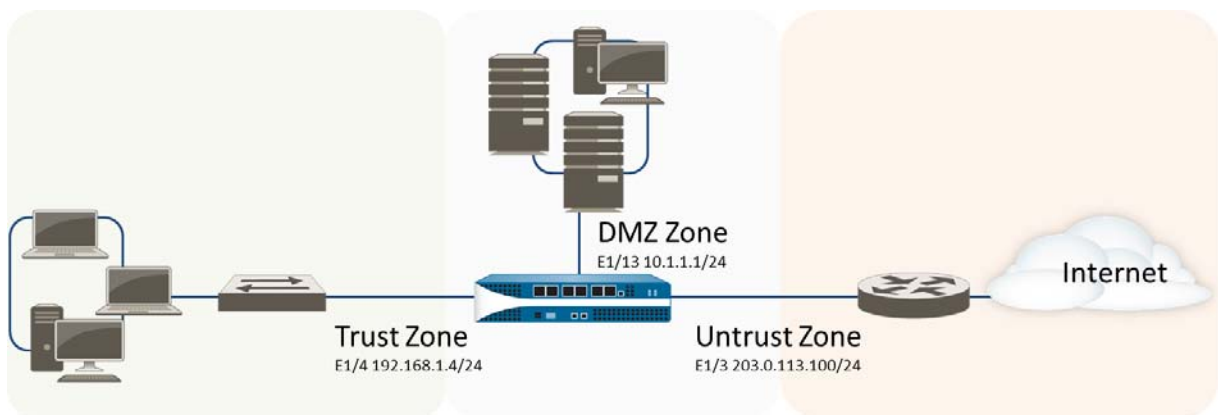
```
Total freelist left: 4096
```

Configure NAT

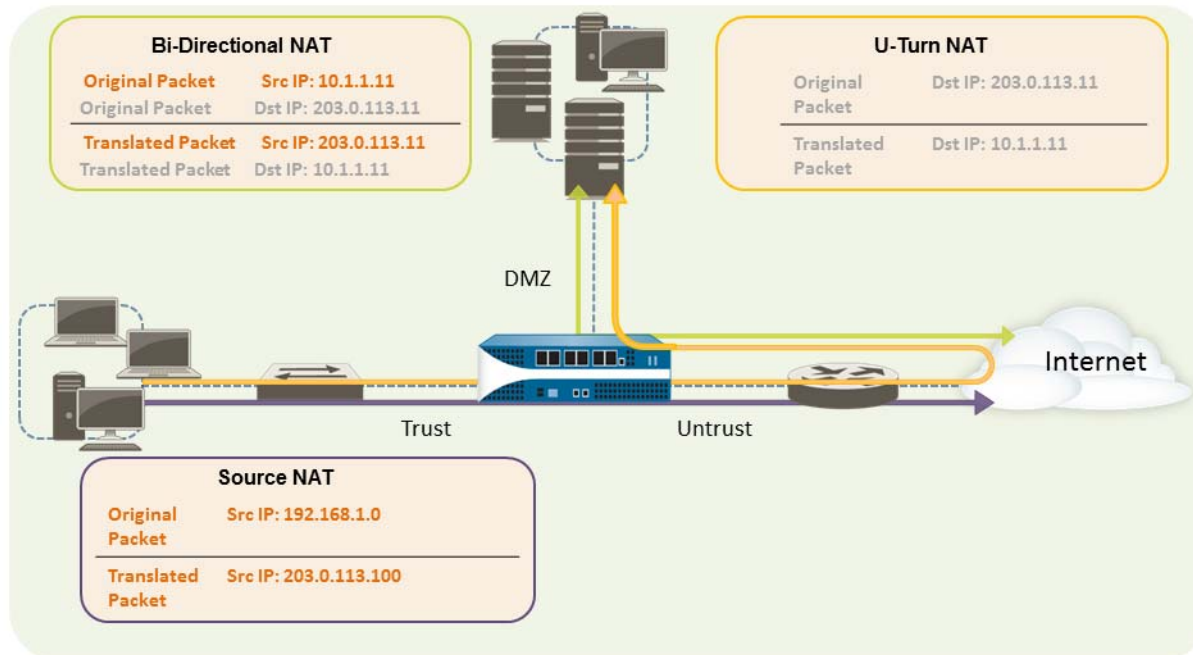
Perform the following tasks to configure various aspects of NAT.

- ▲ Translate Internal Client IP Addresses to Your Public IP Address
- ▲ Enable Clients on the Internal Network to Access your Public Servers
- ▲ Enable Bi-Directional Address Translation for Your Public-Facing Servers
- ▲ Modify the Oversubscription Rate for DIPP NAT
- ▲ Disable NAT for a Specific Host or Interface

The NAT example in this section is based on the following topology, which was also used in [Getting Started](#) for setting up interfaces and zones:



Based on the topology initially used in [Getting Started](#) to create the interfaces and zones, there are three NAT policies we need to create as follows:



- To enable the clients on the internal network to access resources on the Internet, the internal 192.168.1.0 addresses will need to be translated to publicly routable addresses. In this case, we will configure source NAT, using the egress interface address, 203.0.113.100, as the source address in all packets that leave the firewall from the internal zone. See [Translate Internal Client IP Addresses to Your Public IP Address](#) for instructions.
- To enable clients on the internal network to access the public web server in the DMZ zone, we will need to configure a NAT rule that redirects the packet from the external network, where the original routing table lookup will determine it should go based on the destination address of 203.0.113.11 within the packet, to the actual address of the web server on the DMZ network of 10.1.1.11. To do this you must create a NAT rule from the trust zone (where the source address in the packet is) to the untrust zone (where the original destination address is) to translate the destination address to an address in the DMZ zone. This type of destination NAT is called *U-Turn NAT*. See [Enable Clients on the Internal Network to Access your Public Servers](#) for instructions.
- To enable the web server—which has both a private IP address on the DMZ network and a public-facing address for access by external users—to both send and receive requests, the firewall must translate the incoming packets from the public IP address to the private IP address and the outgoing packets from the private IP address to the public IP address. On the firewall, you can accomplish this with a single bi-directional static source NAT policy. See [Enable Bi-Directional Address Translation for Your Public-Facing Servers](#).

Translate Internal Client IP Addresses to Your Public IP Address

When a client on your internal network sends a request, the source address in the packet contains the IP address for the client on your internal network. If you use private IP address ranges internally, the packets from the client will not be able to be routed on the Internet unless you translate the source IP address in the packets leaving the network into a publicly routable address.

On the firewall you can do this by configuring a source NAT policy that translates the source address and optionally the port into a public address. One way to do this is to translate the source address for all packets to the egress interface on your firewall, as shown in the following procedure.

Configure Source NAT

Step 1 Create an address object for the external IP address you plan to use.

1. Select **Objects > Addresses** and then click **Add**.
2. Enter a **Name** and optionally a **Description** for the object.
3. Select **IP Netmask** from the **Type** drop-down and then enter the IP address of the external interface on the firewall, 203.0.113.100 in this example.

The screenshot shows a dialog box titled "Address" with the following fields:

- Name:** Public Address
- Description:** External-facing IP address
- Type:** IP Netmask (selected from a dropdown menu)
- Value:** 203.0.113.100
- Tags:** (empty field)

Below the "Type" dropdown, there is a small text box with the following text: "Enter an IP address or a network using the slash notation (Ex. 192.168.80.150 or 192.168.80.0/24). You can also enter an IPv6 address or an IPv6 address with its prefix (Ex. 2001:db8:123:1::1 or 2001:db8:123:1::/64)".

At the bottom right of the dialog box are "OK" and "Cancel" buttons.

4. To save the address object, click **OK**.

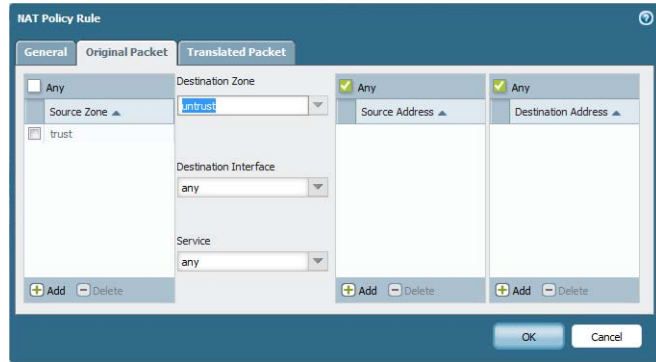


Although you do not have to use address objects in your policies, it is a best practice because it simplifies administration by allowing you to make updates in one place rather than having to update every policy where the address is referenced.

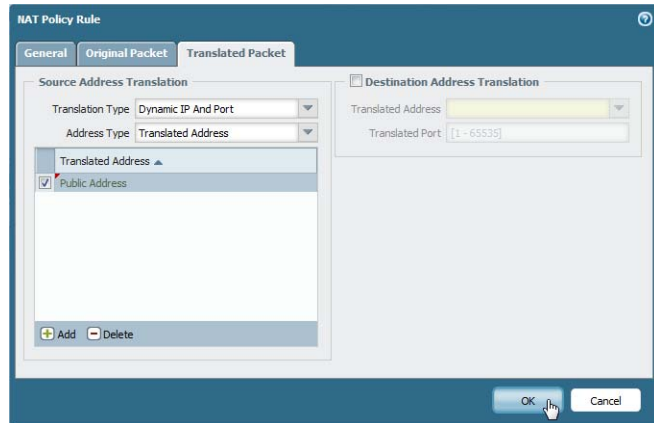
Configure Source NAT (Continued)

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. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.



4. On the **Translated Packet** tab, select **Dynamic IP And Port** from the **Translation Type** drop-down in the **Source Address Translation** section of the screen and then click **Add**.
5. Select the address object you just created.



6. Click **OK** to save the NAT policy.

Step 3 Save the configuration.

Click **Commit**.

Enable Clients on the Internal Network to Access your Public Servers

When a user on the internal network sends a request for access to the corporate web server in the DMZ, the DNS server will resolve to the public IP address. When processing the request, the firewall will use the original destination in the packet (the public IP address) and route the packet to the egress interface for the untrust zone. In order for the firewall to know that it must translate the public IP address of the web server to an address on the DMZ network when it receives requests from users on the trust zone, you must create a destination NAT rule that will enable the firewall to send the request to the egress interface for the DMZ zone as follows.

Configure U-Turn NAT

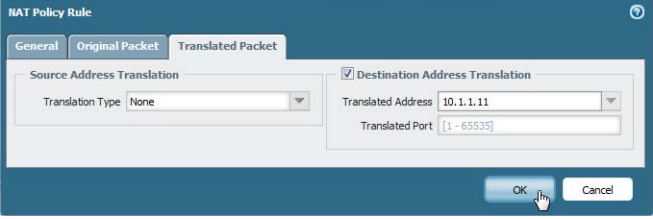
Step 1 Create an address object for the web server.

1. Select **Objects > Addresses** and click **Add**.
2. Enter a **Name** and optionally a **Description** for the object.
3. Select **IP Netmask** from the **Type** drop-down and enter the public IP address of the web server, 203.0.113.11 in this example.

4. Click **OK**.

Step 1 Create the NAT policy.

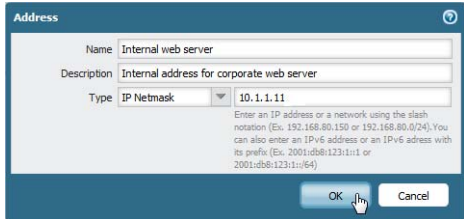

1. Select **Policies > NAT** and click **Add**.
2. On the **General** tab, enter a descriptive **Name** for the NAT rule.
3. On the **Original Packet** tab, select the zone you created for your internal network in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.
4. In the **Destination Address** section, click **Add** and select the address object you created for your public web server.

Configure U-Turn NAT (Continued)	
	<p>5. On the Translated Packet tab, select the Destination Address Translation check box and then enter the IP address that is assigned to the web server interface on the DMZ network, 10.1.1.11 in this example.</p>  <p>6. Click OK to save the NAT policy.</p>
<p>Step 2 Save the configuration.</p>	<p>Click Commit.</p>

Enable Bi-Directional Address Translation for Your Public-Facing Servers

When your public-facing servers have private IP addresses assigned on the network segment where they are physically located, you need a source NAT rule to translate the source address of the server to the external address upon egress. You create a static NAT rule to translate the internal source address, 10.1.1.11, to the external web server address, 203.0.113.11 in our example.

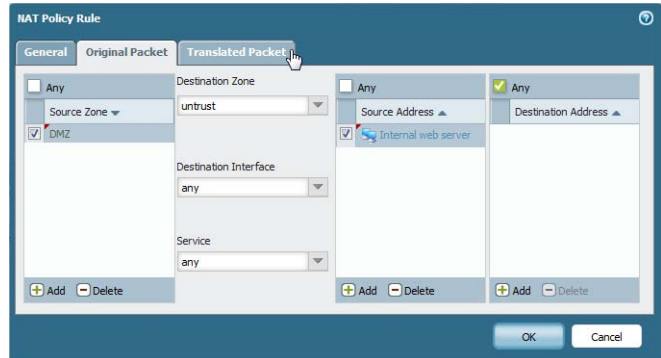
However, a public-facing server must be able to both send and receive packets. You need a reciprocal policy that translates the public address (the destination IP address in incoming packets from Internet users) into the private address so that the firewall can route the packet to your DMZ network. You create a bi-directional static NAT rule, as described in the following procedure.

Configure Bi-Directional NAT	
<p>Step 1 Create an address object for the web server's internal IP address.</p>	<ol style="list-style-type: none"> 1. Select Objects > Addresses and click Add. 2. Enter a Name and optionally a Description for the object. 3. Select IP Netmask from the Type drop-down and enter the IP address of the web server on the DMZ network, 10.1.1.11 in this example.  <ol style="list-style-type: none"> 4. Click OK. <p> If you did not already create an address object for the <i>public</i> address of your web server, you should create that object now.</p>

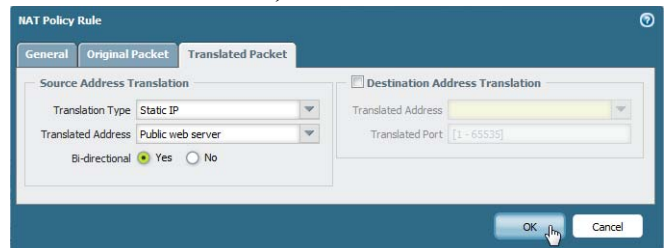
Configure Bi-Directional NAT (Continued)

Step 2 Create the NAT policy.

1. Select **Policies > NAT** and click **Add**.
2. On the **General** tab, enter a descriptive **Name** for the NAT rule.
3. On the **Original Packet** tab, select the zone you created for your DMZ in the **Source Zone** section (click **Add** and then select the zone) and the zone you created for the external network from the **Destination Zone** drop-down.
4. In the **Source Address** section, click **Add** and select the address object you created for your internal web server address.



5. On the **Translated Packet** tab, select **Static IP** from the **Translation Type** drop-down in the **Source Address Translation** section and then select the address object you created for your external web server address from the **Translated Address** drop-down.
6. In the **Bi-directional** field, select **Yes**.




7. Click **OK** to save the NAT policy.

Step 3 Save the configuration.

Click **Commit**.

Modify the Oversubscription Rate for DIPP NAT

If you have enough public IP addresses that you do not need to use DIPP NAT oversubscription, you can reduce the oversubscription rate and thereby gain more DIP and DIPP NAT rules allowed.

Set NAT Oversubscription	
<p>Step 1 View the DIPP NAT oversubscription rate.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Session > Session Settings. View the NAT Oversubscription Rate setting.
<p>Step 2 Set the DIPP NAT oversubscription rate.</p>	<ol style="list-style-type: none"> 1. Click the Edit icon in the Session Settings section. 2. In the NAT Oversubscription Rate drop-down, select 1x, 2x, 4x, or 8x, depending on which ratio you want. <div style="margin-left: 20px;">  The Platform Default setting applies the default oversubscription setting for the platform. If you want no oversubscription, select 1x. </div> 3. Click OK and Commit the change.

Disable NAT for a Specific Host or Interface

You may have exceptions where you do not want NAT to occur for a certain host in a subnet or for traffic exiting a specific interface. Source NAT and destination NAT rules can be configured to disable translation in those cases. The following procedure shows how to disable source NAT for a host.

Create a Source NAT Exemption	
<p>Step 1 Create the NAT policy.</p>	<ol style="list-style-type: none"> 1. Select Policies > NAT and click Add. 2. Enter a descriptive Name for the policy. 3. On the Original Packet tab, select the zone you created for your internal network in the Source Zone section (click Add and then select the zone) and the zone you created for the external network from the Destination Zone drop-down. 4. For Source Address, click Add and enter the host address. Click OK. 5. On the Translated Packet tab, select None from the Translation Type drop-down in the Source Address Translation section of the screen. 6. Click OK to save the NAT policy.
<p>Step 2 Save the configuration.</p>	<p>Click Commit.</p>

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](#)
- ▲ [Configure NAT64](#)

NAT64 Overview

Palo Alto Networks 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.)

Palo Alto Networks also supports IPv4-initiated communication with a static binding that maps an IPv4 address and port number to an IPv6 address. 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.

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](#).

[Configure NAT64](#) for IPv6-initiated communication or IPv4-initiated communication.

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 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.

Configure NAT64

You can configure two types of [NAT64](#) translation on the firewall; each one is doing a bi-directional translation between the two IP address families. Configure one of the following, depending on whether the initial translation is from IPv6 to IPv4 or from IPv4 to IPv6. IPv4-initiated communication offers an additional option to translate port numbers.

- ▲ [IPv6-Initiated Communication](#)
- ▲ [IPv4-Initiated Communication](#)
- ▲ [IPv4-Initiated Communication with Port Translation](#)

IPv6-Initiated Communication

As shown in the example topology below, IPv6-initiated communication requires a [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).

IPv6-initiated communication to the firewall is similar to source NAT for an IPv4 topology. 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.)

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.

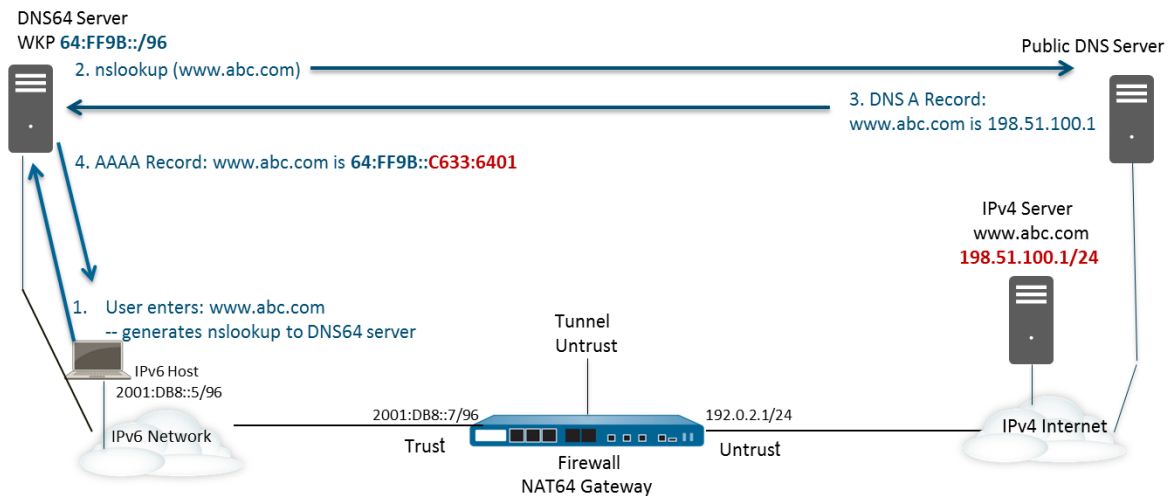
You must also configure a tunnel interface (with no termination point). You apply the NAT64 prefix to the tunnel and apply the appropriate zone to ensure that IPv6 traffic with the NAT64 prefix is assigned to the proper destination zone. The tunnel also has the advantage of dropping IPv6 traffic with 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 figure below 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.

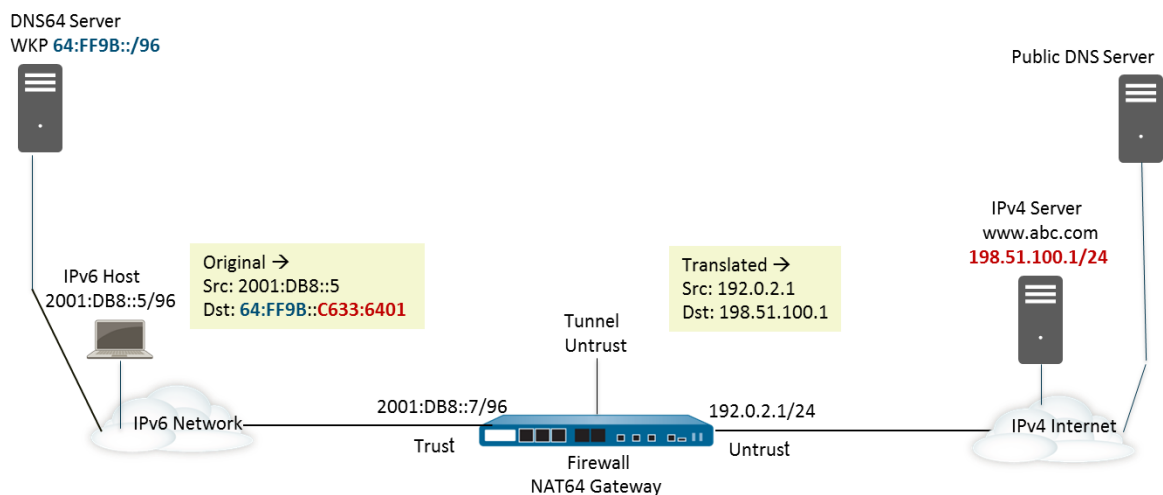
1. A user at the IPv6 host enters the URL `www.abc.com`, which generates a name server lookup (`nslookup`) to the DNS64 server.
2. The DNS64 Server sends an `nslookup` to the public DNS server for `www.abc.com`, requesting its IPv4 address.
3. The DNS server returns an A record that provides the IPv4 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

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**.

Configure NAT64 for IPv6-Initiated Communication (Continued)	
<p>Step 2 Create an address object for the IPv6 destination address (pre-translation).</p>	<ol style="list-style-type: none"> 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. NOTE: 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.
<p>Step 3 (Optional) Create an address object for the IPv6 source address (pre-translation).</p>	<ol style="list-style-type: none"> 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.
<p>Step 4 (Optional) Create an address object for the IPv4 source address (translated).</p>	<ol style="list-style-type: none"> 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.
<p>Step 5 Create the NAT64 rule.</p>	<ol style="list-style-type: none"> 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.
<p>Step 6 Specify the original source and destination information.</p>	<ol style="list-style-type: none"> 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.

Configure NAT64 for IPv6-Initiated Communication (Continued)	
Step 7 Specify the translated packet information.	<ol style="list-style-type: none"> 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: <ul style="list-style-type: none"> • 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.	<ol style="list-style-type: none"> 1. Select Network > Tunnels and Add a tunnel. 2. On the General tab, enter a Name for the tunnel. 3. For the Tunnel Interface, select New Tunnel Interface. 4. In the Interface Name field, enter a numeric suffix, such as .2. 5. On the Config tab, select the Virtual Router where you are configuring NAT64. 6. For Security Zone, select the destination zone associated with the IPv4 server destination (Trust zone). 7. On the IPv6 tab, select Enable IPv6 on the interface. 8. Click Add and for the Address, select New Address. 9. Enter a Name for the address. 10. (Optional) Enter a Description for the tunnel address. 11. For Type, select IP Netmask and enter your IPv6 prefix and prefix length, in this example, 64:FF9B::/96. 12. Click OK to save the new address. 13. Select Enable address on interface and click OK. 14. Click OK to save the tunnel interface. 15. Click OK to save the tunnel.
Step 9 Create a security policy rule to allow NAT traffic from the trust zone.	<ol style="list-style-type: none"> 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.	Click Commit .

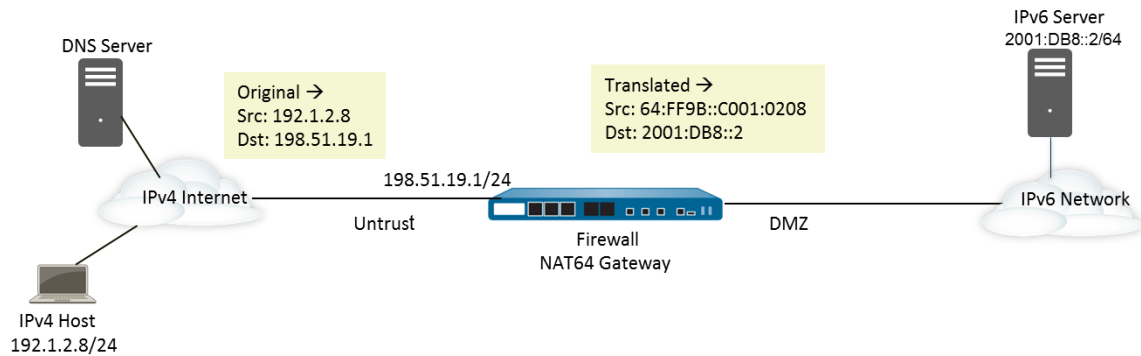
Configure NAT64 for IPv6-Initiated Communication (Continued)

Step 11 Troubleshoot or view a NAT64 session. > show session id <session-id>

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.



Configure NAT64 for IPv4-Initiated Communication

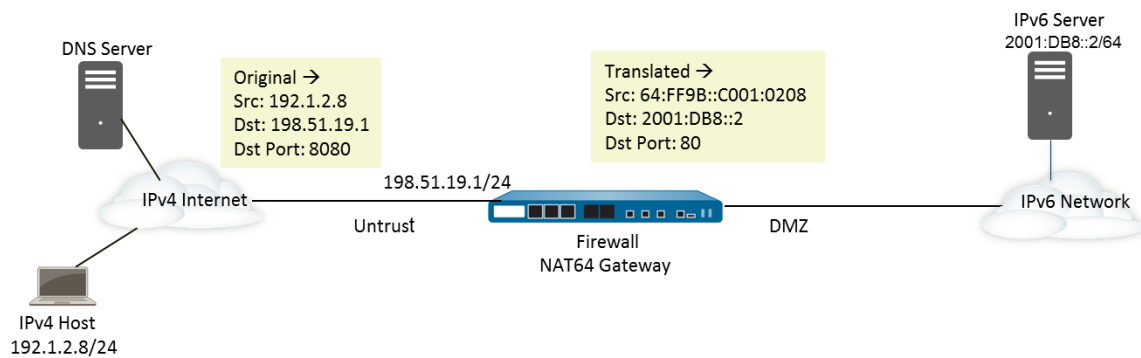
<p>Step 1 Enable IPv6 to operate on the firewall.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Session and edit the Session Settings. 2. Select Enable IPv6 Firewalling. 3. Click OK.
<p>Step 2 (Optional) When an IPv4 packet has its DF bit set to zero (and because IPv6 does not fragment packets), ensure the translated IPv6 packet does not exceed the path MTU for the destination IPv6 network.</p>	<ol style="list-style-type: none"> 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). TIP: 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.

Configure NAT64 for IPv4-Initiated Communication	
<p>Step 3 Create an address object for the IPv4 destination address (pre-translation).</p>	<ol style="list-style-type: none"> 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.
<p>Step 4 Create an address object for the IPv6 source address (translated).</p>	<ol style="list-style-type: none"> 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.
<p>Step 5 Create an address object for the IPv6 destination address (translated).</p>	<ol style="list-style-type: none"> 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. NOTE: The source and destination must have the same netmask (prefix length). 4. Click OK.
<p>Step 6 Create the NAT64 rule.</p>	<ol style="list-style-type: none"> 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.
<p>Step 7 Specify the original source and destination information.</p>	<ol style="list-style-type: none"> 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.

Configure NAT64 for IPv4-Initiated Communication	
<p>Step 8 Specify the translated packet information.</p>	<ol style="list-style-type: none"> For the Translated Packet, in the Source Address Translation, Translation Type, select Static IP. For Translated Address, select the source translated address object you created, nat64_ip6source. 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). Click OK.
<p>Step 9 Create a security policy rule to allow the NAT traffic from the Untrust zone.</p>	<ol style="list-style-type: none"> Select Policies > Security and Add a rule Name. Select Source and Add a Source Zone; select Untrust. For Source Address, select Any. Select Destination and Add a Destination Zone; select DMZ. For Actions, select Allow. Click OK.
<p>Step 10 Commit.</p>	<p>Click Commit.</p>
<p>Step 11 Troubleshoot or view a NAT64 session.</p>	<pre>> show session id <session-id></pre>

IPv4-Initiated Communication with Port Translation

This use case builds on the prior use case, 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.



Configure NAT64 for IPv4-Initiated Communication with Port Translation	
Step 1 Enable IPv6 to operate on the firewall.	<ol style="list-style-type: none"> 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.	<ol style="list-style-type: none"> 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). TIP: 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).	<ol style="list-style-type: none"> 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).	<ol style="list-style-type: none"> 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).	<ol style="list-style-type: none"> 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. NOTE: The source and destination must have the same netmask (prefix length). 4. Click OK.
Step 6 Create the NAT64 rule.	<ol style="list-style-type: none"> 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.

Configure NAT64 for IPv4-Initiated Communication with Port Translation	
<p>Step 7 Specify the original source and destination information, and create a service to limit the translation to a single ingress port number.</p>	<ol style="list-style-type: none"> 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.
<p>Step 8 Specify the translated packet information.</p>	<ol style="list-style-type: none"> 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.
<p>Step 9 Create a security policy rule to allow the NAT traffic from the Untrust zone.</p>	<ol style="list-style-type: none"> 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.
<p>Step 10 Commit.</p>	<p>Click Commit.</p>
<p>Step 11 Troubleshoot or view a NAT64 session.</p>	<pre>> show session id <session-id></pre>

LACP

The firewall can now use Link Aggregation Control Protocol (LACP) to detect the physical interfaces between itself and a connected device (peer) and manage those interfaces as a single virtual interface (aggregate group). An aggregate group increases the bandwidth between peers. Enabling LACP provides redundancy within the group: the protocol automatically detects interface failures and performs failover to standby interfaces. Without LACP, you must manually identify interface failures occurring at layers above the physical or occurring between peers that are not directly connected.

The following Palo Alto Networks firewalls support LACP: PA-500, PA-3000 Series, PA-4000 Series, PA-5000 Series, and PA-7050. The firewalls support LACP for HA3 (only on the PA-500, PA-3000 Series, PA-4000 Series, and PA-5000 Series), Layer 2, and Layer 3 interfaces.

The following topics describe LACP and how to configure it for aggregate groups:

- ▲ [LACP Settings](#)
- ▲ [Configure LACP](#)

LACP Settings

To implement LACP, you must configure an aggregate group on each LACP peer that will use the protocol. Before you [Configure LACP](#), determine the optimal settings for each peer, as described in the following topics:

- ▲ [Mode](#)
- ▲ [Transmission Rate](#)
- ▲ [Fast Failover](#)
- ▲ [Port Priority and System Priority](#)

Mode

In LACP active mode, a device actively searches the network for peers and queries their status (available or unresponsive). In passive mode, a device only responds to queries from an active peer. Between any two peers, it is recommended that one be active and the other passive. LACP cannot function if both peers are passive.

Transmission Rate

You can configure a device to detect the state (available or unresponsive) of peers and individual interfaces at fast (every second) or slow (every 30 seconds) intervals. When the device does not receive an LACP update from the peer within a period that is three times the transmission rate (3 seconds or 90 seconds), it flags the peer as unresponsive. Therefore, set the transmission rate according to how much LACP processing your network can support and how quickly a device should detect and resolve interface failures.

For example, in a network where peers have many redundant links, a single interface failure would not disrupt traffic. In this case, a slow rate would suffice for detecting failures and would avoid the extra processing associated with the fast rate. If the network can support the extra processing and requires high availability (for example, datacenters supporting critical business operations), a fast transmission rate would enable the firewall to substitute failed interfaces quicker.



If devices have different transmission rates, each uses the rate of its peer.

Fast Failover

When an interface goes down, it fails over to a standby interface. The IEEE 802.1ax standard that defines LACP specifies a failover process that takes at least three seconds. For networks that require quicker failover, the firewall provides an option that performs fast failover within one second. This option is recommended for deployments in which critical data might be lost during the standard failover interval. For example, if the amount of traffic between peers is close to the maximum that the active interfaces can support and one interface fails, the risk of data being lost while a standby interface becomes active is significantly less for fast failover than for standard failover.



Because both peers perform failover when an interface goes down, it is recommended that you configure fast failover on both so that they complete the process within the same time frame.

Port Priority and System Priority

When you configure an aggregate group, the **Max Ports** parameter determines how many interfaces can be active at any given time. If the number of interfaces you assign to the group exceeds the **Max Ports**, the remaining interfaces are in standby mode. If an active interface fails, a standby interface becomes active. LACP uses the **LACP Port Priority** assigned to each interface to determine which interfaces are initially active and to determine the order in which standby interfaces become active upon failover. (You set port priorities when configuring individual interfaces, not when configuring the aggregate group.)

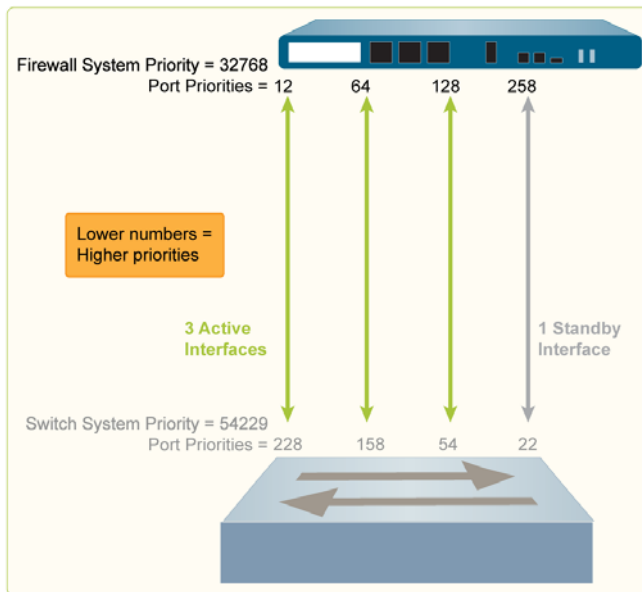


It is recommended that you assign a unique priority number to each interface in an aggregate group. However, if multiple interfaces have the same priority number, the interface with the lowest port number has the highest priority.

To enable LACP between two peers, each requires an aggregate group configuration. If the port priority values for the member interfaces differ in each peer, the values of the peer with the higher **System Priority** will override the other peer in determining active and standby interfaces.


For port and system priorities, the priority level is the inverse of its numerical value: a priority of 1 designates the highest priority level; a value of 65535 designates the lowest priority level. The default is 32768.

The following figure shows an aggregate group instance configured on a firewall and on a switch. The group has four interfaces and the **Max Ports** parameter is set to three. That means three interfaces are active and one is standby. The firewall administrator assigned port priorities 12, 64, 128, and 258 to the interfaces. The switch administrator assigned port priorities 22, 54, 158, and 228 to the interfaces. LACP uses the port priorities of the firewall instance because the **System Priority** of the firewall is higher than that of the switch.

Figure: LACP System Priority

It is a best practice to assign a unique system priority to each peer. However, in some cases, peers might have the same value. This could happen if a different administrator configured the aggregate group on each peer. In such cases, the peer for which the *system ID* is a lower numerical value will override the other peer with respect to port priorities. LACP automatically derives the system ID from the system priority and system MAC address. A numerically lower MAC address produces a numerically lower system ID.

The following figure shows the same peers as in [Figure: LACP System Priority](#) but with identical system priorities. In this case, LACP uses the system IDs to determine port prioritization: the switch overrides the firewall.

Figure: LACP System ID

Firewalls in a high availability (HA) pair have the same system priority value. However, in an active/passive deployment, the system ID for each can be the same or different, depending on whether you assign the same MAC address. (Firewalls in an active/active deployment require unique MAC addresses so PAN-OS automatically assigns them.) When the LACP peers (also in HA mode) are virtualized (appearing to the network as a single device), selecting the **Same System MAC Address for Active-Passive HA** option for the firewalls is a best practice to minimize latency during failover. When the LACP peers are not virtualized, using the unique MAC address of each firewall is the best practice to minimize failover latency. In the latter case, if the HA firewall pair and the LACP peer devices all have the same system priority but each firewall in the HA pair has a unique system ID, one firewall might have a lower system ID than the LACP peers while the other firewall has a higher system ID than the LACP peers. In this case, when failover occurs on the firewalls, port prioritization switches between the LACP peers and the firewall that becomes active.

Configure LACP

Before starting this procedure:

- Determine which physical interfaces connect the LACP peers. This procedure assumes the cabling is complete.
- Determine the optimal [LACP Settings](#) for the peers.






This procedure only covers the configuration steps for an LACP peer that is a Palo Alto Networks firewall. Other devices have proprietary methods for configuring LACP.

Perform the following steps to configure LACP on a firewall. In a high availability deployment, configure the primary (for active/active) or active (for active/passive) firewall; the secondary or passive firewall automatically synchronizes with it.

Configure LACP

Step 1 Add an aggregate group with LACP enabled.

1. Select **Network > Interfaces > Ethernet** and click **Add Aggregate Group**.
2. In the field adjacent to the read-only **Interface Name**, enter a number (1-8) to identify the group.
3. For the **Interface Type**, select **HA, Layer 2**, or **Layer 3**.
 Only select **HA** if the interface is an HA3 link between two firewalls in an active/active deployment.
4. On the **LACP** tab, select **Enable LACP**.
5. For the **LACP Mode**, select whether the firewall is **Passive** (default) or **Active**.
6. For the **Transmission Rate**, select **Fast** or **Slow**.
7. If desired, select **Fast Failover** (under one second) for when interfaces go down. Otherwise, the standard failover (at least three seconds) applies.
8. Enter a **System Priority** number (1-65535, default 32768) that determines whether the firewall or peer overrides the other with respect to port priorities. Note that the lower the number, the higher the priority.
9. For the **Max Ports**, enter the number of interfaces (1-8) that are active. The value cannot exceed the number of interfaces you assign to the group. If the number of assigned interfaces exceeds the number of active interfaces, the remaining interfaces will be in standby mode.
10. By default, each firewall in an HA pair has a unique MAC address. A **Same System MAC Address for Active-Passive HA** configuration is recommended only when the LACP peers are virtualized (appearing to the network as a single device). In such cases, select the check box and select the system-generated **MAC Address**, or enter your own. You must verify the address is globally unique.
 If the firewalls are not in active/passive HA mode, PAN-OS ignores your selections for these fields. Firewalls in an active/active deployment require unique MAC addresses so PAN-OS automatically assigns them. If the **Interface Type** is **HA**, these fields do not appear.
11. Click **OK**.

Configure LACP (Continued)	
<p>Step 2 Assign interfaces to the aggregate group.</p>	<p>Perform the following steps for each physical interface (1-8) that will belong to the aggregate group.</p> <p> During normal operations, assign more than one interface. You would only assign a single interface if a troubleshooting procedure requires it.</p> <ol style="list-style-type: none"> 1. Select Network > Interfaces and select the interface name. It must be the same type (HA, Layer 2, or Layer 3) as the aggregate group. 2. Change the Interface Type to Aggregate Ethernet. 3. Select the Aggregate Group you just defined. 4. Select the Link Speed, Link Duplex, and Link State. It is a best practice to set the same link speed and duplex values for every interface in the group. For non-matching values, the commit operation displays a warning and PAN-OS defaults to the higher speed and full duplex. 5. Enter an LACP Port Priority (1-65535, default 32768). If the number of interfaces you assign exceeds the Max Ports value you configured for the group, the port priorities determine which interfaces are active or standby. The lower the numeric value, the higher the priority. 6. Click OK and Commit.
<p>Step 3 Verify the status of the aggregate group.</p>	<ol style="list-style-type: none"> 1. Select Network > Interfaces > Ethernet. 2. Verify that the Features column displays an LACP enabled icon for the aggregate group. 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 the Link State icon for the aggregate group is yellow or red, select Monitor > Logs > System and review the logs for the LACP subtype to investigate the causes of the interface failures.
<p>Step 4 (Optional) If you selected HA as the Interface Type for the aggregate group, enable forwarding of packets over the HA3 link.</p>	<ol style="list-style-type: none"> 1. Select Device > High Availability > Active/Active Config and edit the Packet Forwarding section. 2. For the HA3 Interface, select the aggregate group you configured. 3. Click OK and Commit.



Policy

Policies allow you to enforce rules and take action. The different types of policy rules that you can create on the firewall are: Security, NAT, Quality of Service (QoS), Policy Based Forwarding (PBF), Decryption, Application Override, Captive Portal, Denial of Service (DoS), and Zone protection policies. All these different policies work together to allow, deny, prioritize, forward, encrypt, decrypt, make exceptions, authenticate access, and reset connections as needed to help secure your network. The following topics describe how to work with policy:

- ▲ [Policy Types](#)
- ▲ [Security Policy](#)
- ▲ [Policy Objects](#)
- ▲ [Security Profiles](#)
- ▲ [Enumeration of Rules Within a Rulebase](#)
- ▲ [Use Tags to Visually Distinguish Objects](#)
- ▲ [Tips and Tricks for Searching for Objects in Security Policy](#)
- ▲ [Use a Dynamic Block List in Policy](#)
- ▲ [Register IP Addresses and Tags Dynamically](#)
- ▲ [Monitor Changes in the Virtual Environment](#)
- ▲ [CLI Commands for Dynamic IP Addresses and Tags](#)
- ▲ [Policy-Based Forwarding](#)
- ▲ [DoS Protection Against Flooding of New Sessions](#)

Policy Types

The Palo Alto Networks next-generation firewall supports a variety of policy types that work together to safely enable applications on your network.

Policy Type	Description
Security	Determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service. For more details, see Security Policy .
NAT	Instruct the firewall which packets need translation and how to do the translation. The firewall supports both source address and/or port translation and destination address and/or port translation. For more details, see NAT .
QoS	Identify traffic requiring QoS treatment (either preferential treatment or bandwidth-limiting) using a defined parameter or multiple parameters and assigns it a class. For more details, see Quality of Service .
Policy Based Forwarding	Identify traffic that should use a different egress interface than the one that would normally be used based on the routing table. For details, see Policy-Based Forwarding .
Decryption	Identify encrypted traffic that you want to inspect for visibility, control, and granular security. For more details, see Decryption .
Application Override	Identify sessions that you do not want processed by the App-ID engine, which is a Layer-7 inspection. Traffic matching an application override policy forces the firewall to handle the session as a regular stateful inspection firewall at Layer-4. For more details, see Manage Custom or Unknown Applications .
Captive Portal	Identify traffic that requires the user to be known. The captive portal policy is only triggered if other User-ID mechanisms did not identify a user to associate with the source IP address. For more details, see Captive Portal .
DoS Protection	Identify potential denial-of-service (DoS) attacks and take protective action in response to rule matches. DoS Protection Profile .

Security Policy

Security policies protect network assets from threats and disruptions and aid in optimally allocating network resources for enhancing productivity and efficiency in business processes. On the Palo Alto Networks firewall, security policies determine whether to block or allow a session based on traffic attributes such as the source and destination security zone, the source and destination IP address, the application, user, and the service.

All traffic passing through the firewall is matched against a session and each session is matched against a security policy. When a session match occurs, the security policy is applied to bi-directional traffic (client to server and server to client) in that session. For traffic that doesn't match any defined rules, the default rules apply. The default rules—displayed at the bottom of the security rulebase—are predefined to allow all intrazone (within the zone) traffic and deny all interzone (between zones) traffic. Although these rules are part of the pre-defined configuration and are read-only by default, you can override them and change a limited number of settings, including the tags, action (allow or deny), log settings, and security profiles.

Security policies are evaluated left to right and from top to bottom. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated. Therefore, the more specific rules must precede more generic ones in order to enforce the best match criteria. Traffic that matches a rule generates a log entry at the end of the session in the traffic log, if logging is enabled for that rule. The logging options are configurable for each rule, and can for example be configured to log at the start of a session instead of, or in addition to, logging at the end of a session.

- ▲ [Components of a Security Policy Rule](#)
- ▲ [Security Policy Best Practices](#)

Components of a Security Policy Rule

The security policy rule construct permits a combination of the required and optional fields as detailed in the following tables:



▲ Required Fields

▲ Optional Fields

Required Fields

Required Field	Description
Name	A label that supports up to 31 characters, used to identify the rule.
Rule Type	Specifies whether the rule applies to traffic within a zone, between zones, or both: <ul style="list-style-type: none"> • universal (default)—Applies the rule to all matching interzone and intrazone traffic in the specified source and destination zones. For example, if you create a universal rule with source zones A and B and destination zones A and B, the rule would apply to all traffic within zone A, all traffic within zone B, and all traffic from zone A to zone B and all traffic from zone B to zone A. • intrazone—Applies the rule to all matching traffic within the specified source zones (you cannot specify a destination zone for intrazone rules). For example, if you set the source zone to A and B, the rule would apply to all traffic within zone A and all traffic within zone B, but not to traffic between zones A and B. • interzone—Applies the rule to all matching traffic between the specified source and destination zones. For example, if you set the source zone to A, B, and C and the destination zone to A and B, the rule would apply to traffic from zone A to zone B, from zone B to zone A, from zone C to zone A, and from zone C to zone B, but not traffic within zones A, B, or C.
Source Zone	The zone from which the traffic originates.
Destination Zone	The zone at which the traffic terminates. If you use NAT, make sure to always reference the post-NAT zone.
Application	The application which you wish to control. The firewall uses App-ID, the traffic classification technology, to identify traffic on your network. App-ID provides application control and visibility in creating security policies that block unknown applications, while enabling, inspecting, and shaping those that are allowed.
Action	Specifies an <i>Allow</i> or <i>Deny</i> action for the traffic based on the criteria you define in the rule.

Optional Fields

Optional Field	Description
Tag	A keyword or phrase that allows you to filter security rules. This is handy when you have defined many rules and wish to then review those that are tagged with a particular keyword, for example <i>Inbound to DMZ</i> .
Description	A text field, up to 255 characters, used to describe the rule.
Source IP Address	Define host IP or FQDN, subnet, named groups, or country-based enforcement. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address).
Destination IP Address	The location or destination for the traffic. If you use NAT, make sure to always refer to the original IP addresses in the packet (i.e. the pre-NAT IP address).
User	The user or group of users for whom the policy applies. You must have User-ID enabled on the zone. To enable User-ID, see User-ID Overview .
URL Category	<p>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.</p> <p>Although you can manually configure URL categories on your device, to take advantage of the dynamic URL categorization updates available on the Palo Alto Networks firewalls, you must purchase a URL filtering license.</p> <p> 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 <i>Any</i> and attach a URL Filtering profile to the security policy. See Set Up Basic Security Policies for information on using the default profiles in your security policy and see Control Access to Web Content for more details.</p>
Service	<p>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.</p> <p> For inbound allow rules (for example, from untrust to trust), using <i>application-default</i> prevents applications from running on unusual ports and protocols. Application-default is the default option; while the device still checks for all applications on all ports, with this configuration, applications are only allowed on their standard ports/protocols.</p>
Security Profiles	Provide additional protection from threats, vulnerabilities, and data leaks. Security profiles are only evaluated for rules that have an <i>allow</i> action.

Optional Field	Description (Continued)
HIP Profile (for GlobalProtect)	Allows you to identify clients with Host Information Profile (HIP) and then enforce access privileges.
Options	Allow you to define logging for the session, log forwarding settings, change Quality of Service (QoS) markings for packets that match the rule, and schedule when (day and time) the security rule should be in effect.

Security Policy Best Practices


The task of safely enabling Internet access and preventing misuse of web access privileges, and exposure to vulnerabilities and attacks is a continuous process. The key principle when defining policy on the Palo Alto Networks firewall is to use a positive enforcement approach. Positive enforcement implies that you selectively allow what is required for day-to-day business operations as opposed to a negative enforcement approach where you would selectively block everything that is not allowed. Consider the following suggestions when creating policy:

- ❑ If you have two or more zones with identical security requirements, combine them into one security rule.
- ❑ The ordering of rules is crucial to ensure the best match criteria. Because policy is evaluated top down, the more specific policy must precede the ones that are more general, so that the more specific rule is not *shadowed*. The term shadow refers to a rule that is not evaluated or is skipped because it is placed lower in the policy list. When the rule is placed lower, it is not evaluated because the match criteria was met by another rule that preceded it, thereby shadowing the rule from policy evaluation.
- ❑ To restrict and control access to inbound applications, in the security policy, explicitly define the port that the service/application will be listening on.
- ❑ Logging for broad allow rules—for example access to well known servers like DNS—can generate a lot of traffic. Hence it is not recommended unless absolutely necessary.
- ❑ By default, the firewall creates a log entry at the end of a session. However, you can modify this default behavior and configure the firewall to log at the start of the session. Because this significantly increases the log volume, logging at session start is recommended only when you are troubleshooting an issue. Another alternative for troubleshooting without enabling logging at session start is to use the session browser (**Monitor > Session Browser**) to view the sessions in real time.

Policy Objects

A *policy object* is a single object or a collective unit that groups discrete identities such as IP addresses, URLs, applications, or users. With policy objects that are a collective unit, you can reference the object in security policy instead of manually selecting multiple objects one at a time. Typically, when creating a policy object, you group objects that require similar permissions in policy. For example, if your organization uses a set of server IP addresses for authenticating users, you can group the set of server IP addresses as an *address group* policy object and reference the address group in the security policy. By grouping objects, you can significantly reduce the administrative overhead in creating policies.

You can create the following policy objects on the firewall:

Policy Object	Description
Address/Address Group, Region	<p>Allow you to group specific source or destination addresses that require the same policy enforcement. The address object can include an IPv4 or IPv6 address (single IP, range, subnet) or the FQDN. Alternatively, a region can be defined by the latitude and longitude coordinates or you can select a country and define an IP address or IP range. You can then group a collection of address objects to create an <i>address group</i> object.</p> <p>You can also use dynamic address groups to dynamically update IP addresses in environments where host IP addresses change frequently.</p>
User/User Group	<p>Allow you to create a list of users from the local database or an external database and group them.</p>
Application Group and Application Filter	<p>An <i>Application Filter</i> allows you to filter applications dynamically. It allows you to filter, and save a group of applications using the attributes defined in the application database on the firewall. For example, you can filter by one or more attributes—category, sub-category, technology, risk, characteristics—and save your application filter. With an application filter, when a PAN-OS content update occurs, any new applications that match your filter criteria are automatically added to your saved application filter.</p> <p>An <i>Application Group</i> allows you to create a static group of specific applications that you wish to group together for a group of users or for a particular service.</p>
Service/Service Groups	<p>Allows you to specify the source and destination ports and protocol that a service can use. The firewall includes two pre-defined services—service-http and service-https—that use TCP ports 80 and 8080 for HTTP, and TCP port 443 for HTTPS. You can however, create any custom service on any TCP/UDP port of your choice to restrict application usage to specific ports on your network (in other words, you can define the default port for the application).</p> <p> To view the standard ports used by an application, in Objects > Applications search for the application and click the link. A succinct description displays.</p>

Security Profiles

While security policies enable you to allow or deny traffic on your network, security profiles help you define an *allow but scan* rule, which scan allowed applications for threats, such as viruses, malware, spyware, and DDOS attacks. When traffic matches the allow rule defined in the security policy, the security profile(s) that are attached to the rule are applied for further content inspection rules such as antivirus checks and data filtering.



Security profiles are not used in the match criteria of a traffic flow. The security profile is applied to scan traffic after the application or category is allowed by the security policy.

The firewall provides default security profiles that you can use out of the box to begin protecting your network from threats. See [Set Up Basic Security Policies](#) for information on using the default profiles in your security policy. As you get a better understanding about the security needs on your network, you can create custom profiles. See [Scan Traffic for Threats](#) for more information.

You can add security profiles that are commonly applied together to a security profile group; this set of profiles can be treated as a unit and added to security policies in one step (or included in security policies by default, if you choose to set up a default security profile group).

The following topics provide more detailed information about each type of security profile and how to set up a security profile group:

- ▲ [Antivirus Profiles](#)
- ▲ [Anti-spyware Profiles](#)
- ▲ [Vulnerability Protection Profiles](#)
- ▲ [URL Filtering Profiles](#)
- ▲ [Data Filtering Profiles](#)
- ▲ [File Blocking Profiles](#)
- ▲ [DoS Protection Profile](#)
- ▲ [Zone Protection Profiles](#)
- ▲ [Security Profile Group](#)

Antivirus Profiles

Antivirus profiles protect against viruses, worms, and trojans as well as spyware downloads. Using a stream-based malware prevention engine, which inspects traffic the moment the first packet is received, the Palo Alto Networks antivirus solution can provide protection for clients without significantly impacting the performance of the firewall. This profile scans for a wide variety of malware in executables, PDF files, HTML and JavaScript viruses, including support for scanning inside compressed files and data encoding schemes. If you have enabled [Decryption](#) on the firewall, the profile also enables scanning of decrypted content.

The default profile inspects all of the listed protocol decoders for viruses, and generates alerts for SMTP, IMAP, and POP3 protocols while blocking for FTP, HTTP, and SMB protocols. Customized profiles can be used to minimize antivirus inspection for traffic between trusted security zones, and to maximize the inspection of traffic received from untrusted zones, such as the Internet, as well as the traffic sent to highly sensitive destinations, such as server farms.

The Palo Alto Networks WildFire system also provides signatures for persistent threats that are more evasive and have not yet been discovered by other antivirus solutions. As threats are discovered by WildFire, signatures are quickly created and then integrated into the standard antivirus signatures that can be downloaded by Threat Prevention subscribers on a daily basis (sub-hourly for WildFire subscribers).

Anti-spyware Profiles

Anti-spyware profiles blocks spyware on compromised hosts from trying to phone-home or beacon out to external command-and-control (C2) servers, allowing you to detect malicious traffic leaving the network from infected clients. You can apply various levels of protection between zones. For example, you may want to have custom anti-spyware profiles that minimize inspection between trusted zones, while maximizing inspection on traffic received from an untrusted zone, such as Internet facing zones.

You can define your own custom Anti-spyware profiles, or choose one of the following predefined profiles when applying anti-spyware to a security policy:

- **Default**—Uses the default action for every signature, as specified by Palo Alto Networks when the signature is created.
- **Strict**—Overrides the default action of critical, high, and medium severity threats to the block action, regardless of the action defined in the signature file. This profile still uses the default action for medium and informational severity signatures.

In addition, you can enable the [DNS Sinkholing](#) action in anti-spyware profiles to enable the firewall to forge a response to a DNS query for a known malicious domain, causing the malicious domain name to resolve to an IP address that you define. This feature helps to identify infected hosts on the protected network using DNS traffic. Infected hosts can then be easily identified in the traffic and threat logs because any host that attempts to connect to the sinkhole IP address are most likely infected with malware.

Anti-spyware and vulnerability protection profiles are configured similarly.

Vulnerability Protection Profiles

Vulnerability protection profiles stop attempts to exploit system flaws or gain unauthorized access to systems. While anti-spyware profiles help identify infected hosts as traffic leaves the network, vulnerability protection profiles protect against threats entering the network. For example, vulnerability protection profiles help protect against buffer overflows, illegal code execution, and other attempts to exploit system vulnerabilities. The default vulnerability protection profile protects clients and servers from all known critical, high, and medium-severity threats. You can also create exceptions, which allow you to change the response to a specific signature.

URL Filtering Profiles

[URL Filtering](#) profiles enable you to monitor and control how users access the web over HTTP and HTTPS. The firewall comes with a default profile that is configured to block websites such as known malware sites, phishing sites, and adult content sites. You can use the default profile in a security policy, clone it to be used as a starting point for new URL filtering profiles, or add a new URL profile that will have all categories set to allow for visibility into the traffic on your network. You can then customize the newly added URL profiles and add lists of specific websites that should always be blocked or allowed, which provides more granular control over URL categories.

Data Filtering Profiles

Data filtering profiles prevent sensitive information such as credit card or social security numbers from leaving a protected network. The data filtering profile also allows you to filter on key words, such as a sensitive project name or the word confidential. It is important to focus your profile on the desired file types to reduce false positives. For example, you may only want to search Word documents or Excel spreadsheets. You may also only want to scan web-browsing traffic, or FTP.

You can use default profiles, or create custom data patterns. There are two default profiles:

- **CC# (Credit Card)**—Identifies credit card numbers using a hash algorithm. The content must match the hash algorithm in order for data to be detected as a credit card number. This method will reduce false positives.
- **SSN# (Social Security Number)**—Uses an algorithm to detect nine digit numbers, regardless of format. There are two fields: **SSN#** and **SSN# (no dash)**.

Weight and Threshold Values

It is important to understand how the weight of an object (SSN, CC#, pattern) is calculated in order to set the appropriate threshold for a condition you are trying to filter. Each occurrence multiplied by the weight value will be added together in order to reach an action threshold (alert or block).

Example: Filter for Social Security Numbers Only

For simplicity, if you only want to filter files with Social Security Numbers (SSN) and you define a weight of 3 for **SSN#**, you would use the following formula: each instance of a SSN x weight = threshold increment. In this case, if a Word document has 10 social security numbers you multiply that by the weight of 3, so $10 \times 3 = 30$. In order to take action for a file that contains 10 social security numbers you would set the threshold to 30. You may want to set an alert at 30 and then block at 60. You may also want to set a weight in the field **SSN# (no dash)** for Social Security Numbers that do not contain dashes. If multiple settings are used, they will accumulate to reach a given threshold.

Example: Filter for Social Security Numbers and a Custom Pattern

In this example, we will filter on files that contain Social Security Numbers and the custom pattern confidential. In other words, if a file has Social Security Numbers in addition to the word confidential and the combined instances of those items hit the threshold, the file will trigger an alert or block, depending on the action setting.

SSN# weight = 3

Custom Pattern confidential weight = 20

The custom pattern is case sensitive.

If the file contains 20 Social Security Numbers and a weight of 3 is configured, that is $20 \times 3 = 60$. If the file also contains one instance of the term confidential and a weight of 20 is configured, that is $1 \times 20 = 20$ for a total of 80. If your threshold for block is set to 80, this scenario would block the file. The alert or block action will be triggered as soon as the threshold is hit.

File Blocking Profiles

The firewall uses file blocking profiles two ways: to forward files to WildFire for analysis or to block specified file types over specified applications and in the specified session flow direction (inbound/outbound/both). You can set the profile to alert or block on upload and/or download and you can specify which applications will be subject to the file blocking profile. You can also configure custom block pages that will appear when a user attempts to download the specified file type. This allows the user to take a moment to consider whether or not they want to download a file.

You can configure a file blocking profile with the following actions:

- **Alert**—When the specified file type is detected, a log is generated in the data filtering log.
- **Block**—When the specified file type is detected, the file is blocked and a customizable block page is presented to the user. A log is also generated in the data filtering log.
- **Continue**—When the specified file type is detected, a customizable response page is presented to the user. The user can click through the page to download the file. A log is also generated in the data filtering log. Because this type of forwarding action requires user interaction, it is only applicable for web traffic.
- **Forward**—When the specified file type is detected, the file is sent to WildFire for analysis. A log is also generated in the data filtering log.
- **Continue-and-forward**—When the specified file type is detected, a customizable continuation page is presented to the user. The user can click through the page to download the file. If the user clicks through the continue page to download the file, the file is sent to WildFire for analysis. A log is also generated in the data filtering log.

DoS Protection Profile

DoS protection profiles provide detailed control for Denial of Service (DoS) protection policies. DoS policies allow you to control the number of sessions between interfaces, zones, addresses, and countries based on aggregate sessions or source and/or destination IP addresses. There are two DoS protection mechanisms that the Palo Alto Networks firewalls support.

- **Flood Protection**—Detects and prevents attacks where the network is flooded with packets resulting in too many half-open sessions and/or services being unable to respond to each request. In this case the source address of the attack is usually spoofed. See [DoS Protection Against Flooding of New Sessions](#).
- **Resource Protection**— Detects and prevent session exhaustion attacks. In this type of attack, a large number of hosts (bots) are used to establish as many fully established sessions as possible to consume all of a system's resources.

You can enable both types of protection mechanisms in a single DoS protection profile.

The DoS profile is used to specify the type of action to take and details on matching criteria for the DoS policy. The DoS profile defines settings for SYN, UDP, and ICMP floods, can enable resource protect and defines the maximum number of concurrent connections. After you configure the DoS protection profile, you then attach it to a DoS policy.

When configuring DoS protection, it is important to analyze your environment in order to set the correct thresholds and due to some of the complexities of defining DoS protection policies, this guide will not go into detailed examples. For more information, refer to the [Threat Prevention Tech Note](#).

Zone Protection Profiles

Zone protection profiles provide additional protection between specific network zones in order to protect the zones against attack. The profile must be applied to the entire zone, so it is important to carefully test the profiles in order to prevent issues that may arise with the normal traffic traversing the zones. When defining packets per second (pps) thresholds limits for zone protection profiles, the threshold is based on the packets per second that do not match a previously established session. For more information, refer to the [Threat Prevention Tech Note](#).

Security Profile Group

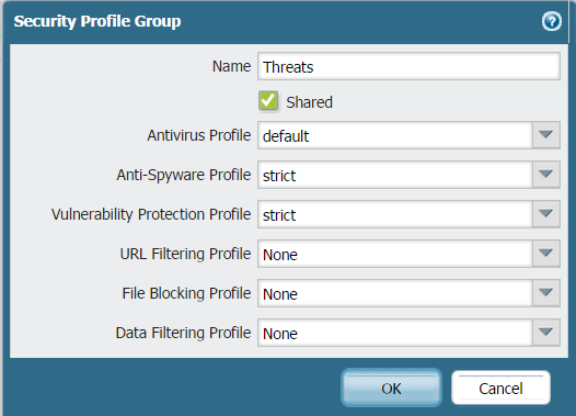
A security profile group is a set of security profiles that can be treated as a unit and then easily added to security policies. Profiles that are often assigned together can be added to profile groups to simplify the creation of security policies. You can also setup a default security profile group—new security policies will use the settings defined in the default profile group to check and control traffic that matches the security policy. Name a security profile group *default* to allow the profiles in that group to be added to new security policies by default. This allows you to consistently include your organization's preferred profile settings in new policies automatically, without administrators having to manually add them each time they create a new policy.

The following sections show how to create a security profile group and how to enable a profile group to be used by default in new security policies:

- ▲ [Create a Security Profile Group](#)
- ▲ [Set Up or Override a Default Security Profile Group](#)

Create a Security Profile Group

Use the following steps to create a security profile group and add it to a security policy.

Create a Security Profile Group	
<p>Step 1 Create a security profile group.</p>	<ol style="list-style-type: none"> 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. <div style="text-align: center;">  </div> <ol style="list-style-type: none"> 5. Click OK to save the profile group.

Create a Security Profile Group	
<p>Step 2 Add a security profile group to a security policy.</p>	<ol style="list-style-type: none"> 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 Threats group): <div data-bbox="743 499 1221 646" style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p style="text-align: center; margin: 0;">Profile Setting</p> <p>Profile Type <input type="text" value="Group"/></p> <p>Group Profile <input type="text" value="Threats"/></p> </div> 5. Click OK to save the policy and Commit your changes.
<p>Step 3 Save your changes.</p>	<p>Click Commit.</p>

Set Up or Override a Default Security Profile Group

Use the following options to set up a default security profile group to be used in new security policies, or to override an existing default group. When an administrator creates a new security policy, the default profile group will be automatically selected as the policy's profile settings, and traffic matching the policy will be checked according to the settings defined in the profile group (the administrator can choose to manually select different profile settings if desired). Use the following options to set up a default security profile group or to override your default settings.



If no default security profile exists, the profile settings for a new security policy are set to **None** by default.

Set Up or Override a Default Security Profile Group

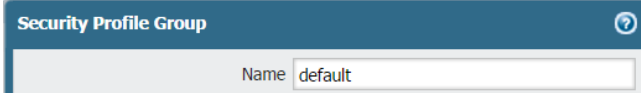
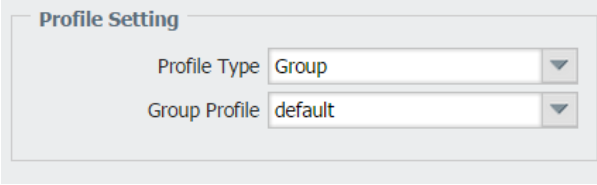
- Create a security profile group.

1. Select **Objects > Security Profile Groups** and Add a new security profile group.
2. Give the profile group a descriptive **Name**, for example, Threats.
3. If the firewall is in Multiple Virtual System Mode, enable the profile to be **Shared** by all virtual systems.
4. Add existing profiles to the group. For details on creating profiles, see [Security Profiles](#).

5. Click **OK** to save the profile group.
6. **Add the security profile group to a security policy.**
7. **Add** or modify a security policy rule and select the **Actions** tab.
8. Select **Group** for the **Profile Type**.
9. In the **Group Profile** drop-down, select the group you created (for example, select the Threats group):

10. Click **OK** to save the policy and **Commit** your changes.

Set Up or Override a Default Security Profile Group

<ul style="list-style-type: none"> Set up a default security profile group. 	<ol style="list-style-type: none"> Select Objects > Security Profile Groups and add a new security profile group or modify an existing security profile group. Name the security profile group <i>default</i>:  Click OK and Commit. Confirm that the <i>default</i> security profile group is included in new security policies by default: <ol style="list-style-type: none"> Select Policies > Security and Add a new security policy. Select the Actions tab and view the Profile Setting fields:  <p>By default, the new security policy correctly shows the Profile Type set to Group and the <i>default</i> Group Profile is selected.</p>
<ul style="list-style-type: none"> Override a default security profile group. 	<p>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).</p>

Enumeration of Rules Within a Rulebase

Each rule within a rulebase is automatically numbered and the ordering adjusts as rules are moved or reordered. When filtering rules to find rules that match the specified filter(s), each rule is listed with its number in the context of the complete set of rules in the rulebase and its place in the evaluation order.

On Panorama, pre-rules and post-rules are independently numbered. When rules are pushed from Panorama to a managed firewall, the rule numbering incorporates hierarchy in pre-rules, device rules, and post-rules within a rulebase and reflects the rule sequence and its evaluation order. The **Preview** option in Panorama offers an ordered list view of the total number of rules on a managed device.

View the Ordered List of Rules Within a Rulebase

- View the numbered list of rules on the firewall.

Select **Policies** and any rulebase under it. For example, **Policies > QoS**. The left-most column in the table displays the rule number.

Virtual System vsys1						
	Name	Source	Destination			
	Name	Zone	Zone	Address	Application	Action
1	allow RIP	any	Zone5	any	rip	✓
2	allow BGP	any	Zone5	any	bgp	✓
3	allow FTP	Zone5	Untrusted	any	ftp	✓
4	allow web-browsing	Zone5	Untrusted	any	web-browsing	✓

- View the numbered list of rules on Panorama.

Select **Policies** and any rulebase under it. For example, **Policies > Security > Pre-rules**.

11 item						
	Name	Location	Tags	Zone	Address	User
1	allow RIP	Shared		any	any	any
2	allow BGP	Shared		any	any	any
3	Trusted_2_All	5050_vsys1		DMZ Trusted Untrusted	any	any
4	DMZ_2_Outside	5050_vsys1		Trusted	any	any
5	Internet_2_Trusted_...	5050_vsys1		Internet	any	any
6	DMZ_2_Trusted_DNS...	5050_vsys1		DMZ	any	any
7	Internet_2_DMZ_svr	5050_vsys1		Internet	any	any
8	allow FTP	5050_vsys1		Zone5	any	any
9	allow web-browsing	5050_vsys1		Zone5	any	any
10	ssl permit	5050_vsys1		Zone5	any	any
11	allow LDAP	5050_vsys1		Zone5	any	any

View the Ordered List of Rules Within a Rulebase (Continued)

- After you push the rules from Panorama, view the complete list of rules with numbers on the managed device. From the web interface of the managed device, select **Policies** and pick any rulebase under it. For example, select **Policies > Security** and view the complete set of numbered rules that will be evaluated on the device.

Virtual System vsys1							
Security							
	Name	Source Zone	Destination Zone	Address	Application	Action	Profile
1	allow RIP	any	Zone5	any	rip	✓	none
2	allow BGP	any	Zone5	any	bgp	✓	none
3	allow FTP	Zone5	Untrusted	any	ftp	✓	none
4	allow web-browsing	Zone5	Untrusted	any	web-browsing	✓	none
5	ssl permit	Zone5	Untrusted	any	ssl	✓	none
6	allow LDAP	Zone5	Untrusted	any	ldap	✓	none
7	tap_policy	Tap_zone	Tap_zone	any	any	✓	
8	vwire_pol	vWire_Trust	vWire_Untrust	any	any	✓	
9	Trusted_2_All	DMZ Trusted Untrusted	Internet	any	any	✓	
10	DMZ_2_Outside	Trusted	DMZ Untrusted	any	any	✓	none
11	temp UserID rule	Trusted	DMZ Untrusted	any	any	✗	none
12	Internet_2_Trusted_...	Internet	Trusted	10.5.204.43	ms-rdp	✓	none
13	DMZ_2_Trusted_DNS...	DMZ	Trusted	40.40.40.9	any	✓	none
14	Internet_2_DMZ_svr	Internet	DMZ	10.5.204.42	any	✓	none
15	post1	Zone5	Untrusted	any	ftp	✓	none
16	post2	Zone5	Internet	any	tftp	✓	none
17	post3	Zone5	Untrusted	any	ssl	✓	none
18	post4	Trusted	Zone5	any	dns	✓	none
19	post5	Zone5	Internet	any	ms-rdp	✓	none
20	allow cisco-nac	any	Zone5	any	cisco-nac	✓	none
21	allow rtsp	any	Zone5	any	rtsp	✓	none

Use Tags to Visually Distinguish Objects

You can tag objects and add color to the tag in order to visually distinguish tagged objects. Tags can be added to the following objects: address objects, address groups, zones, service groups, and policy rules.

The firewall and Panorama support both static tags and dynamic tags, dynamic tags are registered from a variety of sources and are not displayed with the static tags, because dynamic tags are not part of the device configuration. See [Register IP Addresses and Tags Dynamically](#) for information on registering tags dynamically. The tags discussed in this section are statically added and are part of the device configuration.

One or more tags can be applied to objects and to policy rules; a maximum of 64 tags can be applied to an object. Panorama supports a maximum of 10,000 tags that can be apportioned across Panorama (shared and device groups) and the managed devices (including devices with multiple virtual systems).

Tag Objects

Step 1 Create tags.



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 which this object must belong.
3. Click **Add** and enter a **Name** to identify the tag. The maximum length is 127 characters.
4. (Optional) Select **Shared** to create the object in a shared location for access as a shared object in Panorama or for use across all virtual systems in a multiple virtual system firewall.
5. (Optional) Assign one of the 16 predefined colors to the tag. By default, no color is selected.

6. Click **OK** and **Commit** to save the changes.

Step 2 View tags in policy.

1. Select **Policies** and any rulebase under it.
2. Click **Add** to create a policy rule and use the tagged objects you created in Step 1.
3. Verify that the tags are in use.

Source		Destination			
Zone	User	Zone	Address	Application	Service
Trust-VSYS2	known-user	Untrust-VSYS2	FTP	ftp	applica

Tag Objects (Continued)

Step 3 Working with tags.

- Select **Objects > Tags** to perform any of the following operations with tags:
 - Click the link in the **Name** column to edit the properties of a tag.
 - Select a tag in the table, and click **Delete** to remove the tag from the firewall.
 - Click **Clone** to create a duplicate tag with the same properties. A numerical suffix is added to the tag name. For example, FTP-1.
- To apply a tag to an address object, address group, service or service group:
 - Create the object. Click **Add**. For example to create a service group, select **Objects > Service Groups > Add**.
 - Select the tag(s) from the **Tag** drop-down or enter a phrase to create a new tag.



To edit a tag or add color to the tag, select **Objects > Tags**.

Tips and Tricks for Searching for Objects in Security Policy

The policy filtering mechanism is a powerful tool for searching for objects in even the most complex security policy rulebase. The following topics show examples to help you learn to search the security policy rulebase (**Policies > Security**) efficiently so that you can streamline your management tasks:

- ▲ [Search Security Policy Rules for Security Profiles](#)
- ▲ [Search for Disabled Rules in the Security Policy](#)
- ▲ [Search for Log Forwarding Details in Security Policy](#)
- ▲ [Search for Log Details in Security Policy](#)
- ▲ [Search for Scheduled Rules in Security Policy](#)

Search Security Policy Rules for Security Profiles

- ▲ [Search for an Antivirus Profile](#)
- ▲ [Search for a Specific URL Filtering Profile](#)
- ▲ [Search for an Anti-Spyware Profile](#)
- ▲ [Search for a Vulnerability Protection Profile](#)
- ▲ [Search for a File Blocking Profile](#)
- ▲ [Search for a Data Filtering Profile](#)

Search for an Antivirus Profile

Use the `profile-setting/profiles/virus/member` syntax to search for an Antivirus security profile. For example, suppose you have an Antivirus profile named **BlockAll** that you use to block viruses in all FTP, HTTP, IMAP, POP3, SMB, and SMTP application traffic. To search for all security policy rules that include this Antivirus profile you would enter the following search string in the filter text box on the **Policy** tab:

```
profile-setting/profiles/virus/member eq BlockAll
```



The search filter is case sensitive. Therefore, when filtering you must enter the object name exactly as it displays on the **Objects** tab. In this case, filtering on `blockall` or `Blockall` would not yield a match because the exact capitalization of the object is `BlockAll`.

After applying the search filter, the rulebase now only displays security policy rules with the BlockAll Antivirus security profile attached as follows:

profile-setting/profiles/virus/member eq BlockAll												
Source						Destination						
Name	Tag	Zone	Address	User	HP Profile	Zone	Address	Application	Service	Action	Profile	Options
PHLS Web-Browsing ...	Lync ERP-CRM	Trust	any	pancademo\account operators pancademo\finance pancademo\gary.stevenson	any	Untrust	Server 1 bad actors misc test	Webmail bittorrent freegate ssl web-browsing	any	✓	Antivirus Profiles: BlockAll	
Sharepoint access to ...	none	any	any	pancademo\finance steve	any	any	any	sharepoint-c... sharepoint-d...	any	✓		
Browsing	none	Trust	any	pancademo\administrators	any	Untrust	any	ssl web-browsing	any	✓		
test rule	none	Trust	any	pancademo\alicia.hale pancademo\finance	any	Untrust	any	facebook-pos... google-docs nofs	any	✗		
Allow Google docs to ...	none	Trust	any	pancademo\administrators pancademo\carlton.steuer	any	Untrust	any	google-docs	any	✓		
idsfd	none	DC-DMZ	any	pancademo\alicia.hale pancademo\marketing	any	DC-Database	any	facebook-pos... google-docs sharepoint-b... sharepoint-c...	application-d...	✓		
Alert all but block bad	none	Trust-L3	any	pancademo\andy.miller pancademo\finance pancademo\steven.sharma	any	Untrust-L3	any	facebook-apps google-drive-...	any	✗		
Demo	none	Trust-L3	any	any	any	Untrust-L3	any	proxy	any	✗		
Safe p2p	none	Trust-L3	any	pancademo\adam.johnson pancademo\domain admins	any	Untrust-L3	any	safe filesharing	any	✓		
rfep2p	none	Trust-L3	any	pancademo\adam.johnson pancademo\administrators	any	Untrust-L3	any	stp2apps1	any	✓		

Search for a Specific URL Filtering Profile

Use the `profile-setting/profiles/url-filtering/member` syntax to search for an antivirus security profile. For example, to search for all security policy rules that have a URL Filtering profile named **block malware**, you would enter the following search string in the filter text box:

`profile-setting/profiles/url-filtering/member eq "block malware"`



In this case, because the profile name contains a space, you must enclose the name in quotation marks in order for the filter to match policy rules with the block malware URL Filtering profile attached.

After applying the search filter, the rulebase now only displays security policy rules with the block malware URL Filtering security profile attached as follows:

profile-setting/profiles/url-filtering/member eq "block malware"												
Source						Destination						
Name	Tag	Zone	Address	User	HP Profile	Zone	Address	Application	Service	Action	Profile	Options
suffid	none	DC-LPZ	any	pancademo\alicia.hale pancademo\marketing	any	DC-Database	any	facebook-pos... google-docs sharepoint-b... sharepoint-c...	application-d...	✓	URL Filtering Profile: block malware	
facebook	none	Trust	any	any	any	Untrust	any	facebook-base	any	✓		

Search for an Anti-Spyware Profile

Use the `profile-setting/profiles/spyware/member` syntax to search for an Anti-Spyware security profile. For example, to search for all security policy rules that have the **default** Anti-Spyware security profile attached, you would enter the following search string in the filter text box:

```
profile-setting/profiles/spyware/member eq default
```

After applying the search filter, the rulebase now only displays security policy rules with the default Anti-Spyware security profile attached as follows:

The screenshot shows the Palo Alto Networks PAN-OS interface. The search filter is `profile-setting/profiles/spyware/member eq default`. The results table is as follows:

Name	Tag	Zone	Address	User	HP Profile	Zone	Address	Application	Service	Action	Profile	Options
Facebook Hiking	none	Trust	test1	pancademo/marketing pancademo/venedy.Rynn	any	Untrust	Ysa	ic-enterspr facebook facebook-base facebook-pos	SSH service-http	Deny	Anti-Spyware Profiles: default	
Peer to Peer Appli...	LYMC ERP-COM	Trust	333.44.0.0-32	pancademo/administrato pancademo/zany.publishers pancademo/linus.l.srivasa pancademo/julie pancademo/steven.shama	Mobile Wi-Fi	any	any	Encrypted Fu... Gaming Peer to Peer Photos facebook ms-office-365 unknown-asp more...	any	Deny		
nA12	LYMC ERP-COM	Trust	333.44.0.0-32	pancademo/administrato pancademo/zany.publishers pancademo/linus.l.srivasa pancademo/julie pancademo/steven.shama	Mobile Wi-Fi	any	any	Encrypted Fu... Gaming Peer to Peer Photos facebook ms-office-365 unknown-asp more...	any	Deny		
Sharma 1	none	DMZ	any	pancademo/steven.shama	any	any	any	BitTorrent	any	Deny		
facebook-test	none	Trust	any	pancademo/marketing	any	Untrust	any	linkedin-base	any	Deny		
gFDI	none	Management	any	pancademo/fax operators	any	Untrust	any	storage	application-d...	Deny		

Search for a Vulnerability Protection Profile

Use the `profile-setting/profiles/vulnerability/member` syntax to search for a Vulnerability Protection security profile. For example, to search for all security policy rules that have the **strict** Vulnerability Protection security profile attached, you would enter the following search string in the filter text box:

```
profile-setting/profiles/vulnerability/member eq strict
```

After applying the search filter, the rulebase now only displays security policy rules with the strict Vulnerability Protection security profile attached as follows:

The screenshot shows the Palo Alto Networks PAN-OS interface. The search filter is `profile-setting/profiles/vulnerability/member eq strict`. The results table is as follows:

Name	Tag	Zone	Address	User	HP Profile	Zone	Address	Application	Service	Action	Profile	Options
facebook_test	none	Trust	any	any	any	Untrust	any	linkedin-base	application-d...	Deny	Vulnerability Protection Profiles: strict	

Search for a File Blocking Profile

Use the `profile-setting/profiles/file-blocking/member` syntax to search for a File Blocking security profile. For example, suppose you created a File Blocking profile named Wildfire that you will use to forward all PE files to WildFire for analysis. You can quickly verify that all security policy rules have the File Blocking profile attached by searching for all security policy rules that have the **Wildfire** File Blocking security profile attached by entering the following search string in the filter text box:

`profile-setting/profiles/file-blocking/member eq Wildfire`

After applying the search filter, the rulebase now displays all security policy rules that have the Wildfire File Blocking security profile attached as follows:



Search for a Data Filtering Profile

Use the `profile-setting/profiles/data-filtering/member` syntax to search for security policy rules that have a specific Data Filtering security profile attached. For example, suppose you created a Data Filtering profile named **credit cards**. You can search for all security policy rules that have this profile attached by searching for all security policy rules that have the **credit cards** Data Filtering security profile attached by entering the following search string in the filter text box:

`profile-setting/profiles/data-filtering/member eq "credit cards"`

Notice once again that the profile name must be enclosed in quotation marks because it contains a space.

After applying the search filter, the rulebase now displays only security policy rules that have the credit cards Data Filtering security profile attached as follows:



Search for Disabled Rules in the Security Policy

If you want to get a quick view of all security rules that are currently disabled in your policy, you can display these by entering the following command on the search bar on the Security Policy:

```
disabled eq yes
```

The screenshot shows the Palo Alto Networks Security Policy interface. The search bar contains the command 'disabled eq yes'. The results table displays the following data:

Name	Tag	Zone	Address	User	HP Profile	Zone	Address	Application	Service	Action	Profile
Peer-to-Peer Applicat...	LYNC	int-Trust	200.64.0.0/20...	panorama:admins@panorama.com	Mobile WP	any	any	Encrypted Tu...	any	Deny	any
any	LYNC	int-Trust	200.64.0.0/20...	panorama:admins@panorama.com	Mobile WP	any	any	Encrypted Tu...	any	Deny	any
filtering_rndoc	none	int-Trust	any	any	any	int-Trust	any	filtering_rndoc	application/...	Deny	any

Search for Log Forwarding Details in Security Policy

If you want to get a quick view of all security rules that are currently sending their logs using a particular log forwarding profile, you can filter the rulebase using the `logsetting` parameter. For example, to filter for all rules that are forwarding their logs using a log forwarding profile named **panorama**, you would enter the following command on the search bar on the **Policies > Security** page:

```
logsetting eq panorama
```

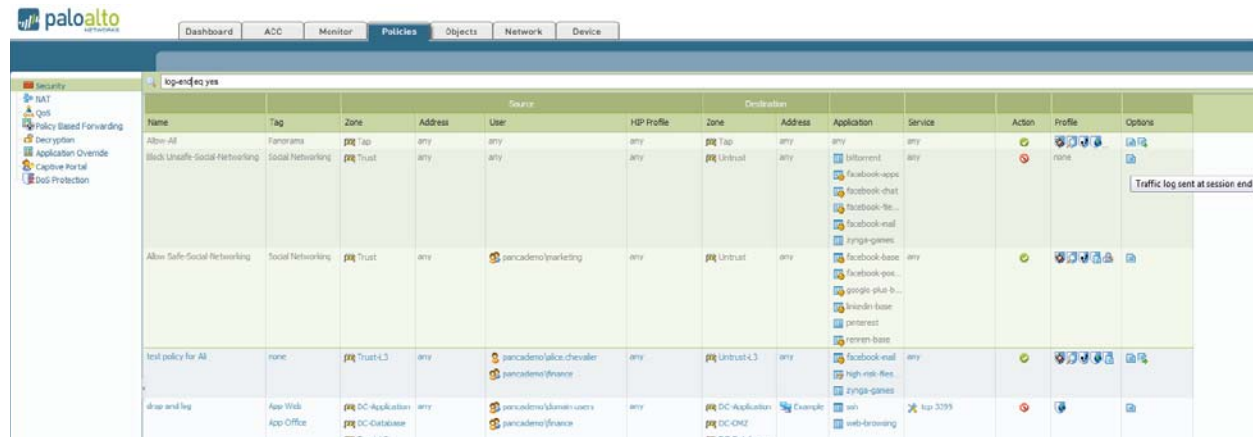
The screenshot shows the Palo Alto Networks Security Policy interface. The search bar contains the command 'logsetting eq panorama'. The results table displays the following data:

Name	Tag	Zone	Address	User	HP Profile	Zone	Address	Application	Service	Action	Profile	Options
allow-all	Panorama	int-Trust	any	any	any	int-Trust	any	any	any	Deny	any	Log Forwarding Profile setting: panorama

Search for Log Details in Security Policy

If you want to get a quick view of all security rules that are currently sending logs, you can filter using the log-end and log-start filters. For example, to view all policy rules that generate logs at the end of a session, you can filter the rulebase using the following filter:

log-end eq yes



Similarly, to view all policy rules that generate logs at the start and the end of a session, you can filter the rulebase as follows:

log-start eq yes



Search for Scheduled Rules in Security Policy

If you want to get a quick view of all security rules that are set to be active based on a particular time schedule, such as during Lunch time, you can display these by entering the following command on the search bar on the **Policies > Security** page:

```
schedule eq "Lunch time"
```



The search field is case sensitive, so you must enter the object name exactly as it appears in the firewall configuration. In addition, you must enclose object names that contain a space in quotation marks.

The screenshot shows the Palo Alto Networks management console interface. The top navigation bar includes 'Dashboard', 'ACC', 'Monitor', 'Policies', 'Objects', 'Network', and 'Device'. The 'Policies' tab is selected, and the 'Security' sub-tab is active. A search bar at the top of the main content area contains the command 'schedule eq "Lunch time"'. Below the search bar, a table displays the results of the search. The table has columns for Name, Tag, Zone, Address, User, HP Profile, Zone, Address, Application, Service, Action, Profile, and Options. The search results show a single rule named 'SAYPE' with a tag of 'Acc Office', a zone of 'DMZ Trust', and a user of 'pancademo/peter.braut'. The rule is associated with the 'HP_Security' profile and is scheduled for 'Lunch time'.

Name	Tag	Zone	Address	User	HP Profile	Zone	Address	Application	Service	Action	Profile	Options
SAYPE	Acc Office	DMZ Trust	192.168.1.2...	pancademo/peter.braut	HP_Security	DMZ Trust	any	HTTP DA	any			Schedule: Lunch time

Use a Dynamic Block List in Policy

The firewall or Panorama typically enforces policy for a source or destination IP address that is defined as a static object on the firewall. If you need agility in enforcing policy for a list of source/destination IP addresses that emerge ad hoc, you can use dynamic block lists.

A dynamic block list is a text file that contains a list of IP addresses, IP ranges, or IP subnets, and is hosted on a web server. The dynamic block list can be used to deny or allow access to the IP addresses (IPv4 and IPv6) included in the list. For example, you can use it as a whitelist for allowing a set of IP addresses or as a blacklist to disallow access to the specified IP addresses. At a configured interval, the firewall dynamically imports the list and enforces policy for the IP addresses included in the list. When you modify the list, the firewall retrieves the updates; a configuration change or commit is not required on the firewall. If the web server is unreachable, the firewall or Panorama will use the last successfully retrieved list for enforcing policy until the connection is restored with the web server that hosts the list.

- ▲ [View the IP Address Limit For Your Firewall Model](#)
- ▲ [Formatting Guidelines for Dynamic Block Lists](#)
- ▲ [Enforce Policy with a Dynamic Block List](#)
- ▲ [View the List of IP addresses in the Dynamic Block List](#)
- ▲ [Retrieve a Dynamic Block List from Web Server](#)

View the IP Address Limit For Your Firewall Model

Irrespective of the firewall model, each firewall supports a maximum of 10 Dynamic Block Lists.

To find the maximum number of addresses, address groups, and IP addresses per group, for your model of the firewall, use the following CLI command:

```
show system state | match cfg.general.max-address
```

For example:

```
admin@PA-7050> show system state | match cfg.general.max-address  
cfg.general.max-address: 80000  
cfg.general.max-address-group: 8000  
cfg.general.max-address-per-group: 500
```



Each list can contain the maximum number of addresses supported by your firewall model minus 300. Up to 300 IP addresses are reserved for internal use on the firewall and are deducted from the available limit. Therefore, in the example above, the firewall can have a maximum of 79,700 IP addresses.

Formatting Guidelines for Dynamic Block Lists

The dynamic block list can include individual IP addresses, subnet addresses (address/mask), or range of IP addresses. In addition, the block list can include comments and special characters such as `*`, `:`, `,`, `;`, `#`, or `/`. The syntax for each line in the list is [IP address, IP/Mask, or IP start range-IP end range] [space] [comment].

Because the firewall ignores incorrectly formatted lines, use these guidelines when defining the list:

- Enter each IP address/range/subnet in a new line; URLs are not supported in this list.
- If you add comments, the comment must be on the same line as the IP address/range/subnet. The space at the end of the IP address is the delimiter that separates a comment from the IP address.

An example:


```
192.168.20.10/32
2001:db8:123:1::1 #test IPv6 address
192.168.20.0/24 ; test internal subnet
2001:db8:123:1::/64 test internal IPv6 range
192.168.20.40-192.168.20.50
```



For an IP address that is blocked, you can display a notification page only if the protocol is HTTP.

Enforce Policy with a Dynamic Block List

Enforce Policy with a Dynamic Block List	
<p>Step 1 Create the dynamic block list and host it on a web server, so that the firewall can retrieve the list for policy evaluation.</p>	<ol style="list-style-type: none"> 1. Create a text file and enter the IP addresses for which you want to enforce policy. For syntax, see Formatting Guidelines for Dynamic Block Lists.
<p>Step 2 Create a dynamic block list object on the firewall.</p>	<ol style="list-style-type: none"> 1. Select Objects > Dynamic Block Lists. 2. Click Add and enter a descriptive Name for the list. 3. (Optional) Select Shared, to share the list with all virtual systems on a device that is enabled for multiple virtual systems. By default, the object is created on the virtual system that is currently selected in the Virtual Systems drop-down. 4. Enter the Source URL (hostname or IP address and the path) for the list you just created on the web server. For example, <code>https://1.2.3.4/DBL_2014</code> 5. Click Test Source URL to verify that the firewall or Panorama can connect to the web server. 6. (Optional) Specify the Repeat frequency at which the firewall or Panorama must retrieve the list. By default the list is retrieved every hour. 7. Click OK to save the changes.

Enforce Policy with a Dynamic Block List	
<p>Step 3 Use the dynamic block list as a source or destination address object in policy.</p> <p> Create separate dynamic block lists if you want to specify allow and deny actions for specific IP addresses.</p> <p>The list can be referenced in any policy type. In this example, we attach it as a destination object in security policy.</p>	<ol style="list-style-type: none"> 1. Select Policies > Security. 2. Click Add and give the rule a descriptive name in the General tab. 3. In the Source tab, select the Source Zone. 4. In the Destination tab, select the Destination Zone and select the dynamic block list as the Destination Address. 5. In the Service/ URL Category tab, make sure the Service is set to application-default. 6. In the Actions tab, set the Action Setting to Allow or Deny. 7. Leave all the other options at the default values. 8. Click OK to save the changes. 9. Commit the changes.
<p>Step 4 Test that the policy action is enforced.</p>	<ol style="list-style-type: none"> 1. Access a IP address that is included in the dynamic block list and verify that action you defined is enforced. 2. Select Monitor > Logs > Traffic and see the log entry for the session. 3. To verify the policy rule that matches a flow, use the following CLI command: <pre>test security-policy-match source <IP_address> destination <IP_address> destination port <port_number> protocol <protocol_number></pre>

View the List of IP addresses in the Dynamic Block List

View the IP Addresses Included in the Dynamic Block List

To view the list of IP addresses that the firewall has retrieved from the web server enter the following CLI command:

```
request system external-list show name <name>
```

For example, for a list named case DBL_2014, the output is:

```
vsys1/DBL_2014:
  Next update at: Wed Aug 27 16:00:00 2014
  IPs:
  1.1.1.1
  1.2.2.2/20 #test China
  192.168.255.0; test internal
  192.168.254.0/24 test internal range
```

Retrieve a Dynamic Block List from Web Server

The firewall or Panorama can be configured to retrieve the list from the web server on an hourly, daily, weekly, or monthly basis. If you have added or deleted IP addresses on the list and need to trigger an immediate refresh, you must use the Command Line Interface.

Retrieve a Dynamic Block List

1. Enter the command: **request system external-list refresh name <name>**

For example, `request system external-list refresh name DBL_2014`

2. Get the job ID for the refresh job using the CLI command: **show jobs all**

Look for the last EBL Refresh job in the list.

3. View the details for the job ID. Use the command **show jobs id <number>**

A message indicating the success or failure displays. For example:

```
admin@PA-200> show jobs id 55
Enqueued                               ID                Type              Status Result Completed
-----
2014/08/26 15:34:14                    55                EBLRefresh        FIN      OK 15:34:40
Warnings:
Details:
```

Register IP Addresses and Tags Dynamically

To mitigate the challenges of scale, lack of flexibility and performance, the architecture in networks today allows for clients, servers and applications to be provisioned, changed, and deleted on demand. This agility poses a challenge for security administrators because they have limited visibility into the IP addresses of the dynamically provisioned clients, servers and the plethora of applications that can be enabled on these virtual resources.

The firewall (hardware-based platforms and the VM-Series) supports the ability to register IP addresses and tags dynamically. The IP addresses and tags can be registered on the firewall directly or registered on the firewall through Panorama. This dynamic registration process can be enabled using any of the following options:

- **User-ID agent for Windows**—In an environment where you've deployed the User-ID agent, you can enable the User-ID agent to monitor up to 100 VMware ESXi and/or vCenter Servers. As you provision or modify virtual machines on these VMware servers, the agent can retrieve the IP address changes and share them with the firewall.
- **VM Information Sources**—Allows you to monitor VMware ESXi and vCenter Server, and the AWS-VPC to retrieve IP address changes when you provision or modify virtual machines on these sources. VM Information Sources polls for a predefined set of attributes and does not require external scripts to register the IP addresses through the XML API. See [Monitor Changes in the Virtual Environment](#).
- **VMware Service Manager** (only available for the integrated NSX solution)—The integrated NSX solution is designed for automated provisioning and distribution of Palo Alto Networks next-generation security services and the delivery of dynamic context-based security policies using Panorama. The NSX Manager updates Panorama with the latest information on the IP addresses and tags associated with the virtual machines deployed in this integrated solution. For information on this solution, see [Set Up a VM-Series NSX Edition Firewall](#).
- **XML API**—The firewall and Panorama support an XML API that uses standard HTTP requests to send and receive data. You can use this API to register IP addresses and tags with the firewall or Panorama. API calls can be made directly from command line utilities such as cURL or using any scripting or application framework that supports REST-based services. Refer to the [PAN-OS XML API Usage Guide](#) for details.

For information on creating and using Dynamic Address Groups, see [Use Dynamic Address Groups in Policy](#).

For the CLI commands for registering tags dynamically, see [CLI Commands for Dynamic IP Addresses and Tags](#).

Monitor Changes in the Virtual Environment

To secure applications and prevent threats in an environment where new users and servers are constantly emerging, your security policy must be nimble. To be nimble, the firewall must be able to learn about new or modified IP addresses and consistently apply policy without requiring configuration changes on the firewall.

This capability is provided by the coordination between the **VM Information Sources** and **Dynamic Address Groups** features on the firewall. The firewall and Panorama provide an automated way to gather information on the virtual machine (or guest) inventory on each monitored source and create policy objects that stay in sync with the dynamic changes on the network.

- ▲ [Enable VM Monitoring to Track Changes on the Virtual Network](#)
- ▲ [Attributes Monitored in the AWS and VMware Environments](#)
- ▲ [Use Dynamic Address Groups in Policy](#)

Enable VM Monitoring to Track Changes on the Virtual Network

VM information sources provides an automated way to gather information on the Virtual Machine (VM) inventory on each monitored source (host); the firewall can monitor the VMware ESXi and vCenter Server, and the AWS-VPC. As virtual machines (guests) are deployed or moved, the firewall collects a predefined set of attributes (or metadata elements) as tags; these tags can then be used to define Dynamic Address Groups (see [Use Dynamic Address Groups in Policy](#)) and matched against in policy.

Up to 10 VM information sources can be configured on the firewall or pushed using Panorama templates. By default, the traffic between the firewall and the monitored sources uses the management (MGT) port on the firewall.



VM Information Sources offers easy configuration and enables you to monitor a predefined set of 16 metadata elements or attributes. See [Attributes Monitored in the AWS and VMware Environments](#) for the list.

Set up the VM Monitoring Agent

Step 1 Enable the VM Monitoring Agent.




Up to 10 sources can be configured for each firewall, or for each virtual system on a multiple virtual systems capable firewall.

If your firewalls are configured in a high availability configuration:

- An active/passive setup, only the active firewall monitors the VM sources.
- An active/active setup, only the firewall with the priority value of primary monitors the VM sources.

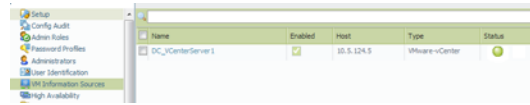
1. Select **Device > VM Information Sources**.
2. Click **Add** and enter the following information:
 - A **Name** to identify the VMware ESX(i) or vCenter Server that you want to monitor.
 - Enter the **Host information for the server**—hostname or IP address and the **Port** on which it is listening.
 - Select the **Type** to indicate whether the source is a **VMware ESX(i)** server or a **VMware vCenter** Server.
 - Add the credentials (**Username** and **Password**) to authenticate to the server specified above.
 - Use the credentials of an administrative user to enable access.
 - (Optional) Modify the **Update interval** to a value between 5-600 seconds. By default, the firewall polls every 5 seconds. The API calls are queued and retrieved within every 60 seconds, so updates may take up to 60 seconds plus the configured polling interval.


- (Optional) Enter the interval in hours when the connection to the monitored source is closed, if the host does not respond. (default: 2 hours, range 2-10 hours)
To change the default value, select the check box to **Enable timeout when the source is disconnected** and specify the value. When the specified limit is reached or if the host cannot be accessed or does not respond, the firewall will close the connection to the source.
- Click **OK**, and **Commit** the changes.
- Verify that the connection **Status** displays as connected .

Set up the VM Monitoring Agent (Continued)

Step 2 Verify the connection status.

Verify that the connection **Status** displays as  connected.



Name	Enabled	Host	Type	Status
DC_CenterServer1	<input checked="" type="checkbox"/>	10.5.124.9	VMware-Center	

If the connection status is pending or disconnected, verify that the source is operational and that the firewall is able to access the source. If you use a port other than the MGT port for communicating with the monitored source, you must change the service route (**Device > Setup > Services**, click the **Service Route Configuration** link and modify the **Source Interface** for the **VM Monitor** service).

Attributes Monitored in the AWS and VMware Environments

Each VM on a monitored ESXi or vCenter server must have VMware Tools installed and running. VMware Tools provide the capability to glean the IP address(es) and other values assigned to each VM.

In order to collect the values assigned to the monitored VMs, the firewall monitors the following predefined set of attributes:

Attributes Monitored on a VMware Source	Attributes Monitored on the AWS-VPC
<ul style="list-style-type: none"> • UUID 	<ul style="list-style-type: none"> • Architecture
<ul style="list-style-type: none"> • Name 	<ul style="list-style-type: none"> • Guest OS
<ul style="list-style-type: none"> • Guest OS 	<ul style="list-style-type: none"> • Image ID
<ul style="list-style-type: none"> • VM State — the power state can be poweredOff, poweredOn, standBy, and unknown. 	<ul style="list-style-type: none"> • Instance ID
<ul style="list-style-type: none"> • Annotation 	<ul style="list-style-type: none"> • Instance State
<ul style="list-style-type: none"> • Version 	<ul style="list-style-type: none"> • Instance Type
<ul style="list-style-type: none"> • Network — Virtual Switch Name, Port Group Name, and VLAN ID 	<ul style="list-style-type: none"> • Key Name
<ul style="list-style-type: none"> • Container Name —vCenter Name, Data Center Object Name, Resource Pool Name, Cluster Name, Host, Host IP address. 	<ul style="list-style-type: none"> • Placement—Tenancy, Group Name, Availability Zone
	<ul style="list-style-type: none"> • Private DNS Name
	<ul style="list-style-type: none"> • Public DNS Name
	<ul style="list-style-type: none"> • Subnet ID
	<ul style="list-style-type: none"> • Tag (key, value) (up to 5 tags supported per instance)
	<ul style="list-style-type: none"> • VPC ID

Use Dynamic Address Groups in Policy

Dynamic address groups are used in policy. They allow you to create policy that automatically adapts to changes—adds, moves, or deletions of servers. It also enables the flexibility to apply different rules to the same server based on *tags* that define its role on the network, the operating system, or the different kinds of traffic it processes.

A dynamic address group uses tags as a filtering criteria to determine its members. The filter uses logical *and* and *or* operators. All IP addresses or address groups that match the filtering criteria become members of the dynamic address group. Tags can be defined statically on the firewall and/or registered (dynamically) to the firewall. The difference between static and dynamic tags is that static tags are part of the configuration on the firewall, and dynamic tags are part of the runtime configuration. This implies that a commit is not required to update dynamic tags; the tags must however be used by Dynamic Address Groups that are referenced in policy, and the policy must be committed on the device.

To dynamically register tags, you can use the XML API or the VM Monitoring agent on the firewall or on the User-ID agent. Each tag is a metadata element or attribute-value pair that is registered on the firewall or Panorama. For example, IP1 {tag1, tag2,.....tag32}, where the IP address and the associated tags are maintained as a list; each registered IP address can have up to 32 tags such as the operating system, the datacenter or the virtual switch to which it belongs. Within 60 seconds of the API call, the firewall registers the IP address and associated tags, and automatically updates the membership information for the dynamic address group(s).

The maximum number of IP addresses that can be registered for each platform is different. Use the following table for specifics on your platform:

Platform	Maximum number of dynamically registered IP addresses
PA-7050, PA-5060, VM-1000-HV	100,000
PA-5050	50,000
PA-5020	25,000
PA-4000 Series, PA-3000 Series	5000
PA-2000 Series, PA-500, PA-200, VM-300, VM-200, VM-100	1000

The following example shows how dynamic address groups can simplify network security enforcement. The example workflow shows how to:

- Enable the VM Monitoring agent on the firewall, to monitor the VMware ESX(i) host or vCenter Server and register VM IP addresses and the associated tags.
- Create dynamic address groups and define the tags to filter. In this example, two address groups are created. One that only filters for dynamic tags and another that filters for both static and dynamic tags to populate the members of the group.
- Validate that the members of the dynamic address group are populated on the firewall.
- Use dynamic address groups in policy. This example uses two different security policies:
 - A security policy for all Linux servers that are deployed as FTP servers; this rule matches on dynamically registered tags.

- A security policy for all Linux servers that are deployed as web servers; this rule matches on a dynamic address group that uses static and dynamic tags.
- Validate that the members of the dynamic address groups are updated as new FTP or web servers are deployed. This ensure that the security rules are enforced on these new virtual machines too.

Use Dynamic Address Groups in Policy

Step 1 Enable VM Source Monitoring.

See [Enable VM Monitoring to Track Changes on the Virtual Network](#).

Step 2 Create dynamic address groups on the firewall.



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**.

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

Use Dynamic Address Groups in Policy (Continued)

Step 3 Use dynamic address groups in policy.



View the [tutorial](#).

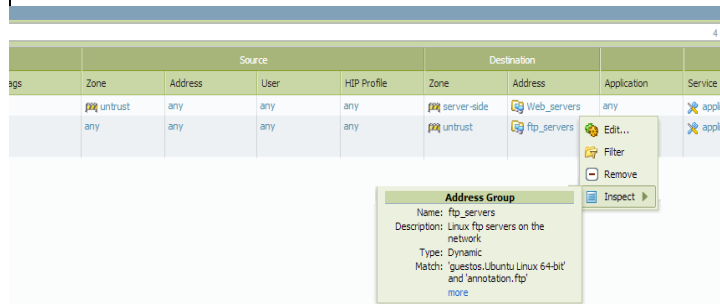
1. Select **Policies > Security**.
2. Click **Add** and enter a **Name** and a **Description** for the policy.
3. Add the **Source Zone** to specify the zone from which the traffic originates.
4. Add the **Destination Zone** at which the traffic is terminating.
5. For the **Destination Address**, select the Dynamic address group you created in [Step 2](#) above.
6. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles to the rule.
7. Repeats Steps 1 through 6 above to create another policy rule.
8. Click **Commit**.

This example shows how to create two policies: one for all access to FTP servers and the other for access to web servers.

	Name	Tags	Zone	Address	User	HP Profile	Zone	Address	Application	Service	Action	Profile	Options
1	Access to web servers		any	any	any	any	untrust	Web_servers	any	application-d...	Allow		
2	Access to FTP servers		any	any	any	any	untrust	ftp_servers	ftp	application.d...	Allow		

Step 4 Validate that the members of the dynamic address group are populated on the firewall.

1. Select **Policies > Security**, and select the rule.
2. Select the drop-down arrow next to the address group link, and select **Inspect**. You can also verify that the match criteria is accurate.



3. Click the **more** link and verify that the list of registered IP addresses is displayed.

Address	Type
10.5.124.45	registered-ip
15.0.0.45	registered-ip
fe80::250:56ff:feb5:beaa	registered-ip
fe80::250:56ff:feb5:cee9	registered-ip

Policy will be enforced for all IP addresses that belong to this address group, and are displayed here.

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</pre> <pre> 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 test registered-ip [<register/unregister>] <ip/netmask> <tag></pre>

Example	CLI Command
View all tags registered from a specific information source.	<pre>show vm-monitor source source-name vmware1 tag all vlanId.4095 vswitch.vSwitch1 host-ip.10.1.5.22 portgroup.TOBEUSED hostname.panserver22 portgroup.VM Network 2 datacenter.ha-datacenter vlanId.0 state.poweredOn vswitch.vSwitch0 vmname.Ubuntu22-100 vmname.win2k8-22-105 resource-pool.Resources vswitch.vSwitch2 guestos.Ubuntu Linux 32-bit guestos.Microsoft Windows Server 2008 32-bit annotation. version.vmx-08 portgroup.VM Network vm-info-source.vmware1 uuid.564d362c-11cd-b27f-271f-c361604dfad7 uuid.564dd337-677a-eb8d-47db-293bd6692f76 Total: 22</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.	<ul style="list-style-type: none"> • 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>

Policy-Based Forwarding

Normally, the firewall uses the destination IP address in a packet to determine the outgoing interface. The firewall uses the routing table associated with the virtual router to which the interface is connected to perform the route lookup. Policy-Based Forwarding (PBF) allows you to override the routing table, and specify the outgoing or *egress* interface based on specific parameters such as source or destination IP address, or type of traffic.


- ▲ [PBF](#)
- ▲ [Create a Policy-Based Forwarding Rule](#)
- ▲ [Use Case: PBF for Outbound Access with Dual ISPs](#)
- ▲ [Use Case: PBF for Routing Traffic Through Virtual Systems](#)

PBF

PBF rules allow traffic to take an alternative path from the next hop specified in the route table, and are typically used to specify an egress interface for security or performance reasons. Let's say your company has two links between the corporate office and the branch office: a cheaper Internet link and a more expensive leased line. The leased line is a high-bandwidth, low-latency link. For enhanced security, you can use PBF to send applications that aren't encrypted traffic, such as FTP traffic, over the private leased line and all other traffic over the Internet link. Or, for performance, you can choose to route business-critical applications over the leased line while sending all other traffic, such as web browsing, over the cheaper link.

Egress Path and Symmetric Return

Using PBF, you can direct traffic to a specific interface on the firewall, drop the traffic, or direct traffic to another virtual system (on systems enabled for multiple virtual systems).



In networks with asymmetric routes, such as in a dual ISP environment, connectivity issues occur when traffic arrives at one interface on the firewall and leaves from another interface. If the route is asymmetrical, where the forward (SYN packet) and return (SYN/ACK) paths are different, the firewall is unable to track the state of the entire session and this causes a connection failure. To ensure that the traffic uses a symmetrical path, which means that the traffic arrives at and leaves from the same interface on which the session was created, you can enable the *Symmetric Return* option.

With symmetric return, the virtual router overrides a routing lookup for return traffic and instead directs the flow back to the MAC address from which it received the SYN packet (or first packet). However, if the destination IP address is on the same subnet as the ingress/egress interface's IP address, a route lookup is performed and symmetric return is not enforced. This behavior prevents traffic from being blackholed.



To determine the next hop for symmetric returns, the firewall uses an Address Resolution Protocol (ARP) table. The maximum number of entries that this ARP table supports is limited by the firewall model and the value is not user configurable. To determine the limit for your model, use the CLI command: **show pbf return-mac all**.

Path Monitoring

Path monitoring allows you to verify connectivity to an IP address so that the firewall can direct traffic through an alternate route, when needed. The firewall uses ICMP pings as *heartbeats* to verify that the specified IP address is reachable.

A monitoring profile allows you to specify the threshold number of heartbeats to determine whether the IP address is reachable. When the monitored IP address is unreachable, you can either disable the PBF rule or specify a *fail-over* or *wait-recover* action. Disabling the PBF rule allows the virtual router to take over the routing decisions.

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
	<input type="checkbox"/> Disable this rule if nexthop/monitor ip is unreachable	<input checked="" type="checkbox"/> Disable this rule if nexthop/monitor ip is unreachable
For an established session	wait-recover —Continue to use egress interface specified in the PBF rule	wait-recover —Continue to use egress interface specified in the PBF rule
	fail-over —Use path determined by routing table (no PBF)	fail-over —Use path determined by routing table (no PBF)
For a new session	wait-recover —Use path determined by routing table (no PBF)	wait-recover —Check the remaining PBF rules. If no match, use the routing table
	fail-over —Use path determined by routing table (no PBF)	fail-over —Check the remaining PBF rules. If no match, use the routing table

Service Versus Applications in PBF

PBF rules are applied either on the first packet (SYN) or the first response to the first packet (SYN/ACK). This means that a PBF rule may be applied before the firewall has enough information to determine the application. Therefore, application-specific rules are not recommended for use with PBF. Whenever possible, use a service object, which is the Layer 4 port (TCP or UDP) used by the protocol or application.

However, if you specify an application in a PBF rule, the firewall performs *App-ID caching*. When an application passes through the firewall for the first time, the firewall does not have enough information to identify the application and therefore cannot enforce the PBF rule. As more packets arrive, the firewall determines the application and creates an entry in the App-ID cache and retains this App-ID for the session. When a new session is created with the same destination IP address, destination port, and protocol ID, the firewall could identify the application as the same from the initial session (based on the App-ID cache) and apply the PBF rule. Therefore, a session that is not an exact match and is not the same application, can be forwarded based on the PBF rule.




Further, applications have dependencies and the identity of the application can change as the firewall receives more packets. Because PBF makes a routing decision at the start of a session, the firewall cannot enforce a change in application identity. YouTube, for example, starts as web-browsing but changes to Flash, RTSP, or YouTube based on the different links and videos included on the page. However with PBF, because the firewall identifies the application as web-browsing at the start of the session, the change in application is not recognized thereafter.



PBF rules cannot be based on domain names; only IP addresses are valid; also, you cannot use custom-applications, application-filters or application groups in PBF rules.


Create a Policy-Based Forwarding Rule

Use a **PBF** rule to direct traffic to a specific egress interface on the firewall, and override the default path for the traffic.

Create a PBF Rule	
<p>Step 1 Create a PBF rule.</p> <p>When creating a PBF rule you must specify a name for the rule, a source zone or interface, and an egress interface. All other components are either optional or have a default value provided.</p>	<ol style="list-style-type: none"> 1. Select Policies > Policy Based Forwarding and click Add. 2. Give the rule a descriptive name in the General tab. 3. In the Source tab, select the following: <ol style="list-style-type: none"> a. Select the Type—Zone or Interface— to which the forwarding policy will be applied, and the relevant zone or interface. <p> PBF is only supported on Layer 3 interfaces.</p> b. (Optional) Specify the Source Address to which PBF will apply. For example, a specific IP address or subnet IP address from which you want to forward traffic to the interface or zone specified in this rule. <p> Use the Negate option to exclude a one or more source IP addresses from the PBF rule. For example, if your PBF rule directs all traffic from the specified zone to the Internet, Negate allows you to exclude internal IP addresses from the PBF rule.</p> <p>The evaluation order is top down. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated.</p> c. (Optional) Add and select the Source User or groups of users to whom the policy applies. 4. In the Destination/Application/Service tab, select the following: <ol style="list-style-type: none"> a. Destination Address. By default the rule applies to Any IP address. Use the Negate option to exclude one or more destination IP addresses from the PBF rule. b. Select the Application(s) or Service(s) that you want to control using PBF. <p> Application-specific rules are not recommended for use with PBF. Whenever possible, use a service object, which is the Layer 4 port (TCP or UDP) used by the protocol or application. For more details, see Service Versus Applications in PBF.</p>

Create a PBF Rule

5. In the **Forwarding** tab, select the following:
 - a. Set the **Action**. The options are as follows:
 - **Forward**—Directs the packet to a specific **Egress Interface**. Enter the **Next Hop** IP address for the packet.
 - **Forward To VSYS**—(On a device enabled for multiple virtual systems) Select the virtual system to which to forward the packet.
 - **Discard**—Drop the packet.
 - **No PBF**—Exclude the packets that match the criteria for source/destination/application/service defined in the rule. Matching packets use the route table instead of PBF; the firewall uses the route table to exclude the matched traffic from the redirected port.

 To trigger the specified action at a daily, weekly or non-recurring frequency, create and attach a **Schedule**.
 - b. (Optional) Enable Monitoring to verify connectivity to a target IP address or to the next hop IP address. Select **Monitor** and attach a monitoring **Profile** (default or custom) that specifies the action when the IP address is unreachable.
 - c. (Optional, required for asymmetric routing environments) Select **Enforce Symmetric Return** and enter one or more IP addresses in the **Next Hop Address List**.

Enabling symmetric return ensures that return traffic (say, from the Trust zone on the LAN to the Internet) is forwarded out through the same interface through which traffic ingresses from the Internet.

Step 2 Save the policies to the running configuration on the device.

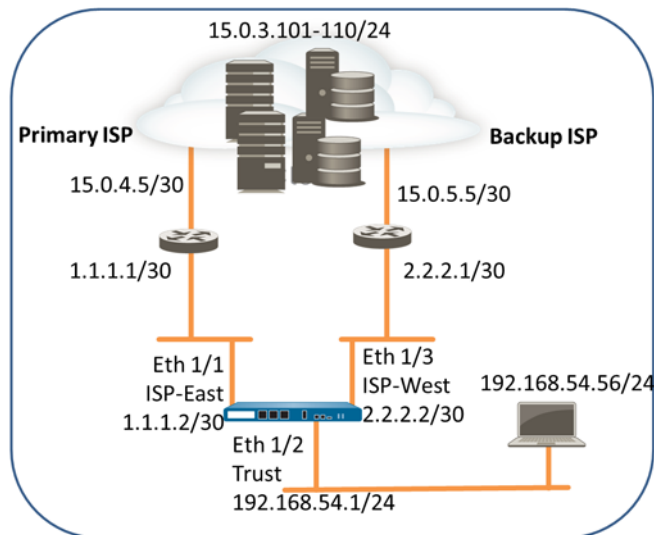
Click **Commit**.
The PBF rule is in effect.

	Name	Source		Destination		Application	Service	Action	Forwarding			Monitoring			
		Zone/Interface	Address	Address	Address				Egress I/F	Next Hop	Enforce Symmetric Return	Profile	Target	Disable If Unreachable	Schedule
1	HTTP to ISP-B	trust	any	HQ-Subnet	any		service-http	forward	ethernet1/4	10.3.4.54	false	default	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.



PBF for Outbound Access with Dual ISPs

Step 1 Configure the ingress and the egress interfaces on the firewall.

Egress interfaces can be in the same zone. In this example we assign the egress interfaces to different zones.

1. Select **Network > Interfaces** and then select the interface you want to configure, for example, Ethernet1/1 and Ethernet1/3. The interface configuration on the firewall used in this example is as follows:
 - Ethernet 1/1 connected to the primary ISP:
 - Zone: ISP-East
 - IP Address:1.1.1.2/30
 - Virtual Router: Default
 - Ethernet 1/3 connected to the backup ISP:
 - Zone: ISP-West
 - IP Address:2.2.2.2/30
 - Virtual Router: Default
 - Ethernet 1/2 is the ingress interface, used by the network clients to connect to the Internet:
 - Zone: Trust
 - IP Address:192.168. 54.1/24
 - Virtual Router: Default
2. To save the interface configuration, click **OK**.

Step 2 On the virtual router, add a static route to the backup ISP.

1. Select **Network > Virtual Router** and then select the **default** link to open the Virtual Router dialog.
2. Select the **Static Routes** tab and click **Add**. Enter a **Name** for the route and specify the **Destination** IP address for which you are defining the static route. In this example, we use 0.0.0.0/0 for all traffic.
3. Select the **IP Address** radio button and set the **Next Hop** IP address for your router that connects to the backup Internet gateway. In this example, 2.2.2.1.
4. Specify a cost metric for the route. In this example, we use 10.



5. Click **OK** twice to save the virtual router configuration.

PBF for Outbound Access with Dual ISPs

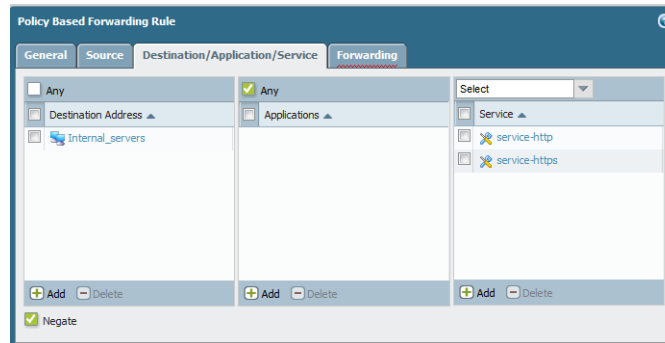
Step 3 Create a PBF rule that directs traffic to the interface that is connected to the primary ISP.

Make sure to exclude traffic destined to internal servers/IP addresses from PBF. Define a negate rule so that traffic destined to internal IP addresses is not routed through the egress interface defined in the PBF rule.

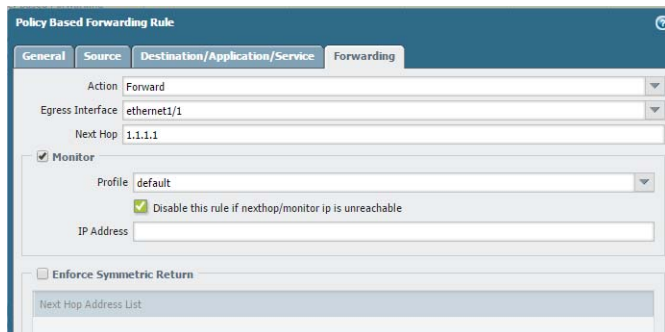
1. Select **Policies > Policy Based Forwarding** and click **Add**.
2. Give the rule a descriptive **Name** in the **General** tab.
3. In the **Source** tab, set the **Source Zone** to Trust.
4. In the **Destination/Application/Service** tab, set the following:
 - a. In the Destination Address section, **Add** the IP addresses or address range for servers on the internal network or create an address object for your internal servers. Select **Negate** to exclude the IP addresses or address object listed above from using this rule.
 - b. In the Service section, **Add** the **service-http** and **service-https** services to allow HTTP and HTTPS traffic to use the default ports. For all other traffic that is allowed by security policy, the default route will be used.



To forward all traffic using PBF, set the Service to **Any**.



5. In the **Forwarding** tab, specify the interface to which you want to forward traffic and enable path monitoring.
 - a. To forward traffic, set the **Action** to **Forward**, and select the **Egress Interface** and specify the **Next Hop**. In this example, the egress interface is ethernet1/1, and the next hop IP address is 1.1.1.1.



PBF for Outbound Access with Dual ISPs

- b. Enable **Monitor** and attach the default monitoring profile, to trigger a failover to the backup ISP. In this example, we do not specify a target IP address to monitor. The firewall will monitor the next hop IP address; if this IP address is unreachable the firewall will direct traffic to the default route specified on the virtual router.
- c. (Required if you have asymmetric routes). Select **Enforce Symmetric Return** to ensure that return traffic from the Trust zone to the Internet is forwarded out on the same interface through which traffic ingressed from the Internet. NAT ensures that the traffic from the Internet is returned to the correct interface/IP address on the firewall.
- d. Click **OK** to save the changes.

	Name	Source	Destination	Service	Action	Forwarding		Monitoring		
		Zone/Interface	Address			Egress I/F	Enforce Symmetric Return	Profile	Target	Disable If Unreachable
1	Use ISP-Primary	Trust	Internal_servers	service-http service-https	forward	ethernet1/1	false	default	1.1.1.2	true

PBF for Outbound Access with Dual ISPs

Step 4 Create NAT rules based on the egress interface and ISP. These rules ensure that the correct source IP address is used for outbound connections.

1. Select **Policies > NAT** and click **Add**.
2. In this example, the NAT rule we create for each ISP is as follows:

NAT for Primary ISP

In the **Original Packet** tab,

- **Source Zone:** Trust
- **Destination Zone:** ISP-West

In the **Translated Packet** tab, under Source Address Translation

- **Translation Type:** Dynamic IP and Port
- **Address Type:** Interface Address
- **Interface:** ethernet1/1
- **IP Address:** 1.1.1.2/30

NAT for Backup ISP

In the **Original Packet** tab,

- **Source Zone:** Trust
- **Destination Zone:** ISP-East

In the **Translated Packet** tab, under Source Address Translation

- **Translation Type:** Dynamic IP and Port
- **Address Type:** Interface Address
- **Interface:** ethernet1/3
- **IP Address:** 2.2.2.2/30

	Name	Original Packet					Translated Packet
		Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Source Translation
1	NAT-Primary ISP	Trust	ISP-West	any	any	any	dynamic-ip-and-port ethernet1/1 1.1.1.2/30
2	NAT-Backup ISP	Trust	ISP-East	any	any	any	dynamic-ip-and-port ethernet1/3 2.2.2.2/30

PBF for Outbound Access with Dual ISPs	
<p>Step 5 Create security policy to allow outbound access to the Internet.</p>	<p>To safely enable applications, create a simple rule that allows access to the Internet and attach the security profiles available on the firewall.</p> <ol style="list-style-type: none"> 1. Select Policies > Security and click Add. 2. Give the rule a descriptive Name in the General tab. 3. In the Source tab, set the Source Zone to Trust. 4. In the Destination tab, Set the Destination Zone to ISP-East and ISP-West. 5. In the Service/ URL Category tab, leave the default application-default. 6. In the Actions tab, complete these tasks: <ol style="list-style-type: none"> a. Set the Action Setting to Allow. b. Attach the default profiles for antivirus, anti-spyware, vulnerability protection and URL filtering, under Profile Setting. 7. Under Options, verify that logging is enabled at the end of a session. Only traffic that matches a security rule is logged.
<p>Step 6 Save the policies to the running configuration on the device.</p>	<p>Click Commit.</p>

		Source	Destination					
	Name	Zone	Zone	Application	Service	Action	Profile	Options
1	Trust2ISP	Trust	ISP-East ISP-West	any	application-default			

PBF for Outbound Access with Dual ISPs

Step 7 Verify that the PBF rule is active and that the primary ISP is used for Internet access.

1. Launch a web browser and access a web server. On the firewall check the traffic log for web-browsing activity.

Traffic is sent through the interface attached to the primary ISP.

Traffic on port 80 is identified as web-browsing.

The security policy that allows the traffic.

Receive Time	Type	From Zone	To Zone	Source	Destination	To Port	Application	Action	Rule	Bytes
11/05 08:24:04	end	Trust	ISP-West	192.168.54.56	204.79.197.200	80	web-browsing	allow	Trust2ISP	3.7

2. From a client on the network, use the ping utility to verify connectivity to a web server on the Internet. and check the traffic log on the firewall.

```
C:\Users\pm-user1>ping 4.2.2.1
Pinging 4.2.2.1 with 32 bytes of data:
Reply from 4.2.2.1: bytes=32 time=34ms TTL=117
Reply from 4.2.2.1: bytes=32 time=13ms TTL=117
Reply from 4.2.2.1: bytes=32 time=25ms TTL=117
Reply from 4.2.2.1: bytes=32 time=3ms TTL=117
Ping statistics for 4.2.2.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 3ms, Maximum = 34ms, Average = 18ms
```

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.

Receive Time	Type	From Zone	To Zone	Source	Destination	To Port	Application	Action	Rule
11/05 09:03:03	end	Trust	ISP-East	192.168.54.56	4.2.2.1	0	ping	allow	Trust2ISP

3. To confirm that the PBF rule is active, use the CLI command **show pbf rule all**

```
admin@PA-NGFW> show pbf rule all
Rule      ID      Rule State Action  Egress IF/VSYS  NextHop
=====  ==  =====  =====  =====  =====
Use ISP-Pr 1    Active    Forward ethernet1/1  1.1.1.1
```

Step 8 Verify that the failover to the backup ISP occurs and that the Source NAT is correctly applied.

1. Unplug the connection to the primary ISP.
2. Confirm that the PBF rule is inactive with the CLI command **show pbf rule all**

```
admin@PA-NGFW> show pbf rule all
Rule      ID      Rule State Action  Egress IF/VSYS  NextHop
=====  ==  =====  =====  =====  =====
Use ISP-Pr 1    Disabled  Forward  ethernet1/1  1.1.1.1
```

3. Access a web server, and check the traffic log to verify that traffic is being forwarded through the backup ISP.

Traffic is sent through the interface attached to the backup ISP.

The security policy that allows the traffic.

Receive Time	Type	From Zone	To Zone	Source	Destination	To Port	Application	Action	Rule
11/05 09:50:44	end	Trust	ISP-East	192.168.54.56	204.79.197.200	443	ssl	allow	Trust2ISP
11/05 09:50:44	end	Trust	ISP-East	192.168.54.56	204.79.197.200	80	web-browsing	allow	Trust2ISP

PBF for Outbound Access with Dual ISPs

4. View the session details to confirm that the NAT rule is working properly.

```
admin@PA-NGFW> show session all
-----
ID Application  State  Type Flag Src[Sport]/Zone/Proto
(translated IP[Port]) Vsys Dst[Dport]/Zone (translated
IP[Port])
-----
87212 ssl ACTIVE FLOW NS 192.168.54.56[53236]/Trust/6
(2.2.2.2[12896]) vsys1 204.79.197.200[443]/ISP-East
(204.79.197.200[443])
```

5. Obtain the session identification number from the output and view the session details. Note that the PBF rule is not used and hence is not listed in the output.

```
admin@PA-NGFW> show session id 87212

Session                87212

c2s flow:
  source:      192.168.54.56 [Trust]
  dst:         204.79.197.200
  proto:       6
  sport:       53236          dport:      443
  state:       ACTIVE        type:       FLOW
  src user:    unknown
  dst user:    unknown

s2c flow:
  source:      204.79.197.200 [ISP-East]
  dst:         2.2.2.2
  proto:       6
  sport:       443          dport:      12896
  state:       ACTIVE        type:       FLOW
  src user:    unknown
  dst user:    unknown

start time           : Wed Nov5 11:16:10 2014
timeout              : 1800 sec
time to live         : 1757 sec
total byte count(c2s) : 1918
total byte count(s2c) : 4333
layer7 packet count(c2s) : 10
layer7 packet count(s2c) : 7
vsys                 : vsys1
application           : ssl
rule                 : Trust2ISP
session to be logged at end : True
session in session ager : True
session synced from HA peer : False
address/port translation : source
nat-rule           : NAT-Backup ISP(vsys1)
  layer7 processing   : enabled
  URL filtering enabled : True
  URL category        : search-engines
  session via syn-cookies : False
  session terminated on host : False
  session traverses tunnel : False
  captive portal session : False
ingress interface : ethernet1/2
egress interface  : ethernet1/3
  session QoS rule    : N/A (class 4)
```

DoS Protection Against Flooding of New Sessions

The following topics describe how to configure DoS protection to better block IP addresses in order to handle high-volume attacks more efficiently.

- ▲ [DoS Protection Against Flooding of New Sessions](#)
- ▲ [Configure DoS Protection Against Flooding of New Sessions](#)
- ▲ [Use the CLI to End a Single Attacking Session](#)
- ▲ [Identify Sessions That Use an Excessive Percentage of the Packet Buffer](#)
- ▲ [Discard a Session Without a Commit](#)

DoS Protection Against Flooding of New Sessions

DoS protection against flooding of new sessions is beneficial against high-volume single-session and multiple-session attacks. In a single-session attack, an attacker uses a single session to target a device behind the firewall. If a Security policy rule allows the traffic, the session is established and the attacker initiates an attack by sending packets at a very high rate with the same source IP address and port number, destination IP address and port number, and protocol, trying to overwhelm the target. In a multiple-session attack, an attacker uses multiple sessions (or connections per second [cps]) from a single host to launch a DoS attack.



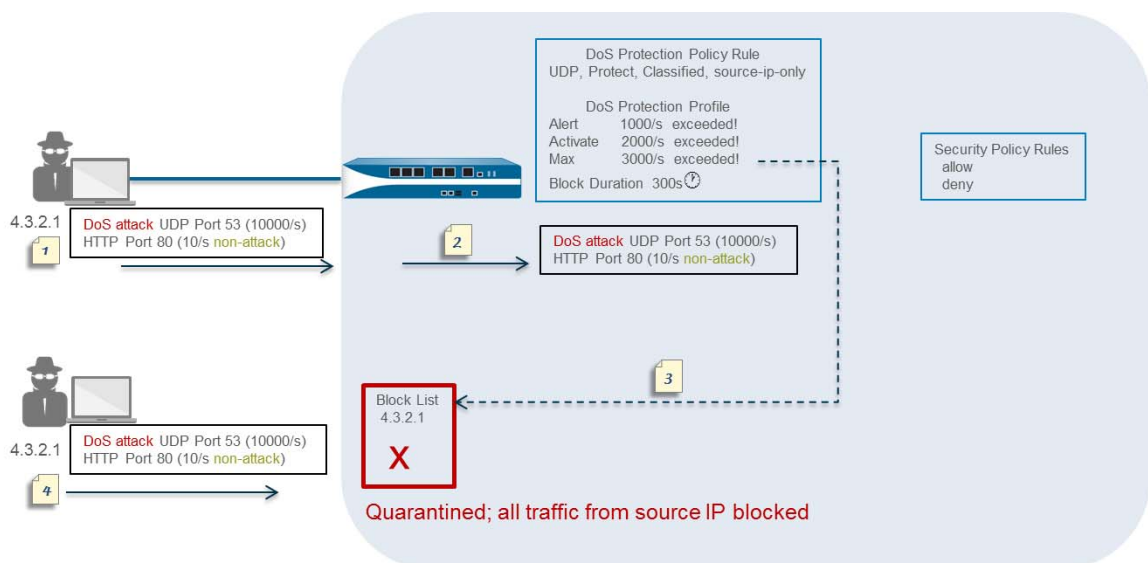
This feature defends only against DoS attacks of new sessions, that is, traffic that has not been offloaded to hardware. An offloaded attack is not protected by this feature. However, this topic describes how you can create a Security policy rule to reset the client; the attacker reinitiates the attack with numerous connections per second and is blocked by the defenses illustrated in this topic.

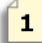
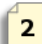



- ▲ Multiple-Session DoS Attack
- ▲ Single-Session DoS Attack

Multiple-Session DoS Attack

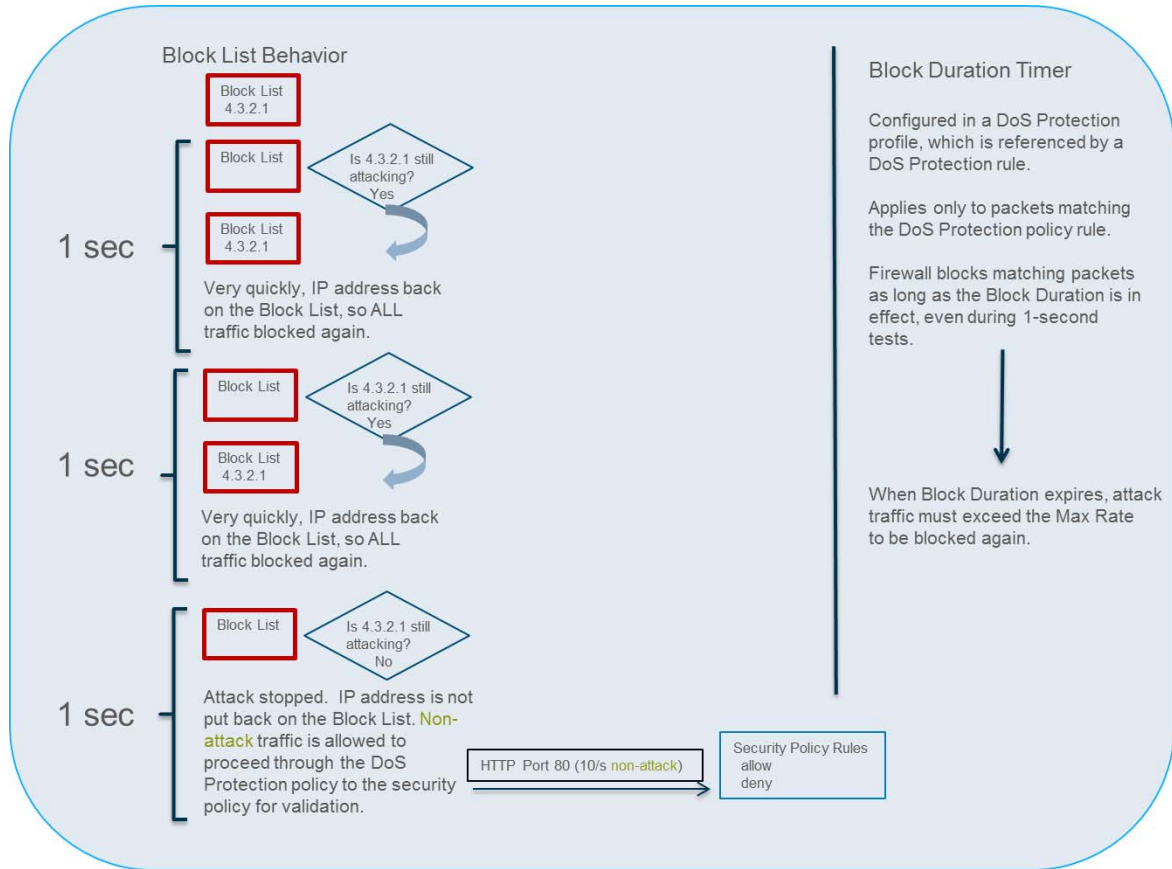
[Configure DoS Protection Against Flooding of New Sessions](#) by configuring a DoS Protection policy rule, which determines the criteria that, when matched by incoming packets, trigger the protect action. The DoS Protection profile counts each new connection toward the Alarm Rate, Activate Rate, and Max Rate thresholds. When the incoming new connections per second exceed the Max Rate allowed, the firewall takes the action specified in the DoS Protection policy rule.

The following figure and table describe how the Security policy rules, DoS Protection policy rules and profile work together in an example.



Sequence of Events as Firewall Quarantines an IP Address	
 1	In this example, an attacker launches a DoS attack at a rate of 10,000 new connections per second to UDP port 53. The attacker also sends 10 new connections per second to HTTP port 80.
 2	<p>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.</p> <p>The DoS rule also specifies the Protect action and Classified, two settings that dynamically put the DoS Protection Profile settings into effect. The DoS Protection Profile specifies that a Max Rate of 3000 packets per second is allowed. When incoming packets match the DoS rule, new connections per second are counted toward the Alert, Activate, and Max Rate thresholds.</p> <p> 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.</p>
 3	<p>The 10,000 new connections per second exceed the Max Rate threshold. When all of the following occur:</p> <ul style="list-style-type: none"> • the threshold is exceeded, • a Block Duration is specified, and • Classified is set to includes source IP address, <p>the firewall puts the offending source IP address on the block list.</p>
 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 do not match the DoS Protection policy criteria (HTTP traffic in this example) through the DoS Protection policy rules to the Security policy for validation. Very few packets, if any, have time to get through because the first attack packet that the firewall receives after the IP address is let off the Block List will match the DoS Protection policy criteria, quickly causing the IP address to be placed back on the block list for another second. The firewall repeats this test each second until the attack stops.
- The firewall blocks all attack traffic from going past the DoS Protection policy rules until the Block Duration expires.

When the attack stops, the firewall does not put the IP address back on the Block List. The firewall allows non-attack traffic to proceed through the DoS Protection policy rules to the Security policy rules for validation. You must configure a Security policy rule because without one, an implicit deny rule denies all traffic.

The block list is based on a source zone and source address combination. This behavior allows duplicate IP addresses to exist as long as they are in different zones belonging to separate virtual routers.

The Block Duration setting in a DoS Protection profile specifies how long the firewall blocks the [offending] packets that exactly match a DoS Protection policy rule. The attack traffic remains blocked until the Block Duration expires, after which the attack traffic must again exceed the Max Rate threshold to be blocked again.



If the attacker uses multiple sessions or bots that initiate multiple attack sessions, the sessions count toward the thresholds in the DoS Protection profile without a Security policy deny rule in place. Hence, a single-session attack requires a Security policy deny rule in order for each packet to count toward the thresholds; a multiple-session attack does not.

Therefore, the DoS protection against flooding of new sessions allows the firewall to efficiently defend against a source IP address while attack traffic is ongoing and to permit non-attack traffic to pass as soon as the attack stops. Putting the offending IP address on the block list allows the DoS protection functionality to take advantage of the block list, which is designed to quarantine all activity. Quarantining the IP address from all activity protects against a modern attacker who attempts a rotating application attack, in which the attacker simply changes applications to start a new attack or uses a combination of different attacks in a hybrid DoS attack.



Beginning with PAN-OS 6.1.7, it is a change in behavior that the firewall places the attacking source IP address on the block list. When the attack stops, non-attack traffic is allowed to proceed to the Security policy rules. The attack traffic that matched the DoS Protection profile and DoS Protection policy rules remains blocked until the Block Duration expires.

Single-Session DoS Attack

A single-session DoS attack typically will not trigger Zone or DoS Protection profiles because they are attacks that are formed after the session is created. These attacks are allowed by the Security policy because a session is allowed to be created, and after the session is created, the attack drives up the packet volume and takes down the target device.

[Configure DoS Protection Against Flooding of New Sessions](#) to protect against flooding of new sessions (single-session and multiple-session flooding). In the event of a single-session attack that is underway, additionally [Use the CLI to End a Single Attacking Session](#).

Configure DoS Protection Against Flooding of New Sessions

Configure DoS Protection Against Flooding of New Sessions

Step 1 (Required for single-session attack mitigation or attacks that have not triggered the DoS Protection policy threshold; optional for multiple-session attack mitigation)

Configure Security policy rules to deny traffic from the attacker's IP address and allow other traffic based on your network needs. You can specify any of the match criteria in a Security policy rule, such as source IP address.



This step is one of the steps typically performed to stop an existing attack. See [Use the CLI to End a Single Attacking Session](#).

- [Set Up Basic Security Policies](#)
- [Components of a Security Policy Rule](#)


Configure DoS Protection Against Flooding of New Sessions (Continued)

Step 2 Configure a DoS Protection profile for flood protection.




Because flood attacks can occur over multiple protocols, as a best practice, activate protection for all of the flood types in the DoS Protection profile.

1. Select **Objects > Security Profiles > DoS Protection** and **Add** a profile **Name**.
2. Select **Classified** as the **Type**.
3. For **Flood Protection**, select the check boxes for all of the following types of flood protection:
 - **SYN Flood**
 - **UDP Flood**
 - **ICMP Flood**
 - **ICMPv6 Flood**
 - **Other IP Flood**
4. (Optional) On each of the flood tabs, change the following thresholds to suit your environment:
 - **Alarm Rate (packets/s)**—Specify the threshold rate (packets per second [pps]) above which a DoS alarm is generated. (Range is 0-2000000; default is 10000.)
 - **Activate Rate (packets/s)**—Specify the threshold rate (pps) above which a DoS response is activated. The DoS response is configured in the **Action** field of the DoS policy where this profile is referenced. When the **Activate Rate** threshold is reached, **Random Early Drop** occurs. (Range is 0-2000000; default is 10000.)
 - **Max Rate (packets/s)**—Specify the threshold rate of incoming packets per second that the firewall allows. When the threshold is exceeded, new packets that arrive are dropped and the **Action** specified in the DoS Policy rule is triggered. (Range is 2-2000000; default is 40000.)



The default threshold values in this step are only starting points and might not be appropriate for your network. You must analyze the behavior of your network to properly set initial threshold values.
5. On each of the flood tabs, specify the **Block Duration** (in seconds), which is the length of time the firewall blocks packets that match the DoS Protection policy rule that references this profile. Specify a value greater than zero. (Range is 1-21600; default is 300.)



Set a low Block Duration value if you are concerned that packets you incorrectly identified as attack traffic will be blocked unnecessarily.

Set a high Block Duration value if you are more concerned about blocking volumetric attacks than you are about incorrectly blocking packets that are not part of an attack.
6. Click **OK**.

Configure DoS Protection Against Flooding of New Sessions (Continued)

<p>Step 3 Configure a DoS Protection policy rule that specifies the criteria for matching the incoming traffic.</p>	<ol style="list-style-type: none"> 1. Select Policies > DoS Protection and Add a Name on the General tab. The name is case-sensitive and can be a maximum of 31 characters, including letters, numbers, spaces, hyphens, and underscores. 2. On the Source tab, choose the Type to be a Zone or Interface, and then Add the zone(s) or interface(s). 3. (Optional) For Source Address, select Any for any incoming IP address to match the rule or Add an address object such as a geographical region. 4. (Optional) For Source User, select any or specify a user. 5. (Optional) Select Negate to match any sources except those you specify. 6. (Optional) On the Destination tab, choose the Type to be a Zone or Interface, and then Add the destination zone(s) or interface(s). For example, enter the security zone you want to protect. 7. (Optional) For Destination Address, select Any or enter the IP address of the device you want to protect. 8. (Optional) On the Option/Protection tab, Add a Service. Select a service or click Service and enter a Name. Select TCP or UDP. Enter a Destination Port. Not specifying a particular service allows the rule to match a flood of any protocol type without regard to an application-specific port. 9. On the Option/Protection tab, for Action, select Protect. 10. Select the Classified check box. 11. For Profile, select the name of the DoS Protection profile you created in Step 2. 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. <ul style="list-style-type: none"> • Specify source-ip-only if you want the firewall to classify only on the source IP address. Because attackers often test the entire network for hosts to attack, source-ip-only is the typical setting for a wider examination. • Specify src-dest-ip-both if you want to protect only against DoS attacks on the server that has a specific destination address and also ensure that every source IP address will not surpass a specific connections-per-second threshold to that server. 13. Click OK.
<p>Step 4 Save the configuration.</p>	<p>Click Commit.</p>

Use the CLI to End a Single Attacking Session

To mitigate a single-session DoS attack, you would still [Configure DoS Protection Against Flooding of New Sessions](#) in advance. At some point after you configure the feature, a session might be established before you realize a DoS attack (from the IP address of that session) is underway. When you see a single-session DoS attack, perform the following task to end the session, so that subsequent connection attempts from that IP address trigger the DoS protection against flooding of new sessions.

Use the CLI to End a Single Attacking Session

Step 1 Identify the source IP address that is causing the attack.

For example, use the firewall Packet Capture feature with a destination filter to collect a sample of the traffic going to the destination IP address. Alternatively, in PAN-OS 7.0 and later, you can use ACC to filter on destination address to view the activity to the target host being attacked.

Step 2 Create a DoS Protection policy rule that will block the attacker's IP address after the attack thresholds are exceeded.

Step 3 Create a Security policy rule to deny the source IP address and its attack traffic.

Step 4 End any existing attacks from the attacking source IP address by executing the `clear session all filter source <ip-address>` operational command.

Alternatively, if you know the session ID, you can execute the `clear session id <value>` command to end that session only.



If you use the `clear session all filter source <ip-address>` command, all sessions matching the source IP address are discarded, which can include both good and bad sessions.

After you end the existing attack session, any subsequent attempts to form an attack session are blocked by the Security policy. The DoS Protection policy counts all connection attempts toward the thresholds. When the Max Rate threshold is exceeded, the source IP address is blocked for the Block Duration, as described in [Sequence of Events as Firewall Quarantines an IP Address](#).

Identify Sessions That Use an Excessive Percentage of the Packet Buffer

When a firewall exhibits signs of resource depletion, it might be experiencing an attack that is sending an overwhelming number of packets. In such events, the firewall starts buffering inbound packets. With PAN-OS 6.1.10 and later PAN-OS 6.1 releases, you can quickly identify the sessions that are using an excessive percentage of the packet buffer and mitigate their impact by discarding them.

Perform the following task on any hardware-based firewall platform (not a VM-Series firewall) to identify, for each slot and dataplane, the packet buffer percentage used, the top five sessions using more than two percent of the packet buffer, and the source IP addresses associated with those sessions. Having that information allows you to take appropriate action.

View Firewall Resource Usage, Top Sessions, and Session Details

Step 1 View firewall resource usage, top sessions, and session details. Execute the following operational command in the CLI (sample output from the command follows):

```
admin@PA-7050> show running resource-monitor ingress-backlogs
-- SLOT:s1, DP:dp1 --
USAGE - ATOMIC: 92% TOTAL: 93%

TOP SESSIONS:
SESS-ID      PCT   GRP-ID  COUNT
6            92%   1        156
              7        1732

SESSION DETAILS
SESS-ID PROTO SZONE SRC          SPORT  DST          DPORT  IGR-IF  EGR-IF  APP
6        6      trust 192.168.2.35 55653  10.1.8.89 80  ethernet1/21 ethernet1/22
undecided
```

The command displays a maximum of the top five sessions that each use 2% or more of the packet buffer.

The sample output above indicates that Session 6 is using 92% of the packet buffers with TCP packets (protocol 6) coming from source IP address 192.168.2.35.

- **GRP-ID**—indicates an internal stage of processing packets.
- **COUNT**—indicates how many packets are in that GRP-ID for that session.
- **SESS-ID**—indicates the global session ID that is used in all other `show session` commands. The global session ID is unique within the firewall.
- **APP**—indicates the application name extracted from the Session information, which can help you determine whether the traffic is legitimate. For example, spoofed packets that use a common TCP or UDP port but indicate an APP of `undecided` are possibly attack traffic.

To restrict the display output:

On a PA-7000 Series platform, you can limit output to a slot, a dataplane, or both. For example:

```
admin@PA-7050> show running resource-monitor ingress-backlogs slot s1
admin@PA-7050> show running resource-monitor ingress-backlogs slot s1 dp dp1
```

On a PA-5000 Series platform, you can limit output to a dataplane. For example:

```
admin@PA-5060> show running resource-monitor ingress-backlogs dp dp1
```

View Firewall Resource Usage, Top Sessions, and Session Details

Step 2 Use the command output to determine whether the source at the source IP address using a high percentage of the packet buffer is sending legitimate or attack traffic.

In the sample output above, a single-session attack is likely occurring. A single session (Session ID 6) is using 92% of the packet buffer for Slot 1, DP 1, and the application at that point is `undecided`.

- If you determine a single user is sending an attack and the traffic is in the slow path, you can [Use the CLI to End a Single Attacking Session](#). At a minimum, you can [Configure DoS Protection Against Flooding of New Sessions](#).
- If the traffic is offloaded to hardware (the traffic is in the fast path), clearing the session does not help because the software then has to handle the barrage of packets. You should [Discard a Session Without a Commit](#) instead.

To see whether a session is offloaded or not, use the `show session id <session-id>` operational command in the CLI as shown in the following example. The `layer7 processing` value indicates completed for sessions offloaded or `enabled` for sessions not offloaded.

```
admin@PA-5060> show session id 68088184

Session      68088184

c2s flow:
  source:    1.1.42.15 [trust]
  dst:       1.2.27.99
  proto:     6
  sport:     55993      dport:    6881
  state:     ACTIVE    type:     FLOW
  src user:  unknown
  dst user:  unknown
  offload:   Yes

s2c flow:
  source:    1.2.27.99 [untrust]
  dst:       1.1.42.15
  proto:     6
  sport:     6881      dport:    55993
  state:     ACTIVE    type:     FLOW
  src user:  unknown
  dst user:  unknown
  offload:   Yes

DP                               : 2
index(local):                   : 979320
start time                       : Tue Oct 27 14:20:09 2015
timeout                           : 1200 sec
time to live                       : 1167 sec
total byte count(c2s)             : 270
total byte count(s2c)             : 270
layer7 packet count(c2s)          : 3
layer7 packet count(s2c)          : 3
vsys                              : vsys1
application                       : bittorrent
rule                               : rule1
session to be logged at end       : True
session in session ager           : True
session updated by HA peer        : False
layer7 processing                  : completed
URL filtering enabled              : False
session via syn-cookies           : False
session terminated on host        : False
session traverses tunnel          : False
captive portal session            : False
ingress interface                 : ethernet1/21
egress interface                  : ethernet1/22
session QoS rule                  : N/A (class 4)
tracker stage l7proc              : ctd decoder bypass
end-reason                        : unknown
```

Discard a Session Without a Commit

With PAN-OS 6.1.10 and later PAN-OS 6.1 releases, you can perform this task to permanently discard a session, such as a session that is overloading the packet buffer. No commit is required; the session is discarded immediately after executing the command. The commands you use for this task apply to both offloaded and non-offloaded sessions.

Discard a Session Without a Commit

Step 1 In the CLI, execute the following operational command on any hardware platform:

```
admin@PA-7050> request session-discard [timeout <seconds>] [reason <reason-string>] id  
<session-id>
```

The default timeout is 3600 seconds.

Step 2 Verify that sessions have been discarded:

```
admin@PA-7050> show session all filter state discard
```



Virtual Systems

This topic describes virtual systems, their benefits, typical use cases, and how to configure them. It also provides links to other topics where virtual systems are documented as they function with other features.

- ▲ [Virtual Systems Overview](#)
- ▲ [Communication Between Virtual Systems](#)
- ▲ [Shared Gateway](#)
- ▲ [Configure Virtual Systems](#)
- ▲ [Configure Inter-Virtual System Communication within the Firewall](#)
- ▲ [Configure a Shared Gateway](#)
- ▲ [Virtual System Functionality with Other Features](#)

Virtual Systems Overview

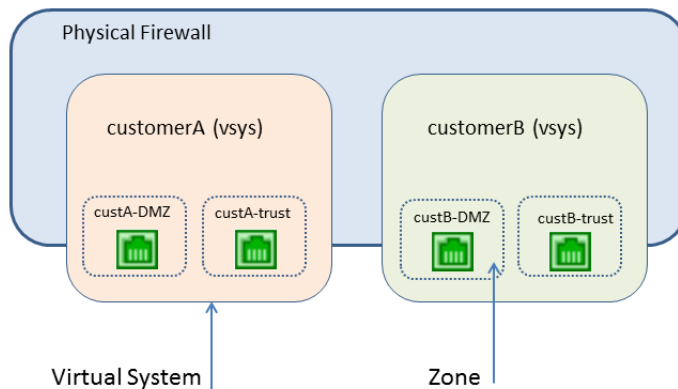
Virtual systems are separate, logical firewall instances within a single physical Palo Alto Networks firewall. Rather than using multiple firewalls, managed service providers and enterprises can use a single pair of firewalls (for high availability) and enable virtual systems on them. Each virtual system (vsys) is an independent, separately-managed firewall with its traffic kept separate from the traffic of other virtual systems.

This topic includes the following:

- ▲ [Virtual System Components and Segmentation](#)
- ▲ [Benefits of Virtual Systems](#)
- ▲ [Use Cases for Virtual Systems](#)
- ▲ [Platform Support and Licensing for Virtual Systems](#)
- ▲ [Restrictions on Virtual Systems](#)
- ▲ [Administrative Roles for Virtual Systems](#)
- ▲ [Shared Objects for Virtual Systems](#)

Virtual System Components and Segmentation

A virtual system is an object that creates an administrative boundary, as shown in the following figure.



A virtual system consists of a set of physical and logical interfaces and subinterfaces (including VLANs and virtual wires), virtual routers, and security zones. You choose the deployment mode(s) (any combination of virtual wire, Layer 2, or Layer 3) of each virtual system. By using virtual systems, you can segment any of the following:

- Administrative access
- The management of all policies (security, NAT, QoS, policy-based forwarding, decryption, application override, captive portal, and DoS protection)
- All objects (such as address objects, application groups and filters, dynamic block lists, security profiles, decryption profiles, custom objects, etc.)
- User-ID

- Certificate management
- Server profiles
- Logging, reporting, and visibility functions

Virtual systems affect the security functions of the firewall, but virtual systems alone do not affect networking functions such as static and dynamic routing. You can segment routing for each virtual system by creating one or more virtual routers for each virtual system, as in the following use cases:

- If you have virtual systems for departments of one organization, and the network traffic for all of the departments is within a common network, you can create a single virtual router for multiple virtual systems.
- If you want routing segmentation and each virtual system's traffic must be isolated from other virtual systems, you can create one or more virtual routers for each virtual system.

Benefits of Virtual Systems

Virtual systems provide the same basic functions as a physical firewall, along with additional benefits:

- **Segmented administration**—Different organizations (or customers or business units) can control (and monitor) a separate firewall instance, so that they have control over their own traffic without interfering with the traffic or policies of another firewall instance on the same physical device.
- **Scalability**—After the physical firewall is configured, adding or removing customers or business units can be done efficiently. An ISP, managed security service provider, or enterprise can provide different security services to each customer.
- **Reduced capital and operational expenses**—Virtual systems eliminate the need to have multiple physical firewalls at one location because virtual systems co-exist on one firewall. By not having to purchase multiple firewalls, an organization can save on the hardware expense, electric bills, and rack space, and can reduce maintenance and management expenses.

Use Cases for Virtual Systems

There are many ways to use virtual systems in a network. One common use case is for an ISP or a managed security service provider (MSSP) to deliver services to multiple customers with a single firewall. Customers can choose from a wide array of services that can be enabled or disabled easily. The firewall's role-based administration allows the ISP or MSSP to control each customer's access to functionality (such as logging and reporting) while hiding or offering read-only capabilities for other functions.

Another common use case is within a large enterprise that requires different firewall instances because of different technical or confidentiality requirements among multiple departments. Like the above case, different groups can have different levels of access while IT manages the firewall itself. Services can be tracked and/or billed back to departments to thereby make separate financial accountability possible within an organization.

Platform Support and Licensing for Virtual Systems

Virtual systems are supported on the PA-2000, PA-3000, PA-4000, PA-5000, and PA-7050 Series firewalls. Each firewall series supports a base number of virtual systems; the number varies by platform. A Virtual Systems license is required in the following cases:

- To support multiple virtual systems on PA-2000 or PA-3000 Series firewalls.
- To create more than the base number of virtual systems supported on a platform.

For license information, see [Activate Licenses and Subscriptions](#). For the base and maximum number of virtual systems supported, see [Compare Firewalls](#) tool.

Multiple virtual systems are not supported on the PA-200, PA-500 or VM-Series firewalls.

Restrictions on Virtual Systems

- Within a Palo Alto Networks firewall, a packet may hop from one virtual system to another virtual system or a shared gateway. A packet may not traverse more than two virtual systems or shared gateways. For example, a packet cannot go from one virtual system to a shared gateway to a second virtual system within the firewall.
- Only one [Shared Gateway](#) is permitted on a firewall.

Administrative Roles for Virtual Systems

A **superuser** administrator can create virtual systems and add a **Device Administrator**, **vsysadmin**, or **vsysreader**. A **Device Administrator** can access all virtual systems, but cannot add administrators. The two types of virtual system administrative roles are:

- **vsysadmin**—Grants full access to a virtual system.
- **vsysreader**—Grants read-only access to a virtual system.

A virtual system administrator can view logs of only the virtual systems assigned to that administrator. Someone with **superuser** or **Device Admin** permission can view all of the logs or select a virtual system to view.

Persons with **vsysadmin** permission can commit configurations for only the virtual systems assigned to them.

Shared Objects for Virtual Systems

If your administrator account extends to multiple virtual systems, you can choose to configure objects (such as an address object) and policies for a specific virtual system or as shared objects, which apply to all of the virtual systems on the firewall. If you try to create a shared object with the same name and type as an existing object in a virtual system, the virtual system object is used.

Communication Between Virtual Systems

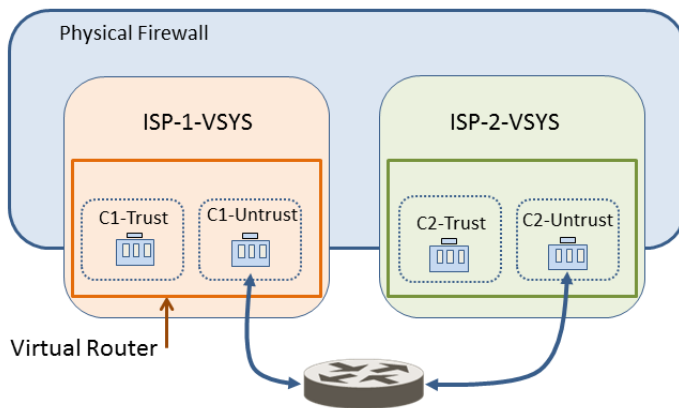
There are two typical scenarios where communication between virtual systems (inter-vsyst traffic) is desirable. In a multi-tenancy environment, communication between virtual systems can occur by having traffic leave the firewall, go through the Internet, and re-enter the firewall. In a single organization 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.

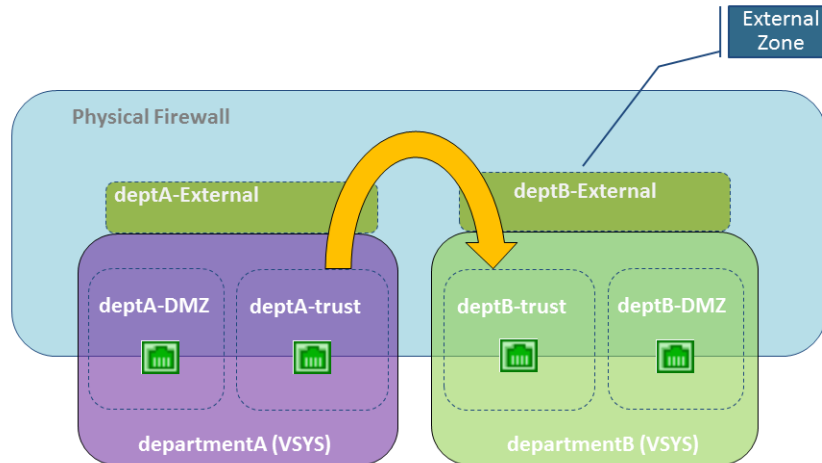
The virtual system administrator configures the security policies needed to allow traffic between two virtual systems. Unlike security zones, an external zone is not associated with an interface; it is associated with a virtual system. The security policy allows or denies traffic between the security (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.



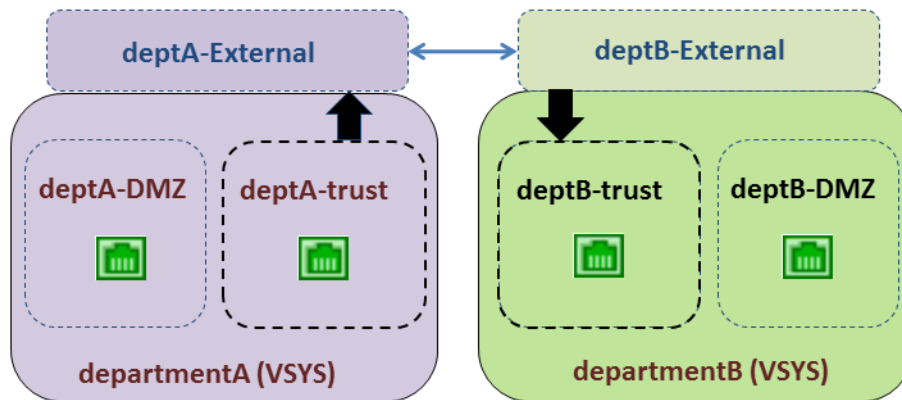
In order to create external zones, the device administrator must configure the virtual systems so that they are *visible* to each other. External zones do not have security policies between them because their virtual systems are visible to each other.

To communicate between virtual systems, the ingress and egress interfaces on the firewall are either assigned to a single virtual router or else they are connected using inter-virtual router static routes. The simpler of these two approaches is to assign all virtual systems that must communicate with each other to a single virtual router.

There might be a reason that the virtual systems need to have their own virtual router, for example, if the virtual systems use overlapping IP address ranges. Traffic can be routed between the virtual systems, but each virtual router must have static routes that point to the other virtual router(s) as the next hop.

Referring to the scenario in the figure above, we have an enterprise with two administrative groups: departmentA and departmentB. The departmentA group manages the local network and the DMZ resources. The departmentB group manages traffic in and out of the sales segment of the network. All traffic is on a local network, so a single virtual router is used. There are two external zones configured for communication between the two virtual systems. The departmentA virtual system has three zones used in security policies: deptA-DMZ, deptA-trust, and deptA-External. The departmentB virtual system also has three zones: deptB-DMZ, deptB-trust, and deptB-External. Both groups can control the traffic passing through their virtual systems.

In order to allow traffic from deptA-trust to deptB-trust, two security policies are required. In the following figure, the two vertical arrows indicate where the security policies (described below the figure) are controlling traffic.



- 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](#)

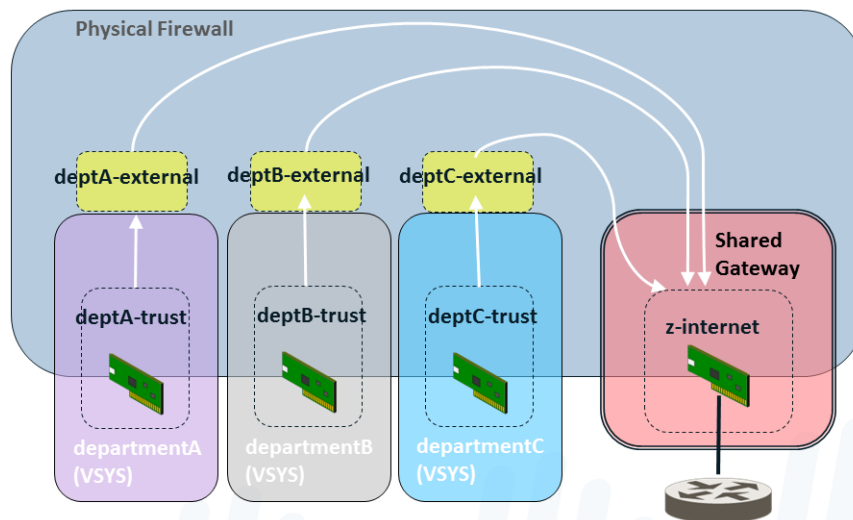
External Zones and Shared Gateway

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 drop-down menu.

A shared gateway is a limited version of a virtual system; it supports NAT, policy-based forwarding (PBF), and DoS policies, but does not support security, QoS, decryption, application override, or captive portal policies.

Networking Considerations for a Shared Gateway

Keep the following in mind while you are configuring a shared gateway.

- The virtual systems in a shared gateway scenario access the Internet through the shared gateway's physical interface, using a single IP address. If the IP addresses of the virtual systems are not globally routable, configure source NAT to translate those addresses to globally-routable IP addresses.
- A virtual router routes the traffic for all of the virtual systems through the shared gateway.
- The default route for the virtual systems should point to the shared gateway.
- Security policies must be configured for each virtual system to allow the traffic between the internal zone and external zone, which is visible to the shared gateway.
- A device administrator should control the virtual router, so that no member of a virtual system can affect the traffic of other virtual systems.

To save configuration time and effort, consider the following advantages of a shared gateway:


- Rather than configure NAT for multiple virtual systems associated with a shared gateway, you can configure NAT for the shared gateway.
- Rather than configure policy-based routing (PBR) for multiple virtual systems associated with a shared gateway, you can configure PBR for the shared gateway.

Configure Virtual Systems


Creating a virtual system requires that you have the following:

- A **superuser** administrative role.
- An interface configured.
- A Virtual Systems license if configuring a PA-2000 or PA-3000 Series firewall, or if creating more than the base number of virtual systems supported on the platform. See [Platform Support and Licensing for Virtual Systems](#).

Perform the following procedure to create and configure virtual systems.

Configure a Virtual System	
<p>Step 1 Enable virtual systems.</p>	<ol style="list-style-type: none"> 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 completing Step 1 will the Device tab display the Virtual Systems and Shared Gateways options.
<p>Step 2 Create a virtual system.</p>	<ol style="list-style-type: none"> 1. Select Device > Virtual Systems, click Add and enter a virtual system ID, which is appended to “vsys.” Range: 1-255.  The default ID is 1, which makes the default virtual system vsys1. This default appears even on platforms that do not support multiple virtual systems. 2. Check the Allow forwarding of decrypted content check box if you want to allow the firewall to forward decrypted content to an outside service. For example, you must enable this option for the firewall to be able to send decrypted content to WildFire for analysis. 3. Enter a descriptive Name for the virtual system. A maximum of 31 alphanumeric, space, and underscore characters is allowed.

Configure a Virtual System	
<p>Step 3 Assign interfaces to the virtual system. The virtual routers, vwires, or VLANs can either be configured already or you can configure them later, at which point you specify the virtual system associated with it. The procedure to configure a virtual router, for example, is in Step 6 below.</p>	<ol style="list-style-type: none"> 1. On the General tab, select a DNS Proxy profile 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: <ul style="list-style-type: none"> • 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.
<p>Step 4 (Optional) Limit the resource allocations for sessions, rules, and VPN tunnels allowed for the virtual system. The flexibility of being able to allocate limits per virtual system allows you to effectively control device resources.</p>	<ol style="list-style-type: none"> 1. On the Resource tab, optionally set limits for a virtual system. There are no default values. <ul style="list-style-type: none"> • Sessions Limit—Range: 1-262144 • Security Rules—Range: 0-2500 • NAT Rules—Range: 0-3000 • Decryption Rules—Range: 0-250 • QoS Rules—Range: 0-1000 • Application Override Rules—Range: 0-250 • Policy Based Forwarding Rules—Range: 0-500 • Captive Portal Rules—Range: 0-1000 • DoS Protection Rules—Range: 0-1000 • Site to Site VPN Tunnels—Range: 0-1024 • Concurrent SSL VPN Tunnels—Range: 0-1024 2. Click OK.
<p>Step 5 Save the configuration.</p>	<p>Click Commit and OK. The virtual system is now an object accessible from the Objects tab.</p>

Configure a Virtual System	
<p>Step 6 Create at least one virtual router for the virtual system in order to make the virtual system capable of networking functions, such as static and dynamic routing.</p> <p>Alternatively, your virtual system might use a VLAN or a virtual wire, depending on your deployment.</p>	<ol style="list-style-type: none"> 1. Select Network > Virtual Routers and Add a virtual router by Name. 2. Add the Interfaces that belong to the virtual router. 3. Click OK.
<p>Step 7 Configure a security zone for each interface in the virtual system.</p>	<p>For at least one interface, create a Layer 3 security zone. See Set Up Interfaces and Zones.</p>
<p>Step 8 Configure the security policies allowing or denying traffic to and from the zones in the virtual system.</p>	<p>See Set Up Basic Security Policies.</p>
<p>Step 9 Save the configuration.</p>	<p>Click Commit and OK.</p> <p> After a virtual system is created, a device administrator can use the CLI to commit a configuration for only a specific virtual system:</p> <pre>PA-5060> commit partial vsys vsys<id></pre>
<p>Step 10 (Optional) View the security policies configured for a virtual system.</p>	<p>Open an SSH session to use the CLI. To view the security policies for a virtual system, in operational mode, use the following commands:</p> <pre>PA-5060> set system setting target-vsys <vsys-id> PA-5060> show running security-policy</pre>

Configure Inter-Virtual System Communication within the Firewall

Perform this task if you have a use case, perhaps within a single enterprise, where you want the virtual systems to be able to communicate with each other within the firewall. Such a scenario is described in [Inter-VSYS Traffic That Remains Within the Firewall](#). This task presumes:


- You completed the task, [Configure Virtual Systems](#).
- When configuring the virtual systems, in the **Visible Virtual System** field, you checked the boxes of all virtual systems that must communicate with each other to be visible to each other.

Configure Inter-Virtual System Communication within the Firewall	
<p>Step 1 Configure an external zone for each virtual system.</p>	<ol style="list-style-type: none"> 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, enter the virtual system that the external zone can reach. 5. Zone Protection Profile—Optionally select a zone protection profile (or configure one later) that provides flood, reconnaissance, or packet-based attack protection. 6. Log Setting—Optionally select a log forwarding profile for forwarding zone protection logs to an external system. 7. Optionally select the Enable User Identification check box to enable User-ID for the external zone. 8. Click OK.
<p>Step 2 Configure the security policies allowing or denying traffic from the internal zones to the external zone of the virtual system, and vice versa.</p>	<ul style="list-style-type: none"> • See Set Up Basic Security Policies. • See Inter-VSYS Traffic That Remains Within the Firewall.
<p>Step 3 Save the configuration.</p>	<p>Click Commit.</p>

Configure a Shared Gateway

Perform this task if you need multiple virtual systems to share an interface (a [Shared Gateway](#)) to the Internet. This task presumes:

- You configured an interface with a globally-routable IP address, which will be the shared gateway.
- You completed the prior task, [Configure Virtual Systems](#). For the interface, you chose the external-facing interface with the globally-routable IP address.
- When configuring the virtual systems, in the **Visible Virtual System** field, you checked the boxes of all virtual systems that must communicate to be visible to each other.

Configure a Shared Gateway	
<p>Step 1 Configure a Shared Gateway.</p>	<ol style="list-style-type: none"> 1. Select Device > Shared Gateway 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 profile if you want to apply DNS proxy rules to the interface. 4. Add an Interface that connects to the outside world. 5. Click OK.
<p>Step 2 Configure the zone for the shared gateway.</p> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="margin-right: 10px;">  </div> <div> <p>When adding objects such as zones or interfaces to a shared gateway, the shared gateway itself will be listed as an available vsys in the VSYS drop-down menu.</p> </div> </div>	<ol style="list-style-type: none"> 1. Select Network > Zones and Add a new zone by Name. 2. For Location, select the shared gateway for which you are creating a zone. 3. For Type, select Layer3. 4. Zone Protection Profile—Optionally select a zone protection profile (or configure one later) that provides flood, reconnaissance, or packet-based attack protection. 5. Log Setting—Optionally select a log forwarding profile for forwarding zone protection logs to an external system. 6. Optionally select the Enable User Identification check box to enable User-ID for the shared gateway. 7. Click OK.
<p>Step 3 Save the configuration.</p>	<p>Click Commit.</p>

Virtual System Functionality with Other Features

Many of the firewall's features and functionality are capable of being configured, viewed, logged, or reported per virtual system. Therefore, virtual systems are mentioned in other relevant locations in the documentation and that information is not repeated here. Some of the specific chapters are the following:

- If you are configuring Active/Passive HA, the two firewalls must have the same virtual system capability (single or multiple virtual system capability). See [High Availability](#).
- To configure QoS for virtual systems, see [Configure QoS for a Virtual System](#).
- For information about configuring a firewall with virtual systems in a virtual wire deployment that uses subinterfaces (and VLAN tags), see the Networking section, [Virtual Wire Subinterfaces](#).