



PAN-OS[®] New Features Guide
Version 6.0

Contact Information

Corporate Headquarters:

Palo Alto Networks
4401 Great America Parkway
Santa Clara, CA 95054

<http://www.paloaltonetworks.com/contact/contact/>

About this Guide

This guide takes you through the configuration and maintenance of your GlobalProtect infrastructure. For additional information, refer to the following resources:

- For information on the additional capabilities and for instructions on configuring the features on the firewall, refer to <https://www.paloaltonetworks.com/documentation>.
- For access to the knowledge base, complete documentation set, discussion forums, and videos, refer to <https://live.paloaltonetworks.com>.
- For contacting support, for information on the support programs, or to manage your account or devices, refer to <https://support.paloaltonetworks.com>.
- For the latest release notes, go to the software downloads page at <https://support.paloaltonetworks.com/Updates/SoftwareUpdates>.

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

Palo Alto Networks, Inc.

www.paloaltonetworks.com

© 2007–2015 Palo Alto Networks, Inc. Palo Alto Networks and PAN-OS are registered trademarks of Palo Alto Networks. All other marks mentioned herein may be trademarks of their respective companies.

Revision Date: March 4, 2016

Table of Contents

Upgrade Your Firewalls to PAN-OS 6.0	7
Before You Upgrade to PAN-OS 6.0	8
Upgrade to PAN-OS 6.0	9
Upgrade Firewalls Using Panorama	9
Upgrade the Firewall to PAN-OS 6.0	13
Upgrade an HA Firewall Pair to PAN-OS 6.0	15
Downgrade from PAN-OS 6.0	18
Downgrade to a Previous Maintenance Release	19
Downgrade to a Previous Feature Release	20
Application Identification Features	21
Support for Hardware Security Modules	22
Set up Connectivity with an HSM	22
Encrypt a Master Key Using an HSM	27
Store Private Keys on an HSM	29
Manage the HSM Deployment	30
Disable the SIP Application-level Gateway (ALG)	31
Content Inspection Features	33
DNS Sinkholing	34
DNS Sinkhole Workflow	34
Configure DNS Sinkholing	35
Verify the Sinkhole Action and Reporting	38
Extended Packet Capture	41
Passive DNS	43
URL Filtering Translation Site Filtering Enhancement	44
URL Filtering Search Engine Cached Site Enhancement	45
URL Filtering Safe Search Enforcement	46
WildFire Report Incorrect Verdict Option	49
WildFire Enhanced File Type and Operating System Support	50
WildFire Analysis Report Enhancement	51
WildFire Submissions Logs Available Without a Subscription	54
WildFire Submissions Log Forwarding	55
GlobalProtect Features	57
GlobalProtect Agent Deployment Customization	58
Push Agent Configuration Settings from the Portal Client Configuration	58
Configure Agent Settings in the Windows Registry or Mac Global plist	60
Deploy Agent Settings Automatically from the Windows Installer (MSIEXEC)	61
Customizable Agent Settings	61

GlobalProtect Agent Update Control	62
Transparent One-Time Password (OTP) Support	64
About the New Authentication Modifier Settings	64
Set up Transparent OTP Authentication	64
Client Certificate Authentication Enforcement	67
HIP Profile Support for Client DLP Products	68
Management Features	73
Enumeration of Rules Within a Rulebase	74
Enhancements in Reports	76
Create Group Activity Reports	76
Disable Predefined Reports	76
Support for Color Coded Tags	78
CLI Find Command	80
Support for Syslog Over TCP and SSL	81
SNMP Resource Monitoring Extensions	84
Enhancement in the Syslog Header	85
Networking Features	87
OSPF v3 Support	88
OSPF Graceful Restart	92
IKE PKI Certificate Authentication for IPsec Site-to-Site VPNs	93
Feature Limitations	93
Prepare a Firewall for IKE PKI Authentication	94
Generate and Authenticate a Certificate	94
Configure a Certificate Profile	95
Configure the IKE Gateway	96
TLS 1.2 Decryption	97
Increase Jumbo Frame Size	98
Decryption Port Mirror	99
Obtain and Install a Decryption Port Mirror License	100
Configure Decryption Port Mirroring	101
Enhanced Use for Address Objects	103
Consolidation of Timers Used in a High Availability Setup	105
Panorama Features	109
Panorama Log Forwarding	110
Scheduled Dynamic Updates	112
Support for Dual URL Filtering Databases	114

User-ID Features	115
User-ID Integration With Syslog.....	116
Configure the PAN-OS Integrated User-ID Agent as a Syslog Listener.....	117
Configure the Windows User-ID Agent as a Syslog Listener.....	123
Support for Custom Terminal Service Solutions	127
About the User-ID XML API Terminal Services Extensions	127
Construct API Scripts to Send User Mapping Information to the Firewall	128
 Virtualization Features	 133
VM Monitoring Agent	134
Dynamic Address Groups	137
Support for the VM-Series Firewall on the Citrix SDX Server	140
About the VM-Series Firewall on the Citrix SDX Server.....	140
Integrate the VM-Series on to the SDX Server.....	142
Supported Deployments.....	142
Install the VM-Series Firewall	147
Support for the VM-Series NSX Edition	149
VM-Series NSX Edition Firewall Overview	149
Deploy the VM-Series NSX Edition Firewall.....	157

Table of Contents



Upgrade Your Firewalls to PAN-OS 6.0

- ▲ [Before You Upgrade to PAN-OS 6.0](#)
- ▲ [Upgrade to PAN-OS 6.0](#)
- ▲ [Downgrade from PAN-OS 6.0](#)

Before You Upgrade to PAN-OS 6.0

You should account for the following information, as needed, before you [Upgrade to PAN-OS 6.0](#):

- Ensure your device is connected to a reliable power source; a loss of power during the upgrade could make the device unusable.
- When upgrading from PAN-OS 5.0 to PAN-OS 6.0, existing threat packet capture (PCAP) files that are stored on the firewall will be deleted (application PCAP files are not impacted). This is due to a change in the way PCAP files are stored. You can perform a bulk export of the PCAP files using the XML API or use the `scp export threat-pcaps` CLI command to export files to an external location before you upgrade the firewall.
- When upgrading firewalls that you manage with Panorama or firewalls that are configured to forward content to a WF-500 appliance, you must first upgrade Panorama and/or the WF-500 appliance before you can upgrade the firewalls.
- If you have asymmetric routes in your network, account for the following to help ensure a successful upgrade to PAN-OS 6.0.5-h3 or later:
 - Before upgrading to PAN-OS 6.0.5-h3 or later, execute the following command to ensure session continuity:

```
set deviceconfig setting tcp asymmetric-path bypass
```
 - If you have asymmetric routes in your network and you also have attached a zone protection profile, you must execute the following command, as well:

```
set network profiles zone-protection-profile <profile-name> asymmetric-path [bypass | global]
```

Upgrade to PAN-OS 6.0

How you upgrade to PAN-OS 6.0 depends on whether you have standalone firewalls or firewalls in a high availability (HA) configuration and whether your firewalls are managed by Panorama. Be sure to review the [Before You Upgrade to PAN-OS 6.0](#) information and the [PAN-OS 6.0 Release Notes](#) and then follow the procedure that matches your configuration:

- ▲ [Upgrade Firewalls Using Panorama](#)
- ▲ [Upgrade the Firewall to PAN-OS 6.0](#)
- ▲ [Upgrade an HA Firewall Pair to PAN-OS 6.0](#)

Upgrade Firewalls Using Panorama

Be sure to review the [Before You Upgrade to PAN-OS 6.0](#) information and the [PAN-OS 6.0 Release Notes](#) and then use the following procedure to upgrade that Panorama manages. This procedure applies to standalone firewalls and firewalls configured in a high availability (HA) configuration.

Upgrade Firewalls Using Panorama

<p>Step 1 Save a backup of the current configuration file on each managed firewall that you plan to upgrade.</p> <p> Although the firewall will automatically create a backup of the configuration, it is a best practice to create a backup prior to upgrade and store it externally.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Operations and click Export Panorama and devices config bundle to generate and export the latest configuration backup of Panorama and of each managed device. 2. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the upgrade.
<p>Step 2 Install the content updates.</p> <p> Make sure the firewalls you plan to upgrade are running content release version 401 or later.</p>	<ol style="list-style-type: none"> 1. Select Panorama > Device Deployment > Dynamic Updates. 2. Click Check Now (located in the lower left-hand corner of the window) to check for the latest updates. If an update is available, the Action column displays a Download link.  <ol style="list-style-type: none"> 3. Download the desired version. After a successful download, the link in the Action column changes from Download to Install. 4. Click Install, select the devices on which you want to install the update, and click OK. When the installation completes, a check mark displays in the Currently Installed column.

Upgrade Firewalls Using Panorama (Continued)	
<p>Step 3 Determine the upgrade path.</p> <p>You cannot skip any major release versions on the path to your desired PAN-OS version. For example, if you want to upgrade from PAN-OS 4.1.8 to PAN-OS 6.0.1, you must:</p> <ul style="list-style-type: none"> • Download and install PAN-OS 5.0.0 and reboot. • Download and install PAN-OS 6.0.1 and reboot. 	<ol style="list-style-type: none"> 1. To access the web interface of the firewall you will upgrade, use the Context drop-down in Panorama or log in to the firewall directly. 2. Select Device > Software. 3. Check which version has a check mark in the Currently Installed column and proceed as follows: <ul style="list-style-type: none"> • If PAN-OS 5.0 or later is currently installed, continue to Step 4. • If a version of PAN-OS prior to 5.0 is currently installed, you must follow the upgrade path to 5.0.0 before upgrading to 6.0. Refer to the Release Notes for your currently installed PAN-OS version for upgrade instructions.
<p>Step 4 Download the software updates.</p>	<ol style="list-style-type: none"> 1. On Panorama, select Panorama > Device Deployment > Software and Check Now for the latest updates. If an update is available, the Action column displays a Download link. 2. Download the files that correspond to the Version to which you want to upgrade and the Platform of the firewalls you are upgrading. You must download a separate installation file for each platform you plan to upgrade. For example, to upgrade your PA-3050 firewalls and PA-5060 firewalls to 6.0.1, download the images that have File Name PanOS_3000-6.0.1 and PanOS_5000-6.0.1. After a successful download, the link in the Action column changes to Install.

Upgrade Firewalls Using Panorama (Continued)

Step 5 Install the software updates on the firewalls.



To avoid downtime when updating the software on HA firewalls, update one peer at a time.

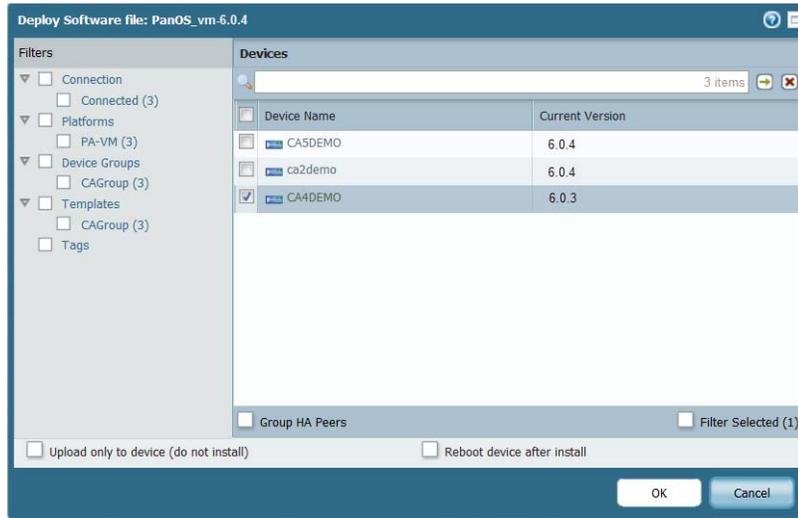
For active/active firewalls, it doesn't matter which HA peer you update first.

For active/passive firewalls, you must update the passive peer first, suspend the active peer (fail over), update the active peer, and then return the active peer to a functional state (fail back).

Perform the steps that apply to your firewall deployment:

- **Non-HA firewalls**—Click the **Install** link for the update in the Action column, select all the firewalls on which you want update the software, select **Reboot device after install**, and click **OK**.
- **Active/active HA firewalls:**
 - a. Click **Install**, clear the **Group HA Peers** check box, select either HA peer, select **Reboot device after install**, and click **OK**. Wait for the firewall to finish rebooting before proceeding.
 - b. Click **Install**, clear the **Group HA Peers** check box, select the HA peer that you didn't update yet, select **Reboot device after install**, and click **OK**.
- **Active/passive HA firewalls**—In this example, the active firewall is named fw1 and the passive firewall is named fw2:
 - a. Click the **Install** link for the update in the Action column, clear the **Group HA Peers** check box, select fw2, select **Reboot device after install**, and click **OK**. Wait for fw2 to finish rebooting before proceeding.
 - b. Access fw1, select **Device > High Availability > Operational Commands**, and click **Suspend local device**.
 - c. Access fw2 and, on the **Dashboard, High Availability** widget, verify that the **Local** firewall state is **active** and the **Peer** firewall is **suspended**.
 - d. Access Panorama, select **Panorama > Device Deployment > Software**, click the **Install** link for the update in the Action column, clear the **Group HA Peers** check box, select fw1, select **Reboot device after install**, and click **OK**. Wait for fw1 to finish rebooting before proceeding.
 - e. Access fw1, select **Device > High Availability > Operational Commands**, and click **Make local device functional**. Wait two minutes before proceeding.
 - f. On fw1, select the **Dashboard** tab and, in the **High Availability** widget, verify that the **Local** firewall state is **active** and the **Peer** firewall is **passive**.

Upgrade Firewalls Using Panorama (Continued)



Step 6 Verify the software and Content Release version running on each managed device.

1. On Panorama, select **Panorama > Managed Devices**.
2. Locate the firewalls and review the content and software versions in the table.

		Status							
Device Group	Device Name	Conn...	Template	Software Version	Apps and Threat	Antivirus	URL Filtering	GlobalProtect Client	WildFire
▼ Branch (1/1 Devices Connected)									
Branch	SupportFW-07	✓	In sync	5.0.0	347-1647	862-1186	4061	1.1.3	15901-23121

Upgrade the Firewall to PAN-OS 6.0

Be sure to review the [Before You Upgrade to PAN-OS 6.0](#) information and the [PAN-OS 6.0 Release Notes](#) and then use the following procedure to upgrade a firewall that is not in an HA configuration to PAN-OS 6.0.

Upgrade PAN-OS																																					
<p>Step 1 Save a backup of the current configuration file.</p> <p> Although the firewall will automatically create a backup of the configuration, it is a best practice to create a backup prior to upgrade and store it externally.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Operations and click Export named configuration snapshot. 2. Select the XML file that contains your running configuration (for example, running-config.xml) and click OK to export the configuration file. 3. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the upgrade. 																																				
<p>Step 2 Make sure the firewall is running Content Release version 401 or later.</p>	<ol style="list-style-type: none"> 1. Select Device > Dynamic Updates. 2. Check the Applications and Threats or Applications 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 desired update and click Download. 5. After the download completes, click Install. 																																				
<p>Step 3 Determine the upgrade path.</p> <p>You cannot skip installing any major release versions on the path to your desired PAN-OS version. Therefore, if you plan to upgrade to a version that is more than one major release away, you must still download, install, and reboot the firewall into all interim PAN-OS versions along the upgrade path.</p> <p>For example, if you want to upgrade from PAN-OS 4.1.8 to PAN-OS 6.0.1, you must:</p> <ul style="list-style-type: none"> • Download and install PAN-OS 5.0.0 and reboot. • Download and install PAN-OS 6.0.1 and reboot. 	<ol style="list-style-type: none"> 1. Select Device > Software. 2. Check which version has a check mark in the Currently Installed column and proceed as follows: <ul style="list-style-type: none"> • If PAN-OS 5.0.0 or later is currently installed, continue to Step 4. • If a version of PAN-OS prior to 5.0.0 is currently installed, you must follow the upgrade path to 5.0.0 before you can upgrade to 6.0. Refer to the release notes for your currently installed PAN-OS version for upgrade instructions. <table border="1" data-bbox="878 1318 1458 1564"> <thead> <tr> <th>Version</th> <th>Size</th> <th>Release Date</th> <th>Downloaded</th> <th>Currently Installed</th> <th>Action</th> </tr> </thead> <tbody> <tr> <td>5.0.0</td> <td>259 MB</td> <td>2012/11/01 19:58:24</td> <td>✓</td> <td></td> <td>Install</td> </tr> <tr> <td>4.1.9</td> <td>169 MB</td> <td>2012/11/05 23:40:31</td> <td></td> <td></td> <td>Download</td> </tr> <tr> <td>4.1.8</td> <td>168 MB</td> <td>2012/09/22 21:01:08</td> <td>✓</td> <td>✓</td> <td>Download</td> </tr> <tr> <td>4.1.8-h3</td> <td>168 MB</td> <td>2012/10/18 23:49:21</td> <td></td> <td></td> <td>Download</td> </tr> <tr> <td>4.1.7</td> <td>152 MB</td> <td>2012/07/29 09:30:58</td> <td></td> <td></td> <td>Download</td> </tr> </tbody> </table>	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
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																																

Upgrade PAN-OS (Continued)	
<p>Step 4 Install PAN-OS 6.0.</p> <p> If your firewall does not have Internet access from the management port, you can download the software update from the Palo Alto Networks Support Site (https://support.paloaltonetworks.com). You can then manually Upload it to your firewall.</p>	<ol style="list-style-type: none"> 1. Click Check Now to check for the latest updates. 2. Locate the version you want to upgrade to and then click Download. 3. After the download completes, click Install. 4. After the install completes, reboot using one of the following methods: <ul style="list-style-type: none"> • 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.
<p>Step 5 Verify that the firewall is passing traffic.</p>	<p>Select Monitor > Session Browser.</p>

Upgrade an HA Firewall Pair to PAN-OS 6.0

Be sure to review the [Before You Upgrade to PAN-OS 6.0](#) information and the [PAN-OS 6.0 Release Notes](#) and then use the following procedure to upgrade a pair of firewalls in a high availability (HA) configuration. This procedure applies to both active/passive and active/active configurations.

Upgrading PAN-OS on firewalls in a HA pair requires that each unit be upgraded separately. Consequently, there is a period of time when PAN-OS versions differ on the individual firewalls in the HA pair and, if a failover occurs before both firewalls are running the same version of PAN-OS, session forwarding could be impacted.



Ensure the devices are connected to a reliable power source as a loss of power during the upgrade could make the devices unusable.

Upgrade PAN-OS

<p>Step 1 Save a backup of the current configuration file.</p> <p> Although the firewall will automatically create a backup of the configuration, it is a best practice to export a backup prior to upgrade and store it externally.</p>	<p>Perform these steps on each firewall in the pair:</p> <ol style="list-style-type: none"> 1. Select Device > Setup > Operations and click Export named configuration snapshot. 2. Select the XML file that contains your running configuration (for example, running-config.xml) and click OK to export the configuration file. 3. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the upgrade.
<p>Step 2 Make sure each device running Content Release version 401 or later.</p>	<ol style="list-style-type: none"> 1. Select Device > Dynamic Updates. 2. Check the Applications and Threats or Applications 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 desired update and click Download. 5. After the download completes, click Install.

Upgrade PAN-OS (Continued)

<p>Step 3 Determine the upgrade path.</p> <p>You cannot skip installing any major release versions on the path to your desired PAN-OS version. Therefore, if you plan to upgrade to a version that is more than one major release away, you must still download, install, and reboot the firewall into all interim PAN-OS versions along the upgrade path.</p> <p>For example, if you want to upgrade from PAN-OS 4.1.8 to PAN-OS 6.0.1, you must:</p> <ul style="list-style-type: none"> • Download and install PAN-OS 5.0.0 and reboot. • Download and install PAN-OS 6.0.1 and reboot. 	<ol style="list-style-type: none"> 1. Select Device > Software. 2. Check which version has a check mark in the Currently Installed column and proceed as follows: <ul style="list-style-type: none"> • If PAN-OS 5.0.0 or later is currently installed, continue to Step 4. • If a version of PAN-OS prior to 5.0.0 is currently installed, follow the upgrade path to 5.0.0 before you can upgrade to 6.0. Refer to the release notes for your currently installed PAN-OS version for upgrade instructions. <table border="1" data-bbox="784 575 1362 821"> <thead> <tr> <th>Version</th> <th>Size</th> <th>Release Date</th> <th>Downloaded</th> <th>Currently Installed</th> <th>Action</th> </tr> </thead> <tbody> <tr> <td>5.0.0</td> <td>259 MB</td> <td>2012/11/01 19:58:24</td> <td>✓</td> <td></td> <td>Install</td> </tr> <tr> <td>4.1.9</td> <td>169 MB</td> <td>2012/11/05 23:40:31</td> <td></td> <td></td> <td>Download</td> </tr> <tr> <td>4.1.8</td> <td>168 MB</td> <td>2012/09/22 21:01:08</td> <td>✓</td> <td>✓</td> <td>Download</td> </tr> <tr> <td>4.1.8-h3</td> <td>168 MB</td> <td>2012/10/18 23:49:21</td> <td></td> <td></td> <td>Download</td> </tr> <tr> <td>4.1.7</td> <td>152 MB</td> <td>2012/07/29 09:30:58</td> <td></td> <td></td> <td>Download</td> </tr> </tbody> </table>	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
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																																
<p>Step 4 Install PAN-OS 6.0 on the passive device (active/passive) or on the active-secondary device (active/active).</p> <p> If your firewall does not have Internet access from the management port, you can download the software update from the Palo Alto Networks Support Site. You can then manually Upload it to your firewall.</p>	<ol style="list-style-type: none"> 1. Click Check Now to check for the latest updates. 2. Locate the version you want to upgrade to and then click Download. 3. After the download completes, click Install. 4. After the install completes, reboot using one of the following methods: <ul style="list-style-type: none"> • 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. After the reboot, the device will not be functional until the active/active-primary device is suspended. 																																				
<p>Step 5 Suspend the active/active-primary firewall.</p>	<ol style="list-style-type: none"> 1. On the active (active/passive) or active-primary (active/active) device, select Device > High Availability > Operational Commands. 2. Click Suspend local device. 3. Select Dashboard and verify that the state of the passive device changes to active in the High Availability widget. 4. Verify that the firewall that took over as active or active-primary is passing traffic by selecting Monitor > Session Browser. 5. If you have session synchronization enabled, run the operational command <code>set session tcp-reject-non-syn no</code>. This will rebuild the session table so that sessions that started prior to the upgrade will continue. 																																				

Upgrade PAN-OS (Continued)

<p>Step 6 Install PAN-OS 6.0 on the other device in the pair.</p> <p> If your firewall does not have Internet access from the management port, you can download the software update from the Palo Alto Networks Support Site. You can then manually Upload it to your firewall.</p>	<ol style="list-style-type: none"> 1. Click Check Now to check for the latest updates. 2. Locate the version you want to upgrade to and then click Download. 3. After the download completes, click Install. 4. After the install completes, reboot using one of the following methods: <ul style="list-style-type: none"> • 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. After the reboot, the device will not be functional until the active (or active-primary) device is suspended. 5. If you configured the firewall to temporarily allow non-syn-tcp traffic in order to enable the firewall to rebuild the session table in Step 5, revert back by running the <code>set session tcp-reject-non-syn yes</code> command.
<p>Step 7 Verify that the devices are passing traffic as expected.</p> <p>In an active/passive deployment, the active device should be passing traffic and in an active/active deployment both devices should be passing traffic.</p>	<p>(Active device(s) only) To verify that the upgrade succeeded and that active devices are passing traffic, run <code>show session all</code>.</p>

Downgrade from PAN-OS 6.0

The way you downgrade from PAN-OS 6.0 depends on whether you are downgrading to a previous feature release (where the first or second digit in the PAN-OS version changes, for example 6.0 to 5.0) or you are downgrading to a maintenance release within the same feature release version (where the third digit in the release version changes, for example, from 6.0.4 to 6.0.2). When downgrading from one feature release to an earlier feature release, the configuration may be migrated to accommodate new features. Therefore, before downgrading you must restore the configuration for the feature release to which you are downgrading. You can downgrade from one maintenance release to another within the same feature release without having to worry about restoring the configuration:

- ▲ [Downgrade to a Previous Maintenance Release](#)
- ▲ [Downgrade to a Previous Feature Release](#)



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

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

Downgrade to a Previous Maintenance Release

Because maintenance releases do not introduce new features, you can downgrade to a previous maintenance release version in the same feature release version without having to restore the previous configuration. A maintenance release is a release in which the third digit in the release version changes, for example a downgrade from 6.0.4 to 6.0.2 is considered a maintenance release downgrade because only the third digit in the release version is different.

Use the following procedure to downgrade to a previous maintenance release within the same feature release version.

Downgrade to a Previous Maintenance Release	
<p>Step 1 Save a backup of the current configuration file.</p> <p> Although the firewall will automatically create a backup of the configuration, it is a best practice to create a backup prior to upgrade and store it externally.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Operations and click Export named configuration snapshot. 2. Select the XML file that contains your running configuration (for example, running-config.xml) and click OK to export the configuration file. 3. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the downgrade.
<p>Step 2 Install the previous maintenance release image.</p> <p> If your firewall does not have Internet access from the management port, you can download the software update from the Palo Alto Networks Support Portal. You can then manually Upload it to your firewall.</p>	<ol style="list-style-type: none"> 1. Select Device > Software and click Check Now. 2. Locate the version to which you want to downgrade. If the image has not yet been downloaded, click Download. 3. After the download completes, click Install. 4. After the install completes, reboot using one of the following methods: <ul style="list-style-type: none"> • 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.

Downgrade to a Previous Feature Release

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



Downgrades from PAN-OS 6.0 to any version earlier than PAN-OS 5.0.5 is not supported because the log management subsystem has been significantly enhanced between PAN-OS 5.0 and PAN-OS 6.0. Because of the changes implemented in the log partitions, a downgrade to PAN-OS 5.0.4 and earlier versions cannot accurately estimate the disk capacity available for storing logs and the log partition could reach maximum capacity without user notification. Such a situation would result in the log partition reaching 100% capacity, resulting in a loss of logs.

Use the following procedure to downgrade to a previous feature release.

Downgrade to a Previous Feature Release	
<p>Step 1 Save a backup of the current configuration file.</p> <p> Although the firewall will automatically create a backup of the configuration, it is a best practice to create a backup prior to upgrade and store it externally.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Operations and click Export named configuration snapshot. 2. Select the XML file that contains your running configuration (for example, running-config.xml) and click OK to export the configuration file. 3. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the downgrade.
<p>Step 2 Install the previous feature release image.</p> <p> Auto-save versions are created when you upgrade to a new release, beginning with PAN-OS 4.1. If you are downgrading to a release prior to PAN-OS 4.1, you may need to do a factory reset and restore the device.</p>	<ol style="list-style-type: none"> 1. Select Device > Software and click Check Now. 2. Locate the version to which you want to downgrade. If the image has not yet been downloaded, click Download. 3. After the download completes, click Install. 4. Select a configuration to load after the device reboots from the Select a Config File for Downgrading drop-down. In most cases, you should select the auto-saved configuration that was created when you upgraded from the release to which you are now downgrading. For example, if you are running PAN-OS 6.0.2 and want to downgrade to PAN-OS 6.0.0, select autosave-6.0.0. 5. After the install completes, reboot using one of the following methods: <ul style="list-style-type: none"> • 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.



Application Identification Features

The following sections describe the new App-ID features and provide instructions for setting them up:

- ▲ [Support for Hardware Security Modules](#)
- ▲ [Disable the SIP Application-level Gateway \(ALG\)](#)

Support for Hardware Security Modules

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.

With this release, HSM clients have been integrated with Palo Alto Networks devices enable enhanced security for the private keys used in SSL 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 now integrated with PA-3000 Series, PA-4000 Series, PA-5000 Series, PA-7050, and VM-Series firewalls and on Panorama (virtual appliance and M-100 appliance) for use with the following HSMs:

- SafeNet Luna SA 5.2.1 or later
- Thales Nshield Connect 11.62 or later



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

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

- ▲ [Set Up Connectivity with a SafeNet Luna SA HSM](#)
- ▲ [Set Up Connectivity with a Thales Nshield Connect HSM](#)

Set Up Connectivity with a SafeNet Luna SA HSM

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



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

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

Set up Connectivity with a SafeNet Luna SA HSM	
<p>Step 1 Configure the firewall to communicate with the SafeNet Luna SA HSM.</p>	<ol style="list-style-type: none"> 1. Log in to the firewall's web interface and select Device > Setup > HSM. 2. Edit the Hardware Security Module Provider section and select SafeNet Luna SA 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.
<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 clear session all CLI command, will clear all existing HSM sessions causing all HSM states to be brought down and then up. During the several seconds required for HSM to recover, all SSL 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.

Set up Connectivity with a SafeNet Luna SA HSM (Continued)	
<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: <pre>client register -c <cl-name> -ip <fw-ip-addr></pre> 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: <pre>client assignpartition -c <cl-name> -p <partition-name></pre> 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.
<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 <code>Authenticated</code> if the HSM is displayed in this table.

Set Up Connectivity with a Thales Nshield Connect HSM

In a Thales Nshield Connect HSM configuration, all key data is stored on the firewall, not on the HSM. A remote filesystem (RFS) is used as a *hub* to sync the HSM and all firewalls configured in the same security world. Each firewall commits its local data to the RFS and can retrieve the data from the RFS by syncing to it. For example, without the RFS if firewall 1 generates a key, only firewall 1 can access the key data. If firewall 1 commits to an RFS and firewall 2 syncs to the same RFS, firewall 2 can also access the same key data.



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"> From the firewall web interface, select Device > Setup > HSM and edit the Hardware Security Module Provider section. Select Thales Nshield Connect as the Provider Configured. Click Add and enter a Module Name. This can be any ASCII string up to 31 characters in length. 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. Enter the IPv4 address of the Remote Filesystem Address. Click OK and Commit.
<p>Step 2 (Optional) Configure a service route for HSM. By default, the firewall uses the Management Interface to communicate with the HSM. To use a different interface, you must configure a service route.</p>	<ol style="list-style-type: none"> Select Device > Setup > Services. Select Service Route Configuration from the Services Features area. Select Customize from the Service Route Configuration area. Select the IPv4 tab. Select HSM from the Service column. Select an interface to use for HSM from the Source Interface drop-down.  If you select a dataplane connected port for HSM, issuing the clear session all CLI command, will clear all existing HSM sessions causing all HSM states to be brought down and then up. During the several seconds required for HSM to recover, all SSL operations will fail. Click OK and Commit.

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

<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. <pre>Client configuration Please enter your client IP address 0.0.0.0 Cancel Next</pre> <ol style="list-style-type: none"> 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.
<p>Step 4 Set up the remote filesystem to accept connections from the firewall.</p>	<ol style="list-style-type: none"> 1. Log in to the remote filesystem (RFS) from a Linux client. 2. Obtain the electronic serial number (ESN) and the hash of the K_{NETI} key. The K_{NETI} key authenticates the module to clients: <pre>anonkneti <ip-address></pre> where ip-address is the IP address of the HSM. <p>The following is an example:</p> <pre>anonkneti 192.0.2.1 B1E2-2D4C-E6A2 5a2e5107e70d525615a903f6391ad72b1c03352c</pre> In this example, B1E2-2D4C-E6A2 is the ESN and 5a2e5107e70d525615a903f6391ad72b1c03352c is the hash of the K_{NETI} key. 3. Use the following command from a superuser to perform the remote filesystem setup: <pre>rfs-setup --force <ip-address> <ESN> <hash-kneti-key></pre> where <ip-address> is the IP address of the HSM, <ESN> is the electronic serial number (ESN) and <hash-kneti-key> is the hash of the K_{NETI} key. <p>The following example uses the values obtained in this procedure:</p> <pre>rfs-setup --force <192.0.2.1> <B1E2-2D4C-E6A2> <5a2e5107e70d525615a903f6391ad72b1c03352c></pre> 4. Use the following command to permit client submit on the Remote Filesystem: <pre>rfs-setup --gang-client --write-noauth <FW-IPaddress></pre> where <FW-IPaddress> 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"> 1. From the firewall's web interface, select Device > Setup > HSM. 2. Select Setup Hardware Security Module in the Hardware Security Operations area. 3. Click OK. <p>The firewall attempts to perform an authentication with the HSM and displays a status message.</p> <ol style="list-style-type: none"> 4. Click OK.

Set up Connectivity with a Thales Nshield Connect HSM (Continued)	
<p>Step 6 Synchronize the firewall with the remote filesystem.</p>	<ol style="list-style-type: none"> 1. Select the Device > Setup > HSM. 2. Select Synchronize with Remote Filesystem in the Hardware Security Operations section.
<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. 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 way you configure master key encryption depends on which HSM you are using:

- ▲ [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 Luna SA and Thales Nshield Connect HSMs.

Refresh the Master Key Encryption

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

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

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

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

The old wrapping key is not deleted by this command.

Store Private Keys on an HSM

For added security, the private keys used to enable SSL decryption—both SSL forward proxy and SSL inbound inspection—can be secured with an HSM as follows:

- **SSL forward proxy**—The private key in the CA certificate that is used to sign certificates in SSL forward proxy operations can be stored on the HSM. The firewall will then send the certificates it generates to the HSM for signing before forwarding them on to the client.
- **SSL inbound inspection**—The private keys for the internal servers for which you are doing SSL inbound inspection can be stored on the HSM.

For instructions on importing the private keys onto the HSM, refer to the documentation from your HSM provider. After the required keys are on the HSM, you can configure the firewall to locate the keys as described in the following sections:

Store Private Keys on an HSM	
Step 1 Import the private keys used in your SSL forward proxy and/or SSL inbound inspection deployments onto the HSM.	For instructions on importing the private keys onto the HSM, refer to the documentation from your HSM provider.
Step 2 (Thales Nshield Connect only) Sync the key data from the HSM remote file system to the firewall.	<ol style="list-style-type: none"> 1. From the firewall's web interface, select Device > Setup > HSM. 2. Select Synchronize with Remote Filesystem in the Hardware Security Operations section.
Step 3 Import the certificate(s) that correspond to the private key(s) you are storing on the HSM onto the firewall.	<ol style="list-style-type: none"> 1. From the firewall's web interface, select Device > Certificate Management > Certificates > Device Certificates. 2. Click Import. 3. Enter the Certificate Name. 4. Enter filename of the Certificate File you imported to the HSM. 5. Select the appropriate file File Format from the drop-down. 6. Select the Private Key resides on Hardware Security Module check box. 7. Click OK and Commit.
Step 4 (Forward trust certificates only) Enable the certificate for use in SSL Forward Proxy.	<ol style="list-style-type: none"> 1. Select the Device > Certificate Management > Certificates > Device Certificates. 2. Locate the certificate you imported in Step 3. 3. Select the Forward Trust Certificate check box.
Step 5 Verify that the certificate has been successfully imported to the firewall.	<ol style="list-style-type: none"> 1. Select Device > Certificate Management > Certificates > Device Certificates. 2. Locate the certificate you imported in Step 3. 3. In the Key column notice the following: <ul style="list-style-type: none"> If a Lock icon is displayed, the private key for the certificate can be found on the HSM. If an Error icon is displayed, the private key is not imported to the HSM or the HSM is not properly authenticated or connected.

Manage the HSM Deployment

Manage HSM	
<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, 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.

Disable the SIP Application-level Gateway (ALG)

The Palo Alto Networks firewall uses the Session Initiation Protocol (SIP) application-level gateway (ALG) to open dynamic pinholes in the firewall where NAT is enabled. However, some applications—such as VoIP—have NAT intelligence embedded in the client application. In these cases, the SIP ALG on the firewall can interfere with the signaling sessions and cause the client application to stop working.

One solution to this problem is to define an Application Override Policy for SIP, but using this approach disables the App-ID and threat detection functionality. A better approach is to disable the SIP ALG, which does not disable App-ID or threat detection.

The following procedure describes how to disable the SIP ALG.

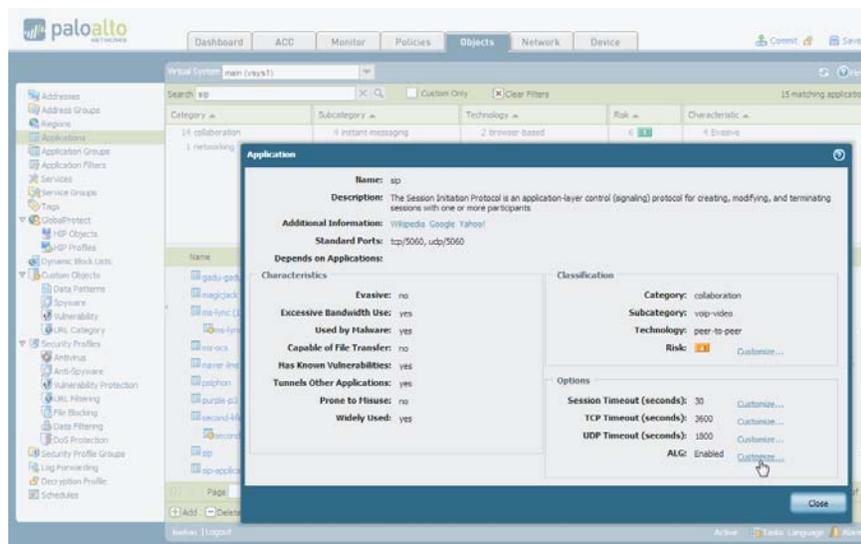
Disable the SIP ALG

Step 1 Select **Objects > Applications**.

Step 2 Select the **sip** application.

You can type **sip** in the **Search** box to help find the sip application.

Step 3 Select **Customize...** for **ALG** in the Options section of the Application dialog box.



Step 4 Select the **Disable ALG** check box in the Application - sip dialog box and click **OK**.



Step 5 **Close** the Application dialog box and **Commit** the change.



Content Inspection Features

The following sections describe the new Content Inspection features and provide instructions for setting them up:

- ▲ [DNS Sinkholing](#)
- ▲ [Extended Packet Capture](#)
- ▲ [Passive DNS](#)
- ▲ [URL Filtering Translation Site Filtering Enhancement](#)
- ▲ [URL Filtering Search Engine Cached Site Enhancement](#)
- ▲ [URL Filtering Safe Search Enforcement](#)
- ▲ [WildFire Report Incorrect Verdict Option](#)
- ▲ [WildFire Enhanced File Type and Operating System Support](#)
- ▲ [WildFire Analysis Report Enhancement](#)
- ▲ [WildFire Submissions Logs Available Without a Subscription](#)
- ▲ [WildFire Submissions Log Forwarding](#)

DNS Sinkholing

The DNS sinkhole action that you can enable in [Anti-Spyware profiles](#) enables the firewall to forge a response to a DNS query for a known malicious domain, causing the malicious domain name to resolve to an IP address that you define. This feature can be used to identify infected hosts on the protected network using DNS traffic in situations where the firewall cannot see the infected client's DNS query (that is, the firewall cannot see the originator of the DNS query). In a typical deployment where the firewall is north of the local DNS server, the threat log will identify the local DNS resolver as the source of the traffic rather than the actual infected host. Sinkholing malware DNS queries solves this visibility problem by forging responses to the client host queries directed at malicious domains, so that clients attempting to connect to malicious domains (for command-and-control, for example) will instead attempt to connect to a sinkhole IP address that you define. Infected hosts can then be easily identified in the traffic logs because any hosts that attempt to connect to the sinkhole IP address are most likely infected with malware.

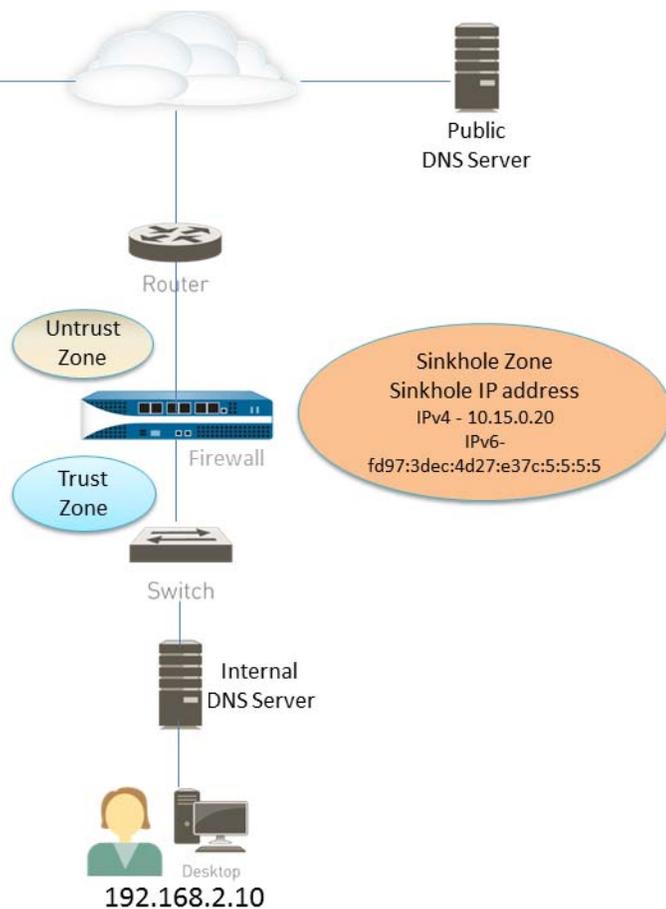
- ▲ [DNS Sinkhole Workflow](#)
- ▲ [Configure DNS Sinkholing](#)

DNS Sinkhole Workflow

The following illustration shows an example of how to identify client hosts that are attempting to communicate with known malicious domains:

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

As described in the [DNS Sinkhole Workflow](#), when a client host attempts to access a malicious domain, the DNS sinkhole option will forge the destination IP address using an IP address that the administrator defines. The following sections describe the procedure required to use this feature, how to configure the DNS sinkhole, and how to verify that the feature is functioning properly:

- ▲ [DNS Sinkhole Configuration Example](#)
- ▲ [Verify the Sinkhole Action and Reporting](#)

DNS Sinkhole Configuration Example

The steps that follow describes how to configure the DNS Sinkhole option in an Anti-Spyware profile and then attach the profile to a security rule.

Configure DNS sinkhole	
<p>Step 1 Obtain an IPv4 and IPv6 address from your network administrator that will be used as the sinkhole IP address. The addresses must be in a different zone than the client hosts, so when the infected host attempts to start a session with the sinkhole IP, 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>Important: This sinkhole addresses must be reserved for this purpose and does not have to be assigned to a physical host. A honey-pot server could also be used as a physical host to further analyze the malicious traffic.</p>	<p>In this example, the sinkhole IPv4 address is <i>10.15.0.20</i> and the IPv6 address is <i>fd97:3dec:4d27:e37c:5:5:5:5</i>.</p>

Configure DNS sinkhole (Continued)	
<p>Step 2 Configure the sinkhole zone if you do not have another zone that can be used.</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>It is recommended that you use a dedicated zone, because the infected host will be sending traffic to this zone.</p>	<ol style="list-style-type: none"> 1. Select Network > Interfaces and choose an available Layer 3 interface to be assigned to the sinkhole zone. For this example, use ethernet1/3. 2. In the Interface Type drop-down, select Layer3. 3. In the Config tab Virtual Router drop-down, select the virtual router that is used for your firewall. For this example, use the default virtual router. 4. To add an IPv4 address, select the IPv4 tab and select Static and then click Add. For this example, enter the IP address 10.15.0.20. 5. Select the IPv6 tab and click Static and then click Add. For this example, use the IP address fd97:3dec:4d27:e37c::/64. 6. Click OK to save. 7. To add a zone for the sinkhole, select Network > Zones and click Add. 8. Enter a name in the Name field. For this example, use the name Sinkhole. 9. In the Type drop-down menu select Layer3. 10. In the Interfaces section, click Add and add an interface. For this example, add ethernet1/3. 11. Click OK.
<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. For this example, clone the strict profile and rename it to <i>strict-dns-sinkhole</i>. 3. (Optional) Enter a description for the profile and select the DNS Signatures tab. 4. In the Action on DNS queries drop-down, select sinkhole. 5. In the Sinkhole IPv4 field enter a sinkhole IP address. For this example, IPv4 is 10.15.0.20 and IPv6 is fd97:3dec:4d27:e37c:5:5:5:5. The default value is the loopback address. 6. (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 and set between 1-50 packets. You can then use the packet captures for further analysis. 7. Click OK to save the profile. <p> The default sinkhole IP address is the loopback address, which will resolve domains to the local host. When the loopback is selected, communication from the infected client to the malware system is cut off/sinkholed.</p>

Configure DNS sinkhole (Continued)

Step 4 Ensure that the sinkhole zone and the zone where the client hosts reside can communicate and then attach the new Anti-Spyware profile to a security policy. This policy should be the policy that allows traffic for the client hosts in the Trust zone to the Untrust zone. Access from Trust to the new Sinkhole zone will also be configured to allow the client hosts to send queries to the sinkhole zone.

1. Select the **Policies** tab and then click **Security**.
2. Select an existing rule that allows traffic from the client host zone to the untrust zone. For this example, Rule1 was used.
3. Add the Sinkhole zone to the rule, so client hosts can reach the zone. For this example, add the Sinkhole zone in the **Destination** tab.
4. Select the **Actions** tab and select the **Log at Session Start** check box to enable it. 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** drop-down and select the Anti-Spyware profile that you configured previously. For this example, use strict-dns-sinkhole.
6. Click **OK** to save the security rule and then **Commit**.

To verify the configuration, see [Verify the Sinkhole Action and Reporting](#).

Verify the Sinkhole Action and Reporting

This section will describe the steps that can be performed to verify the DNS sinkhole feature. The values used are from the [DNS Sinkhole Configuration Example](#). You can perform similar steps for your own configuration.

DNS Sinkhole Verification and Reporting	
<p>Step 1 Verify that traffic is logged properly for traffic going from the client host in the Trust zone to the new Sinkhole zone. 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 the client host, open a command prompt and run the following command: <pre>C:\>ping 10.15.0.20</pre> The following example output shows the ping request and the result, which is <code>Request timed out</code> because there is no physical host assigned to the sinkhole IP address. <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 is traversing the firewall zones. Tip: 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. The following screenshot shows the log with the filter applied.
<p>Step 2 Now that the zones are configured properly, test that the DNS Signatures will perform the sinkhole action when a malware domain is accessed from the client host. 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.</p> <p> In this example, the URL track.bidtrk.com will be used for testing. This is a domain that is listed in the DNS Signatures database and is identified as being malicious, but it is possible that your Antivirus signature DB does not have this domain. To find a valid malicious domain for testing, see the information that follows this step.</p>	<ol style="list-style-type: none"> From the client host, open a command prompt. Perform an NSLOOKUP on the URL, <code>track.bidtrk.com</code>. For example: <pre>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</pre> In the above output, note that the NSLOOKUP to the malicious domain has been forged using the sinkhole IP addresses that were configured. Because the domain matched a malicious DNS signature, the sinkhole action was performed. View the threat log to see if the correct action was taken on the NSLOOKUP request. Select Monitor > Logs > Threat. Perform a ping to <code>track.bidtrk.com</code>, which will generate network traffic to the sinkhole address. This traffic will be used later in our example to generate a report to find infected hosts.

DNS Sinkhole Verification and Reporting (Continued)

How to find a valid malicious domain for testing?

This information will ensure that you are using a valid malicious domain, based on the current version of the antivirus signature database that is installed on your system. The DNS Signatures used to identify malicious domain is only part of the full antivirus signature database, which contains hundreds of thousands of signatures.

Perform the following steps to find a malicious domain for testing:

1. Navigate to **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.
2. In the second column of the release note, locate a line item with a domain extension (com, edu, net, and so on).
3. 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".
4. To test, from a command prompt, run `nslookup tbsbana.net`.
5. A sinkhole action should occur and the domain should resolve to the defined sinkhole address because the domain is verified in the antivirus DB that is installed on the firewall.

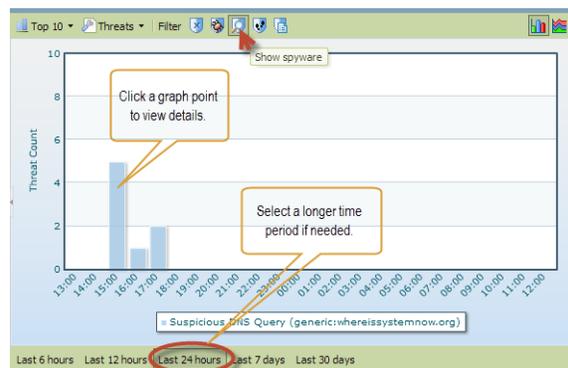
The following shows the content in the release note for this line item:

```
conficker:tbsbanal variants: net
```

Step 3 To view reports, you can use App Scope to view infected client hosts, or create custom reports.

1. To view from App Scope, 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. In this example, select **Last 24 hours**.

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 on the firewall to see more details about the event.



DNS Sinkhole Verification and Reporting (Continued)

Step 4 Configure a custom report that will identify all client hosts that have sent traffic to the sinkhole IP address, which is 10.15.0.20 in this example.



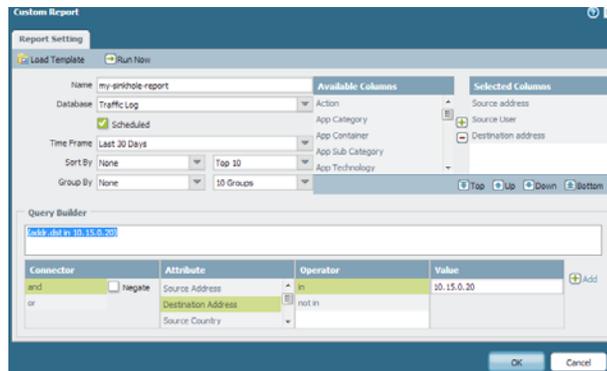
There are several ways to be alerted on these events, such as SNMP traps, sending to a Syslog server and/or Panorama.

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, for example *my-sinkhole-report*.
3. Define the custom report. This example uses the following report definitions:

- **Database**—Choose the detailed threat log, which is displayed as **Traffic Log**.
- **Scheduled**—Enable **Scheduled** and the report will run every night. To view scheduled reports that have run, select **Monitor > Reports**.
- **Time Frame**—30 days
- **Selected Columns**—Source address, Source User, Destination address. The critical fields are **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 the action sinkhole. Either enter the following 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.



- 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. These hosts should be tracked down and checked 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

Extended Packet Capture

A new extended packet capture option has been added to Anti-Spyware and Vulnerability Protection profiles for rules and exceptions defined in the profile. Previously, when selecting packet capture, only the first trigger packet would be captured when a threat was detected in traffic matching the profile. With the extended-capture option enabled, the firewall can capture from 1-50 packets, which provides more context when analyzing the packet captures in the logs.

Packet captures are typically used to further analyze events after receiving an alert or to test custom vulnerability signatures. After capturing the session, you then export it off of the device and then use an application such as Wireshark to replay the capture. Do not leave packet capture enabled after the needed captures have been generated because system performance will be impacted. Also note that packet captures will not occur if the block action is configured because the session is ended before there is anything to capture. A commit error will occur if packet capture is enabled on a rule with the block action. For more information on debugs and packet capture, refer to [Packet Based Troubleshooting](#).

Use the following procedure to enable this feature and to export the packet capture:

Configure Extended Packet Capture	
<p>Step 1 Define the number of packets that will be captured when the extended capture option is selected in an Anti-Spyware or Vulnerability Protection profile.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Content-ID and edit Threat Detection Settings. 2. Modify the Extended Packet Capture Length value. The range is from 1-50 packets and the default is 5. 3. Click OK to save the change.
<p>Step 2 Enable extended packet capture in the Anti-Spyware profile. In this example, the Anti-Spyware DNS Signature will be updated to perform an extended packet capture when a DNS lookup to a malicious site is performed.</p> <p>The extended packet capture can also be enabled in Anti-Spyware and Vulnerability Protection profiles for rules and exceptions.</p>	<ol style="list-style-type: none"> 1. Select Objects > Security Profiles > Anti-Spyware. 2. Select the default profile check box and then click the Clone button at the bottom of the page to create a new profile. Select the new profile and rename it if desired. 3. Select the DNS Signatures tab and select extended-capture from the Packet Capture drop-down. 4. Click OK to save the changes and then Commit the configuration.

Configure Extended Packet Capture (Continued)

Step 3 View the extended packet capture.

Now that extended packet capture is enabled, the next time a DNS Signature match occurs, the logs will have the option to view or download the packet capture.

An NSLOOKUP to a malicious domain was performed on a client host behind the firewall, which triggers a DNS event and the log will now have a packet capture as part of the log.

1. Select **Monitor > Logs > Threat**.
2. Locate the log for this event, which will have a green down arrow icon displayed to the right of the **Detailed Log View** icon. In this example, the client host at 192.168.2.10 performed the NSLOOKUP event that trigged the log and the log type is **spyware**.



3. Click the **Packet Capture** icon to view or export the packet capture.
 -  You can also click the **Detailed Log View** to view more details in the log and then click the **Packet Capture** icon in the **Related Logs** section to view and export the log.
4. To save the packet capture to be replayed in a separate application, click the **Export** button and save the file to your computer.

The capture can now be analyzed or replayed using a third-party application.

Passive DNS

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**.

URL Filtering Translation Site Filtering Enhancement

Translation filtering enhances the [URL filtering](#) engine such that URL filtering policies will also be applied to any URLs that are entered into translation sites such as Google Translate. This will ensure that website translation tools are not used to bypass URL filtering policies.

For example, if you go to `http://translate.google.com` and enter in a URL in the Translate field, such as `http://amazon.com`, the Amazon website would appear with the Google translate URL prepended as follows:

```
http://translate.google.com/translate?sl=fr&tl=en&js=n&prev=_t&hl=en&ie=UTF-8&u=amazon.com&act=url.
```

Without translation filtering, the user would still be able to get to Amazon, even if the corresponding URL category *Shopping* was set to block because the URL filtering engine would have interpreted the category as *Translation*.

URL Filtering Search Engine Cached Site Enhancement

Search engine caching enhances the [URL filtering](#) engine such that URL filtering policies will also be applied when end-users attempt to view the Google and Internet Archive cached copies of websites. There are no configuration changes required for these features.

When performing a Google search, most search results will show a green arrow to the right of the URL to access a cached version of the site as follows:



You can also put the prefix `cache :` before a URL in the Google search field to bring up the cached page. When the cached site comes up, note that the URL is prefixed by the Google cache URL as follows: `http://webcache.googleusercontent.com/`. If a URL filtering policy for *Computer and Internet Info* was set to block, which is the category for the Palo Alto Networks domain, the action would not occur because the URL filtering engine would only look at the `webcache.googleusercontent.com` URL, which is in the *Search Engine* category. With this new feature, the entire URL is analyzed, so in this example, `paloaltonetworks.com` would be blocked.

URL Filtering 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 now 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. With this feature enabled, when an end user attempts to perform a search without first enabling the strictest safe search settings, the firewall blocks the search query results and displays the URL Filtering Safe Search Block Page. By default, this page will provide a URL to the search provider settings for configuring safe search.

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.

The following procedure describes how to enable safe search and how to set strict safe search for each search provider.



As of Content Release version 475 or later, transparent safe search enforcement is also supported. Transparent safe search is a method for enforcing safe search without blocking end users' search results. Instead, when an end user attempts to perform a search without first enabling the strict safe search settings, the firewall redirects the search query to a URL that includes the safe search parameters. You [enable this functionality](#) by importing a new URL Filtering Safe Search Block Page containing the Javascript for rewriting the search URL to include the strict safe search parameters. In this configuration, users will not see the block page, but will instead be automatically redirected to a search query that enforces the strictest safe search options. This safe search enforcement method is only supported for Google, Yahoo, and Bing searches.

Enable Safe Search and Update the Browser Setting

<p>Step 1 Enable the Safe Search option in a URL Filtering profile.</p> <p>This example shows how to configure a new URL Filtering profile. If you have an existing profile, edit it and enable Safe Search.</p>	<ol style="list-style-type: none"> 1. Select Objects > Security Profiles > URL Filtering. 2. Select to modify an existing profile, or clone the default profile to create a new profile. 3. Select the Safe Search Enforcement check box to enable it. 4. (Optional) Modify the Action on any Category you would like to change and add any Block List or Allow List sites. For example, you may want to change alcohol-and-tobacco to block instead of allow or you may want to allow a specific site for a category that is set to block. 5. Click OK to save the new profile.
---	---

Enable Safe Search and Update the Browser Setting (Continued)	
<p>Step 2 Add the URL Filtering profile to the security policy that allows traffic from the client hosts to the Internet. This will activate the policy, so the URL Filtering profile and Safe Search Enforcement will be applied to users.</p> <p> Before activating the Safe Search Enforcement feature, it is recommended that you communicate the policy to your users, so they know what to expect.</p>	<ol style="list-style-type: none"> 1. Select Policies > Security and select a rule to which you want to apply the safe search-enabled URL Filtering profile. 2. On the Actions tab, select the profile you created/modified from the URL Filtering drop-down. 3. Click OK to save and then Commit, or go to the next step to customize the Safe Search Enforcement response page.
<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.* www.google.* search.yahoo.*</pre> d. Click OK to save the custom URL category object. 2. Follow the steps to configure SSL Forward Proxy decryption. 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 4 (Optional) Modify the block page for Safe Search Enforcement.</p> <p>The following lists the links that are used in the block page to directly access search preferences for some of the supported search engine providers:</p> <p>Bing http://www.bing.com/account/general</p> <p>Google https://www.google.com/preferences</p> <p>Yahoo http://search.yahoo.com/preferences/preferences</p>	<ol style="list-style-type: none"> 1. Select Device > Response Pages and select the URL Filtering Safe Search Block Page. 2. Select the Predefined check box and then click the Export button and the text file will be saved to the local system with the filename <code>safe-search-block-page.txt</code>. 3. Modify the file using an HTML or text editor and update as needed. 4. After updating, import the file by clicking the Import button in the URL Filtering Safe Search Block Page window. 5. Browse to the updated response page file, select it, and then click OK and close. 6. Commit the configuration and the new version of the response page will be applied.

Enable Safe Search and Update the Browser Setting (Continued)

Step 5 Verify that Safe Search Enforcement is functioning properly. For this example, use bing.com. Google and Yahoo are very similar, but the wording for the strict setting may vary.



If you are performing a search on Yahoo Japan (yahoo.co.jp) while logged into your Yahoo account, the lock option for the search setting must also be enabled.

1. Open bing.com from a browser on a computer that is behind the firewall.
2. Check the Safe Search setting by clicking the **Preferences** icon on the upper right of the Bing window.



3. The SafeSearch section shows three check boxes: Strict, Moderate, and Off. For this test, select either Off or Moderate.
4. Go back the Bing home page and perform a search. A block page should appear instead of the search results. The block page will look similar to the following:

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.

If the block page appears, the policy and Safe Search feature are working properly.

5. Click the provided link to change the SafeSearch setting. In this example, the link is <http://www.bing.com/account/general>.
6. In the **SafeSearch** section, select the **Strict** option and then click **Save** at the bottom of the page.
7. Perform a search again from Bing and the results should appear because strict safe search is now enabled in the browser.



The link to the SafeSearch settings page will be customized based on the search provider being used.

Step 6 (Optional) To prevent users from bypassing this feature and accessing other search engines other than Bing, Google, Yahoo, Yandex, and YouTube, modify the URL Filtering profile to block the search engine category and to only allow Bing, Google, and Yahoo.

1. Select **Objects > Security Profiles > URL Filtering** and modify the profile that has Safe Search enabled.
2. In the Category section, locate **search-engines** and change the action to block.
3. In the Allow List section, enter the following:
www.bing.com
www.google.com
www.yahoo.com
www.yandex.com
www.youtube.com
4. Click **OK** and then **Commit**.

WildFire Report Incorrect Verdict Option

The [WildFire reports](#) now have a link to submit a sample to the Palo Alto Networks threat team if you feel the verdict is a false positive or false negative. The threat team will perform further analysis on the sample to determine if it should be re-classified. If a malware sample is determined to be safe, the signature for the file will be disabled in an upcoming antivirus signature update. If a benign file is determined to be malicious, a new signature will be generated. After the investigation is complete, an email will be sent to the submitter (if an email address is provided) on the status of the investigation.

Report an Incorrect Verdict	
<p>Step 1 From either the WildFire portal or the Firewall logs, locate the sample that you believe has an incorrect verdict.</p>	<ol style="list-style-type: none"> 1. Select Monitor > Logs > WildFire Submissions. 2. Click the Detailed Log View icon in the first column of the log entry. 3. Select the WildFire Analysis Report tab and the report will appear. 4. Select the Report Incorrect Verdict link at the bottom of the report window. 5. (Optional) Enter your email address and any comments related to why you are suggesting a re-analysis of the sample. If you enter an email address, status of the analysis results will be emailed back to you.
<p>Step 2 The portal can also be used to view and report an incorrect verdict.</p>	<ol style="list-style-type: none"> 1. Go to https://wildfire.paloaltonetworks.com and click the Reports button. 2. Locate the file and click the detailed report icon. You can search by filename or SHA value if needed. 3. At the bottom of the page in the Report to Palo Alto Networks section, click the send the sample link. 4. (Optional) Enter your email address and any comments related to why you are suggesting a re-analysis of the sample. If you enter an email address, status of the analysis results will be emailed back to you.

WildFire Enhanced File Type and Operating System Support

As part of the [WildFire subscription](#), the following advanced file types are now supported in addition to the existing PE file type: Microsoft Office .doc/docx, .xls/xlsx, and .ppt/pptx; Portable Document Format (PDF); Java Applet (jar and class); and Android Application Package (APK). With a WildFire subscription, all listed file types can be submitted to WildFire from a PAN-OS firewall using a file blocking profile in a security policy, the WildFire API, or by manual upload to the [WildFire portal](#). If you do not have a WildFire subscription, the firewall only forwards PE files. You can, however, manually upload any of the other supported file types to the WildFire portal.



The WF-500 WildFire appliance does not support Android APK file analysis.

WildFire Analysis Report Enhancement

The WildFire analysis report has several new enhancements, including the ability to export the full WildFire report to a PDF; download the sample that was analyzed; and view the analysis results for each virtual environment in which the file was analyzed. Each sandbox environment has its own configuration of applications and software used in the file analysis, such as different versions of Adobe Reader, Flash, and MS Office.

The following table describes the report fields:

Report Heading	Description
Download PDF (New)	<ul style="list-style-type: none"> This button is located in the upper right corner of each report. Click the button to download a PDF version of the analysis report.
File Information	<ul style="list-style-type: none"> File Type (New) —PE, PDF, APK, JAR/Class, or MS Office (doc, xls, ppt). File Signer—The entity that signed the file for authenticity purposes. SHA-256—Displays the SHA information for the file. The SHA information is much like a fingerprint that uniquely identifies a file to ensure that the file has not been modified in any way. MD5—Displays the MD5 information for the file. The MD5 information is much like a fingerprint that uniquely identifies a file to ensure that the file has not been modified in any way. File Size—The size (in bytes) of the file that was analyzed. First Seen Timestamp (New) —If the WildFire system has analyzed the file previously, this is the date/time that it was first seen. Verdict—Displays the analysis verdict: <ul style="list-style-type: none"> Benign—The file is safe and does not exhibit malicious behavior. Malware—WildFire identified the file as malware and a signature will be generated to protect against future exposure. Sample File (New) —Click the Download File link to download the sample file to your local system. Virus Coverage—Click this link to see if the file has been previously identified. This will bring up the https://www.virustotal.com/en/ website, which contains information about various antivirus vendors and will show whether or not the vendors have coverage for the infected file. If the file has never been seen by any of the listed vendors, file not found will be displayed.

Report Heading	Description
Session Information	<p>Options used to customize which session information to include in the WildFire reports for files forwarded by a Palo Alto Networks firewall. The settings for these options are defined on the firewall that sends the sample file to WildFire and is configured in Device > Setup > WildFire tab in the Session Information Settings section.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> • Source IP • Source Port • Destination IP • Destination Port • Virtual System (If multi-vsyst is configured on the firewall) • Application • User (If User-ID is configured on the firewall) • URL • Filename
Dynamic Analysis (New)	<p>If a file is low risk and WildFire can easily determine that it is safe, only a static analysis is performed, instead of a dynamic analysis.</p> <p>When a dynamic analysis is performed, this section contains tabs for each virtual environment that the sample was run in when analyzing files in the WildFire cloud. For example, Virtual Machine 1 tab may have Windows XP, Adobe Reader 9.3.3, and Office 2003 and Virtual Machine 2 may have similar attributes, but with Office 2007. When a file goes through a full dynamic analysis, it is run in each virtual machine and the results of each environment can be viewed by clicking any of the Virtual Machine tabs.</p> <p> On the WF-500 WildFire appliance, one virtual machine will be used and is selected by the administrator based on the virtual environment attributes that best matches the local environment. For example, if most users have Windows 7, that virtual machine would be selected.</p>

Report Heading	Description
Behavior Summary	<p>Each Virtual Machine tab summarizes the behavior of the sample file in the specific environment. Examples include whether the sample created or modified files, started a process, spawned new processes, modified the registry, or installed browser helper objects.</p> <p>The following describes the various behaviors that are analyzed:</p> <ul style="list-style-type: none"> • Network Activity—Shows network activity performed by the sample, such as accessing other hosts on the network, DNS queries, and phone-home activity. (New) A link is provided to download the packet capture. • Host Activity—Lists any registry keys that were set, modified, or deleted. • Process Activity—Lists files that started a parent process, the process name, and the action the process performed. • File—Lists files that started a child processes, the process name, and the action the process performed. • Mutex (New) —If the sample file generates other program threads, the mutex name and parent process will be logged in this field. • Activity Timeline (New) —Provides a play-by-play list of all recorded activity of the sample. This will help in understanding the sequence of events that occurred during the analysis. <p> The time line activity is available in the PDF export of the report.</p>
Report Incorrect Verdict (New)	<p>Click this link to submit the sample file to the Palo Alto Networks threat team if you feel the verdict is a false positive or false negative. The threat team will perform further analysis on the sample to determine if it should be reclassified. If a malware sample is determined to be safe, the signature for the file will be disabled in an upcoming antivirus signature update or if a benign file is determined to be malicious, a new signature will be generated. After the investigation is complete, an email will be sent to the submitter (if an email address is provided) on the status of the investigation.</p>

WildFire Submissions Logs Available Without a Subscription

Previous to this release, a WildFire subscription was required to receive log results on the firewall and the reports were only available from the portal. With PAN-OS 6.0, when a firewall is configured with a file blocking profile and security policy to forward files to WildFire for analysis, the log results and access to the detailed analysis can be viewed directly from the [WildFire Submissions log](#) without a subscription.

A subscription is still required to forward files to a WF-500 WildFire appliance and/or to forward the advanced file types that are now supported in PAN-OS 6.0 to either the WildFire cloud or to a WildFire appliance.

WildFire Submissions Log Forwarding

Previously, if you wanted to forward **WildFire logs**, you had to enable forwarding of threat logs and forward medium severity logs (which included WildFire logs with a malicious verdict) and/or informational severity logs (which included WildFire logs with a benign verdict). You can now configure the firewall to automatically forward WildFire Submissions logs independently of the threat log forwarding configuration.

The following procedure describes the required steps to change the log forwarding options:

WildFire Log Forwarding

Step 1 Enable the forwarding of WildFire logs.

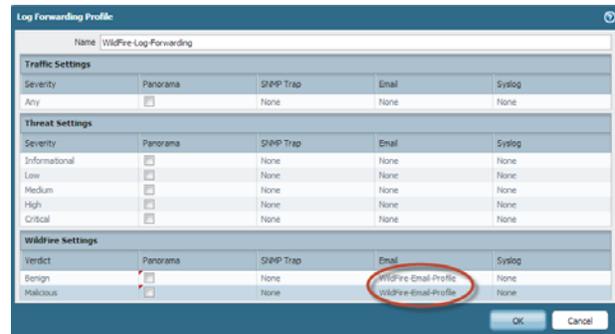


If you have an existing profile that forwards informational and medium threats, you may want to leave that setting, so other threats are still forwarded

1. Select **Objects > Log Forwarding**.
2. Click **Add** to add a new profile or select an existing profile. For this example, the profile was named WildFire-Log-Forwarding.
3. Modify the profile and in the **WildFire Settings** section, select the desired forwarding destination. You can choose Panorama, SNMP Trap, Email and/or Syslog.

If you do not have the destination profiles configured, you can configure them in **Device > Server Profiles**.

The following screenshot shows that Benign and Malicious files will be forwarded to an email profile.



4. Click **OK** and **Commit**.

WildFire Log Forwarding (Continued)

Step 2 Verify the configuration.

1. Attempt to download a file that will match the file blocking profile to forward a sample to WildFire. For example, if PE (executable) files are set to forward, an .exe file needs to be downloaded.

 To avoid downloading real malware, you can configure the reporting of benign files to make it easier to test. To enable this option, run the following command:

```
admin@host# set deviceconfig setting wildfire report-benign-file
```

After committing, the firewall logs Data Filtering logs and WildFire Submissions logs for both benign and malicious files.
2. After downloading the file, check the **Data Filtering (Monitor > Logs)** log and find the file that was downloaded. The **Action** column for the log entry will show **forward**, which indicates that the file was sent to WildFire. If WildFire successfully received the file, a second log will appear with the action **wildfire-upload-success**.
3. After approximately five minutes, the analyses result should appear in the **WildFire Submissions** log and triggers log forwarding. In this case, an email alert should be sent to the destination address defined in the email profile. If Panorama, SNMP Trap, or Syslog was selected as the destination, check those systems for the log.



GlobalProtect Features

The following sections describe the new GlobalProtect features and provide instructions for setting them up:

- ▲ [GlobalProtect Agent Deployment Customization](#)
- ▲ [GlobalProtect Agent Update Control](#)
- ▲ [Transparent One-Time Password \(OTP\) Support](#)
- ▲ [Client Certificate Authentication Enforcement](#)
- ▲ [HIP Profile Support for Client DLP Products](#)

GlobalProtect Agent Deployment Customization

There are now three ways to deploy custom agent configuration settings to your end user systems:

- ▲ [Push Agent Configuration Settings from the Portal Client Configuration](#)
- ▲ [Configure Agent Settings in the Windows Registry or Mac Global plist](#)
- ▲ [Deploy Agent Settings Automatically from the Windows Installer \(MSIEXEC\)](#)

Settings defined in the [GlobalProtect portal client configuration](#) take precedence over settings defined in the Windows Registry or the Mac plist. One setting—`can-prompt-user-credential`—is not available in the portal client configuration and must be set through the [Windows Registry](#) (applicable to Windows clients only). This setting is used in conjunction with single sign-on and indicates whether or not to prompt the user for credentials if SSO fails.

For a list of all configurable agent settings and the option name/syntax to use for each deployment method, see [Table: Customizable Agent Settings](#).

Push Agent Configuration Settings from the Portal Client Configuration

The portal client configuration in the web interface includes the following new settings:

- **Show GlobalProtect icon**—Disabling this option makes GlobalProtect invisible to the end user by removing the icon from the system tray and preventing the user from having any interaction with the GlobalProtect agent user interface.
- **Allow user to change portal address**—Disabling this option prevents users from manually changing the portal address on Settings panel in the GlobalProtect agent or app.
- **Allow user to continue if portal server certificate is invalid**—Disabling this option prevents users from continuing if there is a warning screen that might indicate a man-in-the-middle (MITM) attack.

The following workflow provides instructions on configuring these new options from the web interface.

Customize the Portal Client Configuration	
<p>Step 1 Go to the Agent tab in the client configuration you want to customize.</p>	<ol style="list-style-type: none"> 1. Select Network > GlobalProtect > Portals and select the portal configuration for which you want to add a client configuration (or click Add to add a new configuration). 2. Select the Client Configuration tab and select the client configuration you want to modify (or click Add to add a new configuration). 3. Select the Agent tab.

Customize the Portal Client Configuration

Step 2 Define what the end users with this configuration can do from the agent.



All of the settings on the **Agent** tab can also be configured in the end client via group policy by adding settings to the Windows Registry/Mac plist. On Windows systems, you can also set them using the Msiexec utility from the command line during the agent installation. However, settings defined in the web interface or the CLI take precedence over Registry/plist settings.



Another **NEW OPTION**—`can-prompt-user-credential`—is available only through the Windows command line (Msiexec) or Windows Registry (`HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect\Settings`). This setting controls whether or not to allow the agent to prompt the end user for credentials if Windows SSO fails. By default, the agent will prompt for credentials (default setting=`yes`).

By default, the agent functionality is fully enabled (meaning all check boxes are selected). To remove functionality, clear the corresponding check box for any or all of the following options:

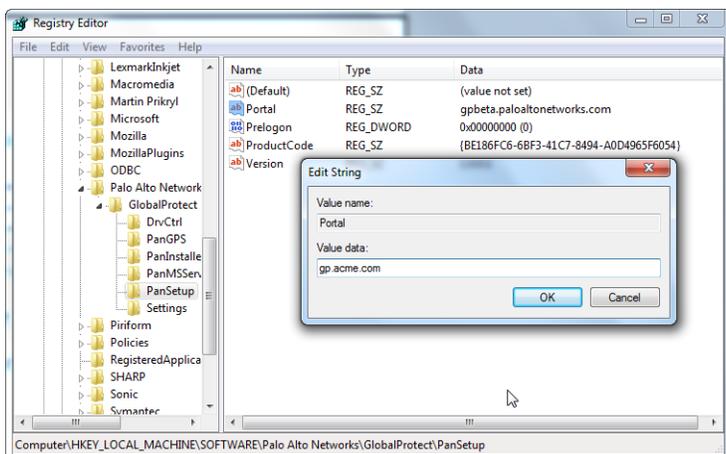
- If you want users to only be able to see basic status information from within the application, clear the **Enable advanced view** check box. By default, the advanced view is enabled, which allows end users to see detailed statistical, host, and troubleshooting information and perform tasks such as changing their passwords.
- **NEW OPTION** If you want hide the GlobalProtect agent on the end user systems, clear the **Show GlobalProtect icon** check box. When the icon is hidden, users cannot perform other tasks such as changing passwords, rediscovering the network, resubmitting host information, viewing troubleshooting information, or performing an on-demand connection. However, HIP notification messages, login prompts, and certificate dialogs will still display as necessary for interacting with the end user.
- **NEW OPTION** Clear the **Allow user to change portal address** check box to disable the **Portal** field on the **Settings** panel in the GlobalProtect agent/app.
- If you do not want users to be able to save their passwords on the agent (that is, you want to force them to provide the password—either transparently through the browser or by manually entering one—each time they connect), clear the **Allow user to save password** check box.
- To prevent users from performing a network rediscovery, clear the **Enable Rediscover Network option** check box.
- To prevent users from manually resubmitting HIP data to the gateway, clear the **Enable Resubmit Host Profile option** check box. This option is enabled by default, and is useful in cases where HIP-based security policy prevents users from accessing resources because it allows the user to fix the compliance issue on the computer and then resubmit the HIP.
- **NEW OPTION** To prevent users from continuing with a connection if the portal certificate presented is invalid (which may indicate a MITM attack), clear the **Allow user to continue if portal certificate is invalid** check box. For the strongest protection against MITM attacks, use a certificate from a well-known CA on the portal and pre-deploy this setting through the Windows Registry or Mac plist.

Customize the Portal Client Configuration	
<p>Step 3 Specify which users to deploy this configuration to. There are two ways to specify who will get the configuration: by user/group name and/or the operating system the agent is running on.</p>	<p>Select the User/User Group tab and then specify the user/user groups and/or operating systems to which this configuration should apply:</p> <ul style="list-style-type: none"> To restrict this configuration to a specific user or group, click Add in the User/User Group section of the window and then select the user or group you want to receive this configuration from the drop-down. Repeat this step for each user/group you want to add. NEW OPTION To deliver this configuration to agents or apps running on specific operating systems, click Add in the OS section of the window and then select the OS (Android, iOS, Mac, or Windows) to which this configuration applies.
<p>Step 4 Save the agent configuration settings.</p>	<ol style="list-style-type: none"> Click OK twice to close the dialogs. Commit your changes.

Configure Agent Settings in the Windows Registry or Mac Global plist

You can set the GlobalProtect agent customization settings in the [Windows registry](#) (HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect\Settings\) or the [Mac global plist](#) (/Library/Preferences/com.paloaltonetworks.GlobalProtect.settings.plist). This enables deployment of GlobalProtect agent settings to client systems prior to their first connection to the GlobalProtect portal. For a list of commands and values, see [Table: Customizable Agent Settings](#).

If you do not want the user to manually enter the portal address even for the first connection, you can also pre-deploy the portal address through the Windows Registry: (HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect\PanSetup with key Portal) or the Mac plist (/Library/Preferences/com.paloaltonetworks.GlobalProtect.settings.plist and configure key Portal under dictionary PanSetup):



Deploy Agent Settings Automatically from the Windows Installer (MSIEXEC)

For Windows clients, you can automatically deploy the settings in the Windows Installer ([Msiexec](#)). For example, you could use Msiexec to deploy the setting to prevent users from changing the portal address using the following command:

```
msiexec.exe /i GlobalProtect.msi CANCHANGEPORTAL="no"
```

For a list of Msiexec parameters, see [Table: Customizable Agent Settings](#).

Customizable Agent Settings

The following table summarizes the agent settings that you can customize. For a description of each option, refer to the online help:

Table: Customizable Agent Settings

Portal Client Configuration	Windows Registry/ Mac plist	Msiexec Parameter	Default
Enable advanced view	enable-advanced-view yes no	ENABLEADVANCEDVIEW="yes no"	yes
Show GlobalProtect icon	show-agent-icon yes no	SHOWAGENTICON="yes no"	yes
Allow users to change portal address	can-change-portal yes no	CANCHANGEPORTAL="yes no"	yes
Allow user to save password	can-save-password yes no	CANSAVEPASSWORD="yes no"	yes
Enable rediscover network option	rediscover-network yes no	REDISCOVERNETWORK="yes no"	yes
Enable Resubmit Host Profile option	resubmit-host-info yes no	RESUBMITHOSTINFO="yes no"	yes
Allow user to continue if portal server certificate is invalid	can-continue-if-portal-cert-invalid yes no	CANCONTINUEIFPORTALCERTINVALID="yes no"	yes
Use single sign-on	use-sso yes no	USESSO="yes no"	yes
Config Refresh Interval (hours)	refresh-config-interval <hours>	REFRESHCONFIGINTERVAL="<hours>"	24
Connect Method	connect-method on-demand pre-logon user-logon	CONNECTMETHOD="on-demand pre-logon user-logon"	user-logon
Windows only/not in portal	can-prompt-user-credential yes no	CANPROMPTUSERCREDENTIAL="yes no"	yes

GlobalProtect Agent Update Control

The client configurations delivered by the GlobalProtect portal now have two additional options for controlling when users can [upgrade the GlobalProtect agent](#):

- **disable**—Prevent end users from upgrading the agent. This is useful in environments where you want to test the new agent version on a subset of users before deploying it to all users.
- **manual**—Allows end users to initiate agent upgrades. In this case, the user would select the **Check Version** option in the agent to determine if there is a new agent version and then upgrade if desired.

Because this is a setting within the agent configuration, you can define different configurations for different groups of users or computers and deploy the agent updates appropriately for each group. If you want to control when users can upgrade, for example if you want to test a release on a small group of users before deploying it to your entire user base, you can customize the agent upgrade behavior on a per-configuration basis. In this case, you could create a configuration that applies to users in your IT group only to allow them to upgrade and test and disable upgrade in all other user/group configurations. Then, after you have thoroughly tested the new version, you could modify the agent configurations for the rest of your users to allow the upgrade. The following workflow shows how to define the agent upgrade settings and deploy them to specific groups.

Define Agent Update Settings	
<p>Step 1 Go to the Agent tab in the client configuration for which you want to customize agent updates.</p>	<ol style="list-style-type: none"> 1. Select Network > GlobalProtect > Portals and select the portal configuration. 2. Select the Client Configuration tab and select the client configuration you want to modify (or click Add to add a new configuration). 3. Select the Agent tab.
<p>Step 2 Specify how GlobalProtect agent upgrades will occur.</p>	<p>By default, the Agent Upgrade field is set to prompt the end user to upgrade. To modify this behavior, select one of the following options:</p> <ul style="list-style-type: none"> • If you want upgrades to occur automatically without interaction with the user, select transparent. • To prevent agent upgrades, select disable. • To allow end users to initiate agent upgrades, select manual. In this case, the user would select the Check Version option in the agent to determine if there is a new agent version and then upgrade if desired.

Define Agent Update Settings (Continued)

<p>Step 3 Specify which users to deploy this configuration to. There are two ways to specify who will get the configuration: by user/group name and/or the operating system the agent is running on.</p>	<p>Select the User/User Group tab and then specify the user/user groups and/or operating systems to which this configuration should apply:</p> <ul style="list-style-type: none"> To restrict this configuration to a specific user or group, click Add in the User/User Group section of the window and then select the user or group you want to receive this configuration from the drop-down. Repeat this step for each user/group you want to add. NEW OPTION To deliver this configuration to agents or apps running on specific operating systems, click Add in the OS section of the window and then select the OS (Android, iOS, Mac, or Windows) to which this configuration applies.
<p>Step 4 Save the configuration.</p>	<ol style="list-style-type: none"> Click OK to save the settings and close the Configs dialog. Click OK to save the settings and close the GlobalProtect Portal dialog. Commit your changes.

Transparent One-Time Password (OTP) Support

By default, the GlobalProtect agent forwards the user authentication credentials for the portal on to the gateway and, in the simplest case where the gateway and the portal use the same authentication profile and/or certificate profile, the agent will connect to the gateway transparently. However, if the portal and the gateway use different [authentication methods](#) or if they require different credentials (such as unique OTPs), this default behavior would cause delays in connecting to the gateway because the gateway would not prompt the user to authenticate until after it tried and failed to authenticate using the credentials the agent forwarded. The following sections describe the new configuration settings for enabling transparent OTP support and provide procedures for setting up transparent OTP authentication:

- ▲ [About the New Authentication Modifier Settings](#)
- ▲ [Set up Transparent OTP Authentication](#)

About the New Authentication Modifier Settings

To simplify the authentication process and make it more transparent for the end user, the portal now includes the following new settings for modifying the default authentication behavior on a per-client configuration basis:

- **Cookie authentication for config refresh**—Enables the portal to use an encrypted cookie to authenticate agents when refreshing a configuration that has already been cached (the user will always be required to authenticate for the initial configuration download and upon cookie expiration). This simplifies the authentication process for end users because they will no longer be required to log in to both the portal and the gateway in succession or enter multiple OTPs for authenticating to each.
- **Different password for external gateway**—Disables the forwarding of credentials to some or all gateways, allowing the gateway to immediately prompt for its own set of credentials. This option speeds up the authentication process when the portal and the gateway require different credentials (either different OTPs or different login credentials entirely). Or, you can choose to use a different password on manual gateways only. With this option, the agent will forward its portal credentials to automatic gateways but not to manual gateways, allowing you to have the same security on your portals and automatic gateways, while requiring a second factor OTP or a different password for access to those gateways that provide access to your most sensitive resources.

Set up Transparent OTP Authentication

On the firewall, the process for setting up access to a two-factor authentication service is similar to setting up any other type of authentication: create a server profile (usually to a RADIUS server), add the server profile to an authentication profile, and then reference that authentication profile in the configuration for the device that will be enforcing the authentication—the GlobalProtect portal and/or gateway in this case.

By default, the agent forwards the same authentication credentials it used to authenticate to the portal on to the gateway and the gateway attempts to use these credentials to authenticate the user. In the case of OTP authentication, this behavior will cause the authentication to initially fail on the gateway and, because of the delay this causes in prompting the user for a login, the user's OTP may expire. To prevent this, the portal allows for

modification of this behavior on a per-client configuration basis—either by allowing the portal to authenticate using an encrypted cookie or by disabling the forwarding of credentials. Both of these options solve this problem by enabling the gateway to immediately prompt for the appropriate credentials.

Enable OTP Support Using Authentication Modifiers	
<p>Step 1 Set up your RADIUS server to interact with the firewall.</p> <p>This procedure assumes that your RADIUS service is already configured for OTP- or token-based authentication and that necessary devices (such as hardware tokens) have been deployed to users.</p>	<p>For specific instructions, refer to your the documentation for your RADIUS server. In most cases, you will need to set up an authentication agent and a client configuration on the RADIUS server to enable communication between the firewall and the RADIUS server. You will also define the shared secret that will be used to encrypt sessions between the firewall and the RADIUS server.</p>
<p>Step 2 On the firewall that will act as your gateway and/or portal, create a RADIUS server profile.</p> <p> When creating the RADIUS server profile, always enter a Domain name because this value will be used as the default domain if users don't supply one upon login.</p>	<ol style="list-style-type: none"> 1. Select Device > Server Profiles > RADIUS, click Add and enter a Name for the profile. 2. Enter the RADIUS Domain name. 3. To add a RADIUS server entry, click Add in the Servers section and then enter the following information: <ul style="list-style-type: none"> • A descriptive name to identify this RADIUS Server • The IP Address of the RADIUS Server • The shared Secret used to encrypt sessions between the firewall and the RADIUS server • The Port number on which the RADIUS server will listen for authentication requests (default 1812) 4. Click OK to save the profile.
<p>Step 3 Create an authentication profile.</p>	<ol style="list-style-type: none"> 1. Select Device > Authentication Profile, click Add, and enter a Name for the profile. The authentication profile name cannot contain any spaces. 2. Select RADIUS from the Authentication drop-down. 3. Select the Server Profile you created for accessing your RADIUS server. 4. Click OK to save the authentication profile.
<p>Step 4 Assign the authentication profile to the GlobalProtect gateway(s) and/or portal.</p>	<ol style="list-style-type: none"> 1. Select Network > GlobalProtect > Gateways or Portals and select the configuration (or Add one). 2. On the General tab (on the gateway) or the Portal Configuration tab (on the portal), select the Authentication Profile you just created. 3. Enter an Authentication Message to guide users as to which authentication credentials. This is helpful to users, especially if you require different credentials on the portal and the gateway(s). 4. Click OK to save the configuration.

Enable OTP Support Using Authentication Modifiers (Continued)	
<p>Step 5 (Optional) Modify the default authentication behavior on the portal.</p>	<ol style="list-style-type: none"> 1. Select Network > GlobalProtect > Portals and select the configuration (or Add one). 2. Select the Client Configuration tab and then select or Add a client configuration. 3. On the General tab, select one of the following values from the Authentication Modifier field: <ul style="list-style-type: none"> • Cookie authentication for config refresh—Enables the portal to use an encrypted cookie to authenticate users so they don't have to enter multiple OTPs or credentials. • Different password for external gateway—Prevents the agent from forwarding the user credentials to the gateway to prevent OTP authentication failures. 4. Click OK twice to save the configuration.
<p>Step 6 Save the configuration.</p>	<p>Click Commit.</p>
<p>Step 7 Verify the configuration. This step assumes that your gateway and portal are already configured.</p>	<p>From a client system running the GlobalProtect agent, try to connect to a gateway or portal on which you enabled OTP authentication. You should see two prompts similar to the following:</p> <p>The first will prompt you for a PIN (either a user- or system-generated PIN):</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  </div> <p>The second will prompt you for your token or OTP:</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  </div>

Client Certificate Authentication Enforcement

The portal and gateway require the end-user authentication credentials before the GlobalProtect agent/app will be allowed access to GlobalProtect resources. [Client certificate authentication](#) is one authentication mechanism supported on GlobalProtect. It can be used alone or in conjunction with other authentication mechanisms, such as external authentication against an existing LDAP, Kerberos, or RADIUS service (including support for two-factor token-based authentication mechanisms such as one-time password (OTP) authentication).

With this release, enhancements have been made to how client certificate authentication is enforced in GlobalProtect in various scenarios:

- If a certificate profile is configured on the GlobalProtect portal, the client must present a certificate in order to connect. This means that certificates must be pre-deployed to the end clients before their initial portal connection.
- If the certificate profile specifies **Subject** in the **Username Field**, the certificate presented by the client must contain a common-name in order to connect. If the certificate profile specifies a **Subject-Alt** with an **Email** or **Principal Name** as the **Username Field**, the certificate presented by the client must contain the corresponding fields, which will be used as the username when the GlobalProtect agent authenticates to the portal or gateway.
- If you are using two-factor authentication, users must successfully authenticate both with a certificate profile and an authentication profile. And, if the certificate profile is configured with a **Username Field**, the username field of the portal or gateway login screen will automatically be populated with the value from the corresponding certificate field. For example, if the **Username Field** in the certificate profile is set to **Subject**, the common-name in the certificate will by default be used as the username when the user attempts to authenticate. If you do not want force users to authenticate with a username from the certificate, you must select **None** in the **Username Field** of the certificate profile.
- For agents configured with the **pre-logon** connect method, if the new **Cookie authentication for config refresh** setting is enabled, you no longer need to configure a certificate profile for pre-logon authentication; in this case the portal will use the cookie to authenticate the client prior to user logon. Note, however, that you must still configure a certificate profile on the gateway to enable establishment of the VPN tunnel. For more information on cookie authentication, see [About the New Authentication Modifier Settings](#).

The client certificate authentication enhancements do not require any changes to existing configuration or require any different steps to configure client certificate authentication; they provide behavioral changes only.

HIP Profile Support for Client DLP Products

The GlobalProtect agent by default now collects vendor-specific data about whether data loss prevention (DLP) software is installed and/or enabled on Windows clients. DLP software is used to prevent sensitive corporate information from leaving the corporate network or from being stored on a potentially insecure device. Because this information is now collected from the client systems, you can include DLP as matching criteria for the host information profiles (HIPs) you create, thereby enabling you to use DLP compliance as criteria for your security policies.

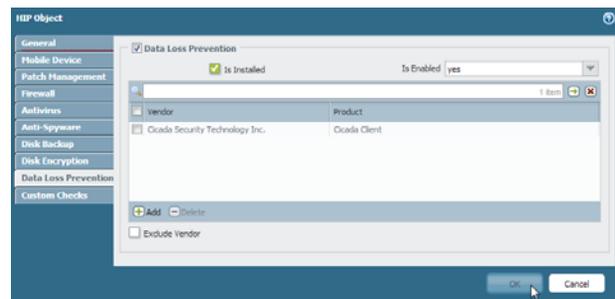
Because host information about DLP client software configuration is automatically collected by the GlobalProtect gateway when the agent connects and submits its host information, there is nothing you need to do to enable this feature. As with all other HIP data submitted by the agent, the gateway will match the host DLP information against any HIP objects and HIP profiles you have defined. If the gateway finds a match, it generates an entry in the HIP Match log. Additionally, if the gateway finds a HIP profile match in a policy rule, it enforces the corresponding security policy. The following example shows how to use the DLP state of the Windows client to enforce security policy:

Enforce Security Policy Based on the DLP state of the Client																	
<p>Step 1 Verify proper licensing for HIP checks.</p>	<p>To use the HIP feature, you must have purchased and installed a GlobalProtect Portal license on the firewall where your portal is configured and a GlobalProtect Gateway subscription license on each gateway that will perform HIP checks. To verify the status of your licenses on each portal and gateway, select Device > Licenses.</p> <div data-bbox="760 1010 1190 1266"><table border="1"><thead><tr><th colspan="2">GlobalProtect Portal</th></tr></thead><tbody><tr><td>Date Issued</td><td>March 23, 2012</td></tr><tr><td>Date Expires</td><td>Never</td></tr><tr><td>Description</td><td>GlobalProtect Portal License</td></tr></tbody></table> <table border="1"><thead><tr><th colspan="2">GlobalProtect Gateway</th></tr></thead><tbody><tr><td>Date Issued</td><td>March 19, 2012</td></tr><tr><td>Date Expires</td><td>March 19, 2015</td></tr><tr><td>Description</td><td>GlobalProtect Gateway License</td></tr></tbody></table></div>	GlobalProtect Portal		Date Issued	March 23, 2012	Date Expires	Never	Description	GlobalProtect Portal License	GlobalProtect Gateway		Date Issued	March 19, 2012	Date Expires	March 19, 2015	Description	GlobalProtect Gateway License
GlobalProtect Portal																	
Date Issued	March 23, 2012																
Date Expires	Never																
Description	GlobalProtect Portal License																
GlobalProtect Gateway																	
Date Issued	March 19, 2012																
Date Expires	March 19, 2015																
Description	GlobalProtect Gateway License																

Enforce Security Policy Based on the DLP state of the Client (Continued)

Step 2 Create the HIP objects to filter the raw host data collected by the agents.

1. On the gateway (or on Panorama if you plan to share the HIP objects among multiple gateways), select **Objects > GlobalProtect > HIP Objects** and click **Add**.
2. On the **General** tab, enter a **Name** for the object.
3. Select the tab that corresponds to the category of host information you are interested in matching against and select the check box to enable the object to match against the category. For example, to create an object that looks for information about DLP software, select the **Data Loss Prevention** tab and then select the **Data Loss Prevention** check box to enable the corresponding fields. Complete the fields to define the desired matching criteria. For example, the following screenshot shows how to create an object that will match if the Cicada Security Technology Inc. Cicada Client application is installed and enabled.



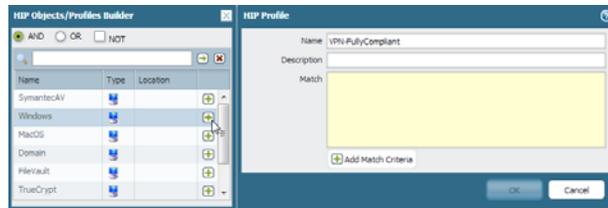
4. Click **OK** to save the HIP object.
5. **Commit** your changes.

Enforce Security Policy Based on the DLP state of the Client (Continued)

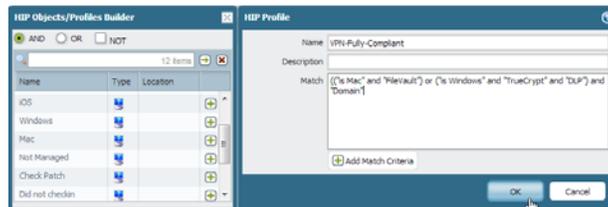
Step 3 Create the HIP profiles that you plan to use in your policies.

When you create your HIP profiles, you can combine the HIP objects you previously created (as well as other HIP profiles) using Boolean logic such that when a traffic flow is evaluated against the resulting HIP profile it will either match or not match. If there is a match, the corresponding policy rule will be enforced; if there is not a match, the flow will be evaluated against the next rule, as with any other policy matching criteria.

1. On the gateway (or on Panorama if you plan to share the HIP profiles among multiple gateways), select **Objects > GlobalProtect > HIP Profiles** and click **Add**.
2. Enter a descriptive **Name** for the profile and optionally a **Description**.
3. Click **Add Match Criteria** to open the HIP Objects/Profiles Builder.
4. Select the first HIP object or profile you want to use as match criteria and then click add **+** to move it over to the **Match** text box on the HIP Profile dialog. Keep in mind that if you want the HIP profile to evaluate the object as a match only when the criteria in the object is not true for a flow, select the **NOT** check box before adding the object.



5. Continue adding match criteria as appropriate for the profile you are building, making sure to select the appropriate Boolean operator radio button (**AND** or **OR**) between each addition (and, again, using the **NOT** check box when appropriate).
6. If you are creating a complex Boolean expression, you must manually add the parenthesis in the proper places in the **Match** text box to ensure that the HIP profile is evaluated using the logic you intend. For example, the following HIP profile will match traffic from a host that has either FileVault disk encryption (for Mac OS systems) or TrueCrypt disk encryption and has the Cicada Client DLP client installed (for Windows systems) and also belongs to the required Domain



7. When you are done adding match criteria, click **OK** to save the profile.
8. **Commit** your changes.

Enforce Security Policy Based on the DLP state of the Client (Continued)

Step 4 Enable User-ID on the source zones that contain the GlobalProtect users that will be sending requests that require HIP-based access controls. You must enable User-ID even if you don't plan on using the user identification feature or the firewall will not generate any HIP Match logs entries.

1. Select **Network > Zones**.
2. Click on 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**.

Name	Type	Interfaces / Virtual Systems	Zone Protection Profile	Log Setting	Enable User Identification
corp-vpn	layer3	ethernet1/2 tunnel.1			<input checked="" type="checkbox"/>

Step 5 Create the HIP-enabled security rules on your gateway(s).

As a best practice, you should create your security rules and test that they match the expected flows based on the source and destination criteria as expected before adding your HIP profiles. By doing this you will also be better able to determine the proper placement of the HIP-enabled rules within the policy.

Add the HIP profile to your security rules:

1. Select **Policies > Security** and select the rule to which you want to add a HIP profile.
2. On the **Source** tab, make sure the **Source Zone** is a zone for which you enabled User-ID in [Step 4](#).
3. On the **User** tab, click **Add** in the **HIP Profiles** section and select the HIP profile(s) you want to add to the rule (you can add up to 63 HIP profiles to a rule).
4. Click **OK** to save the rule.
5. **Commit** your changes.

Name	Tags	Zone	Source			HIP Profile	Destination	
			Address	User	Zone		Address	
1 Full Corporate Access		corp-vpn l3-trust	any	any	VPN-fully-compliant	any	any	



Management Features

The following sections describe the new Management features and provide instructions for setting them up:

- ▲ Enumeration of Rules Within a Rulebase
- ▲ Enhancements in Reports
- ▲ Support for Color Coded Tags
- ▲ CLI Find Command
- ▲ Support for Syslog Over TCP and SSL
- ▲ SNMP Resource Monitoring Extensions
- ▲ Enhancement in the Syslog Header
- ▲ Virtual Machine (VM) Monitoring Agent (See Virtualization Features)
- ▲ Dynamic Address Groups (See Dynamic Address Groups)

Enumeration of Rules Within a Rulebase

The **Policies** tab on the web interface includes a new column for displaying rule numbers. Each rule is automatically numbered and the ordering adjusts as rules are moved or reordered. When filtering rules to find rules that match the specified filter(s), each rule is listed with its number in the context of the complete set of rules in the rulebase and its place in the evaluation order.

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.

	Name	Zone	Destination	Application	Service	Class	Schedule
1	Video	any	any	google-video http-video youtube	any	1	none
2	HTTPS	any	any	web-browsing	any	2	none
3	FTP	any	any	ftp	any	4	none

- View the numbered list of rules on Panorama.

Select **Policies** and any rulebase under it. For example, **Policies > Security > Pre-rules**.

	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.

	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

Enhancements in Reports

The firewall and Panorama allow you to generate reports and view the log data in a tabular format. To enhance the reporting enhancements that are currently available, the following capabilities have been added:

- ▲ Create Group Activity Reports
- ▲ Disable Predefined Reports

Create Group Activity Reports

In addition to generating user activity reports, you can generate group activity reports for user groups on the firewall. Groups activity reports cannot be generated on Panorama because user to group mapping information is not available on Panorama.

To create groups activity reports, you must have set up the User-ID functionality on the firewall so that the user to group mapping information is available.

Generate Group Activity Reports

1. Select **Monitor > PDF Reports > User Activity Report**.
2. Select **Type: Group** and pick the group for which to generate the report.
3. (Optional) Select **Include Detailed Browsing** only if you want an expansive report with detailed URL activity information for each user. With this option selected, the report is voluminous.
4. Click **Run Now** to run the report immediately, or click **OK** to save the report. These reports cannot be scheduled.



5. Click **Commit** to save the changes.

Disable Predefined Reports

The firewall and Panorama include about 40 predefined reports that are automatically generated. If you do not use some or all of these predefined reports, you can disable selected reports and conserve system resources on the firewall and Panorama. 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 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 or **Panorama > Setup > Management** on Panorama.
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.

Pre-Defined Reports

Application Reports	Traffic Reports	Threat Reports	URL Filtering Reports
<input checked="" type="checkbox"/> Applications	<input checked="" type="checkbox"/> Security Rules	<input checked="" type="checkbox"/> Threats	<input checked="" type="checkbox"/> URL Categories
<input checked="" type="checkbox"/> Application Categories	<input checked="" type="checkbox"/> Sources	<input type="checkbox"/> Threat Trend	<input type="checkbox"/> URL Users
<input checked="" type="checkbox"/> Technology Categories	<input checked="" type="checkbox"/> Source Countries	<input checked="" type="checkbox"/> Attackers	<input checked="" type="checkbox"/> URL User Behavior
<input checked="" type="checkbox"/> HTTP Applications	<input checked="" type="checkbox"/> Destinations	<input checked="" type="checkbox"/> Attacker Countries	<input type="checkbox"/> Web Sites
<input checked="" type="checkbox"/> Denied Applications	<input checked="" type="checkbox"/> Destination Countries	<input checked="" type="checkbox"/> Victims	<input type="checkbox"/> Blocked Categories
<input checked="" type="checkbox"/> Risk Trend	<input checked="" type="checkbox"/> Connections	<input checked="" type="checkbox"/> Victim Countries	<input checked="" type="checkbox"/> Blocked Users
<input type="checkbox"/> Bandwidth Trend	<input checked="" type="checkbox"/> Source Zones	<input checked="" type="checkbox"/> Viruses	<input checked="" type="checkbox"/> Blocked User Behavior
	<input checked="" type="checkbox"/> Destination Zones	<input type="checkbox"/> Spyware	<input checked="" type="checkbox"/> Blocked Sites
	<input checked="" type="checkbox"/> Ingress Interfaces	<input checked="" type="checkbox"/> Vulnerabilities	
	<input type="checkbox"/> Egress Interfaces	<input checked="" type="checkbox"/> Spyware Infected Hosts	
	<input checked="" type="checkbox"/> Denied Sources	<input checked="" type="checkbox"/> Top Users	
	<input checked="" type="checkbox"/> Denied Destinations	<input checked="" type="checkbox"/> Wildfire File Digests	
	<input checked="" type="checkbox"/> Unknown TCP Sessions		
	<input checked="" type="checkbox"/> Unknown UDP Sessions		
	<input checked="" type="checkbox"/> Risky Users		

Note: Group Reports and PDF Reports will have no data if a contained pre-defined report is disabled

[Select All](#) [Deselect All](#)

Support for Color Coded Tags

You can now 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.

While the firewall supports both static tags and dynamic tags, dynamic tags are added using the XML API and scripts or by configuring the [Support for the VM-Series Firewall on the Citrix SDX Server](#) on the firewall. Dynamic tags are not displayed along with the static tags, and they are not part of the device configuration. 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).

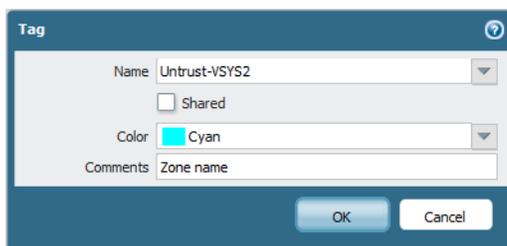
Tag Objects

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**.

CLI Find Command

The new **CLI find** command helps you find a command when you don't know where to start looking in the hierarchy. The command—which is available in all CLI modes—has two forms. You can either use the **find** command alone to display the entire command hierarchy in the current command mode. Or, you can use the **find** command with the **keyword** argument to locate all commands that have the specified keyword. For example, to find all configure mode commands with the **username** keyword you would enter the following:

```
admin@mgmt-ui-4060# find command keyword username
set deviceconfig system log-export-schedule <name> protocol ftp username <value>
set deviceconfig system log-export-schedule <name> protocol scp username <value>
set deviceconfig setting wildfire session-info-select exclude-username <yes|no>
set mgt-config password-complexity block-username-inclusion <yes|no>
set network interface ethernet <name> layer3 pppoe username <value>
set shared certificate-profile <name> username-field
```

Support for Syslog Over TCP and SSL

If you send log data to a syslog server for long-term storage, archival and centralized reporting, both the firewall and Panorama support using TCP and SSL (in addition to UDP) for reliable and secure transport of logs to an external syslog server. The five types of logs that can be forwarded to a syslog server are traffic, threat, HIP match, config, and system.

To enable client authentication for syslog over SSL, you can import or generate a certificate 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 Transport Mechanism for Syslog

Step 1 Create a syslog server profile for connecting to the syslog server.

1. Select **Device > Server Profiles > Syslog**.
2. On Panorama, select the template to which these settings must apply.
3. Click **Add** and then enter a **Name** for the profile.
4. (Optional) Select the virtual system to which this profile applies from the **Location** drop-down.
5. 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):
 - **Name**—Unique name for the server profile.
 - **Syslog Server**—IP address or fully qualified domain name (FQDN) of the Syslog server.
 - **Transport**—UDP, TCP, or SSL as the medium of transport; SSLv3 and TLSv1 are supported.
 - **Port**—The port number on which to send Syslog messages (default is 514 for UDP and 6514 for SSL); you must use the same port number on the firewall and the Syslog server.
 - **Format**—To separate individual syslog messages in a TCP stream, the delimiter formats available are LF- Line Feed (BSD Format, the default), and Message Length (IETF Format).
 - **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.
6. Click **OK** to save the server profile.

Configure the Transport Mechanism for Syslog (Continued)

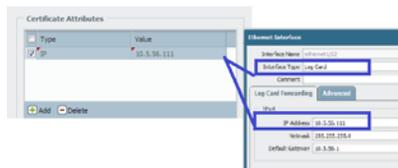
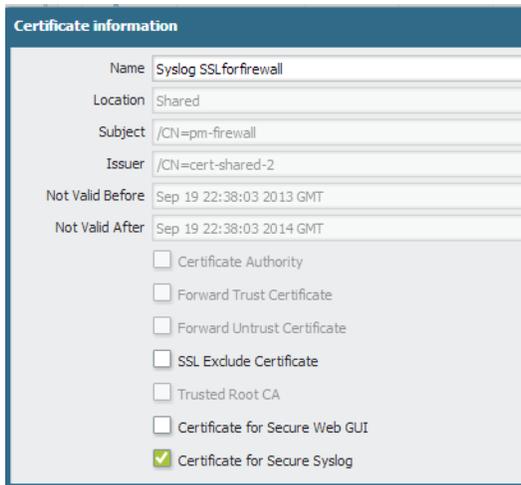
Step 2 If the syslog server requires client authentication, generate the certificate for secure communication.

To verify that the sending device (firewall or Panorama) 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 same trusted CA. Alternatively, you can generate a self-signed certificate on Panorama or the firewall, export the certificate from the firewall/Panorama and import it in to the syslog server.
- Use the trusted 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 or Panorama is authorized to communicate with the syslog server.

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

1. Select **Device (or Panorama) > 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 or FQDN 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. (Required on the PA-7050 for successful SSL negotiation) In the **Certificate Attributes**, **Add** the **IP** address of the **Log Card** as an attribute.



This value is then included as a Subject Alternate Name in the certificate.

8. Click **Generate**. The certificate and the keypair will be generated.
9. Click the link with name of the certificate and enable the option **Certificate for Secure Syslog** for secure access to the syslog server.
10. Verify the certificate details and that it is marked for **Usage as Certificate for Secure Syslog**.

Configure the Transport Mechanism for Syslog (Continued)

Name	Location	Subject	Issuer	CA	Key	Expires	Status	Usage
<input type="checkbox"/> 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
<input checked="" type="checkbox"/> Syslog SSLforfirewall	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

Step 3 (Only for managed collectors) On Panorama, select the certificate to use for secure syslog communication.

You must have imported the trusted CA certificate in to Panorama or generated it on Panorama. The certificate must be enabled for use as a **Certificate for Secure Syslog**.

1. Select **Panorama > Managed Collectors**.
 2. Click **Add** to add a new managed collector or select the link to edit the configuration for a managed collector.
 3. Select **General**, and choose the certificate from the **Certificate for Secure Syslog** drop-down.
- Note** You can only select from the certificate that are available on **Panorama > Certificate Management > Certificates**.

Step 4 Optional) Configure the header format used in Syslog messages.

See [Enhancement in the Syslog Header](#).

SNMP Resource Monitoring Extensions

SNMP allows you to use network management software to poll devices on the network. Monitoring devices helps you find trends in system resource usage, which provides insight in to system health and system utilization levels. This also helps with capacity planning.

In this release, the following SNMP extensions have been added:

- Ability to monitor management plane memory and swap utilization trends for detecting potential failures
- Ability to monitor dataplane packet buffer utilization for detecting failures
- Ability to monitor active VPN tunnel count for GlobalProtect gateway utilization
- Ability to monitor session utilization on a per virtual system basis in multi-tenant environments.

The following MIBS have been enhanced to support the resource monitoring capabilities listed above:

- HOST-RESOURCES-MIB (RFC 2970) hrStorageTable objects (OID .1.3.6.1.2.1.25.2.3)
 - Management plane memory monitoring
 - Dataplane packet buffer monitoring
- PAN-COMMON-MIB (firewalls only)
 - Session utilization on a per virtual system basis (OID: .1.3.6.1.4.1.25461.2.1.2.3.9)
 - GlobalProtect gateway utilization (OID: .1.3.6.1.4.1.25461.2.1.2.5.1)
- [Trap](#) OIDs for reporting failures for power supply, disks or fans on the appliances. The new OIDs range from .1.3.6.1.4.1.25461.2.1.3.2.0.901 to .1.3.6.1.4.1.25461.2.1.3.2.0.916.

Enhancement in the Syslog Header

If you send log data to a [syslog](#) server for centralized reporting, both the firewall and Panorama allow you to select the format of the hostname used in the syslog header entries. Choosing the header format offers more flexibility in filtering and reporting on the log data. The syslog header can now display one of the following: FQDN (hostname and domain name), hostname, the IPv4 address, or the IPv6 address of the sending device; or None to leave the hostname field in the syslog header empty.

Modify the Syslog Header Format

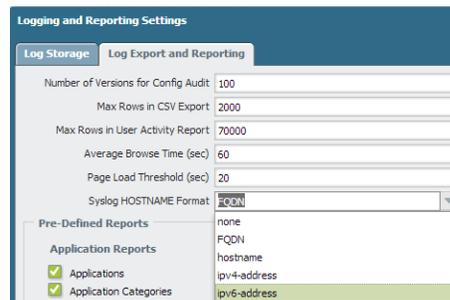
Step 1 Specify header format for Syslog messages.



This is a global setting and applies to all syslog server profiles configured on the appliance.

1. Select **Device > Setup > Management** and click the Edit icon in the Logging and Reporting Settings section.

2. Select **Log Export and Reporting**.



3. Select one of the following options from the **Send Hostname in Syslog** drop-down:

- **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**.



Networking Features

This chapter describes the following new networking features and provides instructions for setting them up:

- ▲ OSPF v3 Support
- ▲ OSPF Graceful Restart
- ▲ IKE PKI Certificate Authentication for IPsec Site-to-Site VPNs
- ▲ TLS 1.2 Decryption
- ▲ Decryption Port Mirror
- ▲ Increase Jumbo Frame Size
- ▲ Enhanced Use for Address Objects
- ▲ Consolidation of Timers Used in a High Availability Setup

OSPF v3 Support

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.

Configure OSPFv3 on the Virtual Router	
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 OSPFv3 internal, and OSPFv3 external.
Step 2 Create the virtual router.	<ol style="list-style-type: none"> 1. Select Network > Virtual Routers. 2. Click Add and enter a name for the virtual router. 3. Click Add in the Interfaces box and select an already defined interface from the drop-down box. 4. Repeat step 3 for all interfaces that you want to add to the virtual router. 5. Click OK.
Step 3 Set Administrative Distances for OSPFv3.	Set Administrative Distances as required. <ul style="list-style-type: none"> • OSPFv3 Internal – Range: 10-240, Default: 30 • OSPFv3 External – Range: 10-240, Default: 110

Configure OSPFv3 on the Virtual Router (Continued)	
<p>Step 4 Configure general OSPFv3 configuration settings.</p>	<ol style="list-style-type: none"> 1. Select the OSPFv3 sub tab. 2. Select the Enable check box to enable the OSPFv3 protocol. 3. Enter a Router ID. 4. 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. Clear the Reject Default Route check box if you want to permit redistribution of default routes through OSPFv3.
<p>Step 5 Configure an Auth Profile for the OSPFv3 protocol.</p> <p>While OSPFv3 doesn't include any authentication capabilities of its own, instead, it relies entirely on IPsec to secure communications between neighbors.</p>	<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. Select the Auth Profiles sub tab. 2. Click Add. 3. Enter a name for the authentication profile to authenticate OSPFv3 messages. 4. Specify a Security Policy Index (SPI). The SPI must match between both ends of the OSPFv3 adjacency. The SPI number must be a HEX value between 00000000 and FFFFFFFF. 5. Select ESP for Protocol. 6. Select a Crypto Algorithm from the drop-down. You can enter none or one of the following algorithms: SHA1, SHA256, SHA384, SHA512 or MD5. 7. If a Crypto Algorithm other than none was selected, enter a value for Key and then confirm. <p>AH OSPFv3 authentication</p> <ol style="list-style-type: none"> 1. Select the Auth Profiles sub tab. 2. Click Add. 3. Enter a name for the authentication profile to authenticate OSPFv3 messages. 4. Specify a Security Policy Index (SPI). The SPI must match between both ends of the OSPFv3 adjacency. The SPI number must be a HEX value between 00000000 and FFFFFFFF. 5. Select AH for Protocol. 6. Select a Crypto Algorithm from the drop down. You must enter one of the following algorithms: SHA1, SHA256, SHA384, SHA512 or MD5. 7. Enter a value for Key and then confirm. 8. Click OK. 9. Click OK again in the Virtual Router - OSPF Auth Profile configuration box.

Configure OSPFv3 on the Virtual Router (Continued)	
<p>Step 6 Configure Areas Type for the OSPFv3 protocol.</p>	<ol style="list-style-type: none"> 1. Select the Areas sub-tab. 2. Click Add. 3. Enter an Area ID. This is the identifier that each neighbor must accept to be part of the same area. 4. 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. • NSSA (Not-So-Stubby Area)—Traffic from 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 7 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. Select a previously defined Authentication Profile (Step 5) from the Authentication drop-down list on the General tab. 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.

Configure OSPFv3 on the Virtual Router (Continued)	
<p>Step 8 (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. The value can be: 1 - 65535. 7. Click OK.
<p>Step 9 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. Configure the Graceful Restart section as described in OSPF Graceful Restart.
<p>Step 10 Save virtual router general settings and commit.</p>	<p>Click OK to save your settings and click Commit.</p>

OSPF Graceful Restart

OSPF Graceful Restart directs OSPF neighbors to continue using routes through a device during a short transition when it is out of service. This increases network stability by reducing the frequency of routing table reconfiguration and the related route flapping that can occur during short periodic down times.

For a Palo Alto Networks firewall this involves the following operations:

- **Firewall as a restarting device**—In a situation where the firewall will be down for a short period of time or is unavailable for short intervals, it sends Grace LSAs to its OSPF neighbors. The neighbors must be configured to run in Graceful Restart Helper mode. In Helper Mode, the neighbors receive the Grace LSAs that inform it that the firewall will perform a graceful restart within a specified period of time defined as the **Grace Period**. During the grace period, the neighbor continues to forward routes through the firewall and to send LSAs that announce routes through the firewall. If the firewall resumes operation before expiration of the grace period, traffic forwarding will continue as before without network disruption. If the firewall does not resume operation after the grace period has expired, the neighbors will exit helper mode and resume normal operation which will involve reconfiguring the routing table to bypass the firewall.
- **Firewall as a Graceful Restart Helper**—In a situation where neighboring routers may be down for a short periods of time the firewall can be configured to operate in Graceful Restart Helper mode. If configured in this mode, the firewall will be configured with a **Max Neighbor Restart Time**. When the firewall receives the Grace LSAs from its OSPF neighbor, it will continue to route traffic to the neighbor and advertise routes through the neighbor until either the grace period or max neighbor restart time expires. If neither expires before the neighbor returns to service, traffic forwarding continues as before without network disruption. If either period expires before the neighbor returns to service, the firewall will exit helper mode and resume normal operation which will involve reconfiguring the routing table to bypass the neighbor.

Configure OSPF Graceful Restart

Step 1 Select **Network > Virtual Routers**.

Step 2 Click on the OSPF Virtual Router that you want to configure for Graceful Restart to bring up the configuration page.

Graceful restart can be configured for OSPF and OSPFv3 although it must be configured for each protocol separately.

Step 3 Select the **OSPF** (or **OSPFv3**) tab.

Step 4 Select the **Advanced** tab.

The following Graceful Restart check boxes are enabled by default:

Enable Graceful Restart, Enable Helper Mode, and Enable Strict LSA checking.

All should remain selected unless required by your topology.

Step 5 Configure a **Grace Period** in seconds.

Default value is 120. Acceptable range is 5 to 1800.

Step 6 Configure a **Max Neighbor Restart Time** in seconds.

Default value is 140. Acceptable range is 5 to 1800.

Step 7 Click **OK** to save your settings and click **Commit**.

IKE PKI Certificate Authentication for IPsec Site-to-Site VPNs

With this release authentication security has been enhanced over previous releases that only supported pre-shared-key VPN authentication. With this release the firewall can use IKE PKI certificate authentication for IP [Site-to-Site VPNs](#).

The PAN-OS implementation of IKE PKI certificate authentication supports the following:

- Conforms to RFC 2409 and RFC 4945.
- Supports the following ID types for authentication:
 - IP address – iPAddress from the SubjectAltName extension.
 - FQDN – dNSName from the SubjectAltName extension.
 - Email Address – User FQDN form frc822Name form the SubjectAltName extension.
 - Distinguished Name (DN) – The following are supported: Cert subject field – supports multiple OU fields, the entire DN when DN is used as the local ID type and DN as the ID type if the DN field of the certificate is empty.
- IKE Fragmentation Support:
 - Supports IKE to ensure IKE packets are not blocked by devices that do not permit UDP fragments.
 - Ensures that IKE initiator and receivers properly negotiate IKE fragmentation.
 - Ensures that IKE receivers can reassemble IKE fragmentation packets.
 - Maximum fragmented packet size is 576 bytes.

Feature Limitations

The following PAN-OS implementation of IKE PKI certificate authentication has the following limitations:

- The maximum length for a certificate chain is 5.
- CRL over LDAP is not supported.
- All IKE gateways configured on the same interface or local IP address must use the same crypto profile
- Authentication using Public Key Encryption and Public Key Encryption revised mode, described in RFC 2409, sections 5.2 and 5.3, respectively, are not currently supported.

Prepare a Firewall for IKE PKI Authentication

Preparing a firewall for IKE PKI certificate authentication, the following steps are required.

- ❑ **Obtain a signed certificate:** Generate a certificate and have it signed by a CA. See [Generate and Authenticate a Certificate](#).
- ❑ **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](#).
- ❑ **Configure the IKE Gateway:** The IKE gateway is configured to specify certificate authentication and sets the parameters for peer and local identification. It also specifies the certificate profile and other settings required for IKE authentication. See [Configure the IKE Gateway](#).

Generate and Authenticate a Certificate

The following process describes how to generate and authenticate a certificate on the Palo Alto Network firewall.

Generate a Certificate	
Step 1	Select Device > Certificate Management > Certificates and click Generate .
Step 2	Enter a Certificate Name such as <i>my-fw-trust</i> .
Step 3	Enter a Common Name , such as <i>192.169.3.1</i> . You can use an IP address or a Fully Qualified Domain Name (FQDN).
Step 4	Select a certificate from the Signed By drop-down to be used to authenticate this certificate. You can select a certificate that has been signed by an external authority and that has been imported into the firewall. The certificate must be currently valid and not expired.
Step 5	Select an OCSP responder to validate the certificate authority. OCSP checks to make sure the certificate has not been revoked.
Step 6	Set the Cryptographic Settings as described: Number of Bits – Select the key length for the certificate from the following: 512, 1024, 2048 3072. Digest – Select one of the following for the digest algorithm: md5, sha1, sha256, sha384, sha512 Expiration (days) – Specify the number of days during which the certificate will be valid. Default value is 365.
Step 7	Click Generate .
Step 8	Click OK to save.

Configure a Certificate Profile

The following process describes how to configure a certificate profile on a Palo Alto Networks firewall.

Create a Certificate Profile

Step 1 Select **Device > Certificates > Certificate Management > Certificate Profile** and click **Add** and enter a profile **Name**.

Step 2 Select a virtual system that you want the profile to apply to or select **Shared** if you want the profile to be shared between virtual systems.



Username Field and **Domain** are not used in this configuration.

Step 3 Click **Add** to enable the Certificate Profile configuration.

1. **CA Certificate** – Select a CA Certificate from the drop down box. See [Generate and Authenticate a Certificate](#).
 2. **Default OCSP URL** – Specify a default OCSP URL to check the revocation status of the CA certificate.
 3. **OCSP Verify CA Certificate** – (Optional) Select a separate OCSP responder to verify certificates.
-

Step 4 Select **Use CRL** to enable Certificate Revocation List (CRL) or **Use OCSP** to enable Online Certificate Status Protocol (OCSP).

You can enable either or both. If both CRL and OCSP are enabled, OCSP takes precedence.

Step 5 Select **Save** and **Commit**.

Configure the IKE Gateway

The following procedure describes how to configure an IKE Gateway for certificate-based authentication on a Palo Alto Network firewall.

Configure the IKE Gateway	
<p>Step 1 Configure the IKE gateway for certificate authentication.</p>	<ol style="list-style-type: none"> 1. Select Network > Network Profiles > IKE Gateways, click Add and enter a gateway Name. 2. Select the outgoing firewall Interface. 3. Select the IP address for the local interface that is the endpoint of the tunnel from the Local IP Address drop down list. 4. Select the Static or Dynamic option for the IP address of the peer on the far end of the tunnel. 5. If the Static option is selected for peer type, specify the IP address for the peer on the far end of the tunnel. 6. Select Certificate for the Authentication method and select the signed certificate from the Local Certificate drop-down. 7. From the Local Identification drop-down, choose one of the following types and enter the value: IP address, FQDN (hostname), User FQDN (email address), Distinguished Name (subject). 8. 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).
<p>Step 2 Configure the IKE gateway for certificate authentication. (Advanced Phase 1 Options tab)</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. 3. Select an existing profile or keep the default profile from IKE Crypto Profile drop-down. 4. Select Passive Mode Select to have the firewall respond only 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. 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) and delay before retrying (2 - 100 seconds). Dead peer detection identifies inactive or unavailable IKE peers through ICMP ping and can help restore resources that are lost when a peer is unavailable. 8. Click Save and Commit.

TLS 1.2 Decryption

Previous releases of PAN-OS only supported TLS version 1.1. This release provides the Palo Alto Networks firewall with the ability to decrypt inbound sessions and forward proxy sessions that negotiate with TLS 1.2. With this release TLS 1.2 is enabled by default and cannot be disabled. This implementation includes the following details:

- TLS 1.2 is supported as defined by RFC 5246.
- The following additional cipher suites are supported: `TLS_RSA_WITH_AES_128_CBC_SHA256` and `TLS_RSA_WITH_AES_256_CBC_SHA256`.
- Newer unsupported versions of TLS (1.3+) will be downgraded to 1.2 when used with the PAN-OS.

Increase Jumbo Frame Size

By default, the maximum transmission unit (MTU) size for packets sent on a Layer 3 interface is 1500 bytes. This size can be manually set to any size from 512 to 1500 bytes on a per-interface basis. Some configurations require Ethernet frames with an MTU value greater than 1500 bytes. These are called jumbo frames.

To use jumbo frames on a firewall you must specifically enable jumbo frames at the global level. When this is enabled, the default MTU size for all Layer 3 interfaces is set to a value of 9192 bytes. This default value can then be set to any value in the range of 512 to 9216 bytes.

After setting a global jumbo frame size it becomes the default value for all Layer 3 interfaces that have not explicitly had an MTU value set at the interface configuration level. This can become a problem if you only want to exchange jumbo frames on some interfaces. In these situations, you must set the MTU value at every Layer 3 interface that you do not want to use the default value.

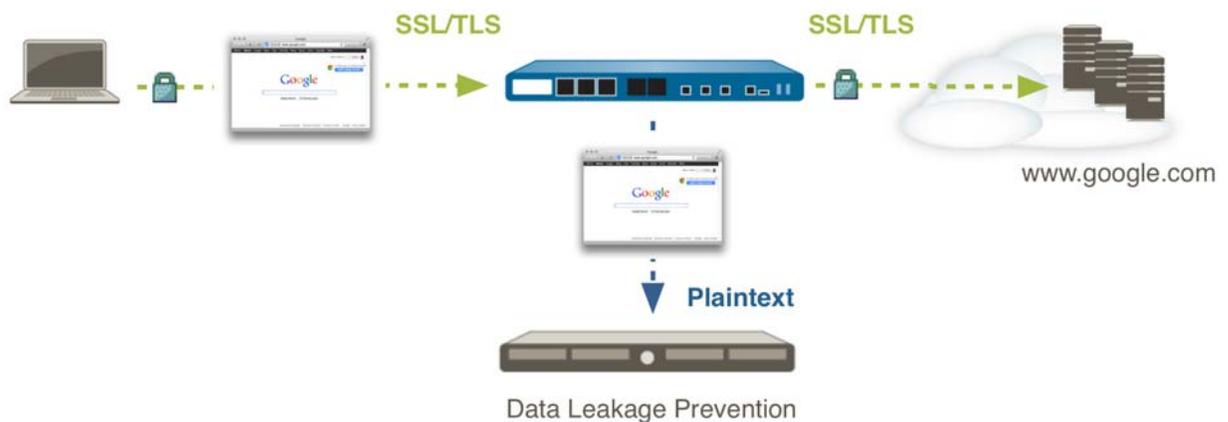
The following procedure describes how to enable jumbo frames on a firewall, set the default MTU value for all Layer 3 interfaces and to then set a different value for a specific interface.

Enable Jumbo Frames and Set MTU Values	
<p>Step 1 Enable jumbo frames and set a default global MTU value.</p>	<ol style="list-style-type: none"> 1. Select Device > Setup > Session and edit the Session Settings section. 2. Select Enable Jumbo Frame. 3. Enter a value for Global MTU. The default value is 9192. The range of acceptable values is: 512 - 9216. 4. Click OK. A message is displayed that informs you that enabling or disabling Jumbo Frame mode requires a reboot and that Layer 3 interfaces inherit the Global MTU value. 5. Click Yes. A message is displayed to inform you that Jumbo Frame support has been enabled and reminds you that a device reboot is required for this change to be activated. 6. Click OK. 7. Click Commit.
<p>Step 2 Set the MTU value for a Layer 3 interface and reboot the firewall.</p> <p> The value set for the interface overrides the global MTU value.</p>	<ol style="list-style-type: none"> 1. Select Network > Interfaces. 2. Select an interface of the Layer3 Interface type. 3. Select Advanced > Other Info. 4. Enter a value for MTU. The default value is 9192. The range of acceptable values is: 512 - 9216. 5. Click OK. 6. Click Commit. 7. Select Device > Setup > Operations and select Reboot Device.

Decryption Port Mirror

The **Decryption Port mirror** feature provides the ability to create a copy of decrypted traffic from a firewall and send it to a traffic collection tool that is capable of receiving raw packet captures—such as NetWitness or Solera—for archiving and analysis. This feature is necessary for organizations that require comprehensive data capture for forensic and historical purposes or data leak prevention (DLP) functionality. 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 may require user consent, 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 via an encrypted channel. Palo Alto Networks recommends that you consult with your corporate council before activating and using this feature in a production environment.



The following sections describe how to license and use this feature:

- ▲ [Obtain and Install a Decryption Port Mirror License](#)
- ▲ [Configure Decryption Port Mirroring](#)

Obtain and Install a Decryption Port Mirror License

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.

Install a Decryption Port Mirror License

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](#) portal 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

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	I6729162	Perpetual	⌵
Premium Support	I7480971	12/29/2015	

AVAILABLE FEATURE LICENSES

Decryption Port Mirror

Step 2 Install the Decryption Port Mirror license on firewall.

1. From the firewall's web interface, select **Device > Licenses**.
2. Click **Retrieve license keys from license server**.
3. Verify that the license has been activated on the firewall.

Decryption Port Mirror

Date Issued: January 06, 2014
Date Expires: Never
Description: Decryption Port Mirror
Active: Yes

4. Reboot the firewall (**Device > Setup > Operations**). This feature will not be available for configuration until PAN-OS reloads.

Configure Decryption Port Mirroring

To enable decryption port mirroring, you must enable the forwarding of decrypted traffic and configure a decrypt mirror interface, which is only available after you install the Decryption Port Mirror license and reboot the firewall as described in [Obtain and Install a Decryption Port Mirror License](#). You can then create a decryption profile that specifies the interface and attach it to a decryption policy.

Configure Decryption Port Mirroring	
<p>Step 1 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"> 1. Select Device > Setup > Content - ID. 2. Select the Allow forwarding of decrypted content check box. 3. Click OK to save. <p>On a firewall with multiple virtual systems:</p> <ol style="list-style-type: none"> 1. Select Device > Virtual System. 2. Select a Virtual System to edit or create a new Virtual System by selecting Add. 3. Select the Allow forwarding of decrypted content check box. 4. Click OK to save.
<p>Step 2 Configure a decrypt mirror interface.</p>	<ol style="list-style-type: none"> 1. Select Network > Interfaces > Ethernet. 2. Select the Ethernet interface that you want to configure for decryption port mirroring. 3. Select Decrypt Mirror as the Interface Type. This interface type will only appear if the Decryption Port Mirror license is installed. If you have not yet installed the license, see Obtain and Install a Decryption Port Mirror License. 4. Click OK to save.
<p>Step 3 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 port 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 Forward 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.

Configure Decryption Port Mirroring (Continued)

Step 4 Set a decryption policy for decryption port mirroring.	<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 the Decryption Profile created in Step 3.4. Click OK to save the policy.
Step 5 Save the configuration.	Click Commit .

Enhanced Use for Address Objects

An address object is a name value pair, it allows you to separate the configuration of the object from its IP address associations. With this release, you can select an address object when configuring a [Layer 3 interface](#) on a firewall, or when configuring a Layer 3 interface in a template on Panorama. The address object can include an IPv4 or IPv6 address (single IP, range, subnet) or the FQDN. It can be defined once and referenced in multiple places in configuration. When the IP address or range defined for the address object changes, you can edit the address object and the change in value is automatically inherited by all instances where the address object is used.

When using Panorama templates, the value for the address object can either be defined locally on the firewall or it can be defined as a shared object or as a device group object on Panorama. The advantage of defining the address object on the firewall is that you can configure a unique IP address for each managed device and do not have to override the value that is pushed to all the managed devices included in the template. In order to prevent a commit failure, you must create the address object on each firewall before pushing the template to the managed firewalls.

Use Address Objects

Step 1 Create an address object.

1. Select **Objects > Addresses**, and click **Add**.
2. Enter a **Name** and a **Description** for the address object.
3. Select **Type**—IP Netmask, IP range or FQDN.

4. (Optional) Enter or select a **Tag**. Tags allow you to filter and visually distinguish an object using color. For information on using tags, see [Support for Color Coded Tags](#).

Name	Location	Type	Address	Tags
Eth1-IP_1by7	Shared	IP Netmask	10.22.0.4/23	Ethernet1
ISP Provider	VSYS2 (vsys2)	FQDN	comcast.net	

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

Use Address Objects (Continued)

Step 2 Use the Layer 3 interface address object in place of an IP address when configuring an interface.

You can do this on the firewall or on Panorama. On Panorama, you can create a shared address object or a device group address object. This example shows how an address object is used to configure the interface in a Panorama template without creating the address object on Panorama. The value of the address object is defined locally on each managed device.

On the Firewall:

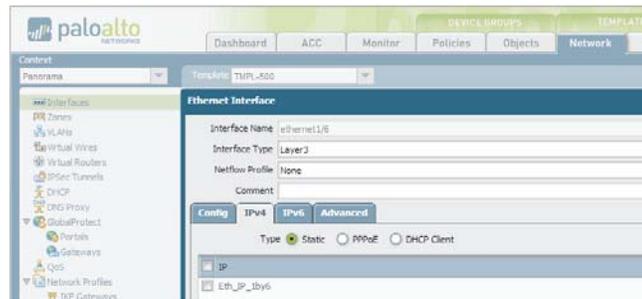
1. Select **Network > Interfaces > Ethernet**.
2. Select an interface, set the **Interface Type** as Layer 3.
3. Select the **IPv4** or **IPv6** tab, click **Add** and select the address object you created in Step 1.



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

In a Panorama Template:

1. Select **Network > Interfaces > Ethernet**.
2. Select the **Template** from the Template drop-down.
3. To configure an interface, click the link that corresponds to the interface and set the **Interface Type** as Layer 3.
4. Select the **IPv4** or **IPv6** tab, click **Add** and either select an address object you have already defined on Panorama or enter the name of the address object.



If you enter the name of the address object, you must use the same name and create the address object on the managed firewall before you push the configuration to the managed device from Panorama.

5. (Required only if you want to define the address object on the managed device) Define the address object locally on each managed firewall included in the template.
6. **Commit** the changes on **Panorama** and the **Template**.
7. Verify that the template settings have been applied to the managed device(s) included in the template and that you can view the IP address that you defined for the address object.

Interface	Interface Type	IP Address	Virtual Router	Security Zone	Feature
ethernet1/6	Layer3	Eth_IP_tby6			
ethernet1/7	Layer3	SOME_ADDRESS_OBJ	vr1		
ethernet1/8	Layer3	TMRI_SHAREDF_ADDR1	4-Fa/0/1		

Address	
Name:	Eth_IP_tby6
IP Netmask:	5.1.1.52/24

Consolidation of Timers Used in a High Availability Setup

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

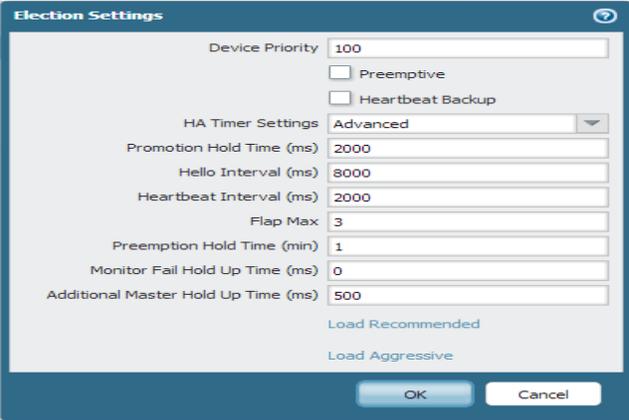
Use the **Recommended** profile for typical failover timer settings and the **Aggressive** profile for faster failover timer settings. The **Advanced** profile allows you to customize the timer values to suit your network requirements.

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.

Current 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	The interval during which the firewall will remain active following a path monitor or link monitor failure. This setting is recommended to avoid an HA failover due to the occasional flapping of neighboring devices.	0/0	0/0	0/0
Preemption hold time	Time a passive or active-secondary device will wait before taking over as the active or active-primary device.	1/1	1/1	1/1
Heartbeat interval	Specify how frequently the HA peers exchange heartbeat messages in the form of an ICMP ping.	1000/1000	2000/1000	2000/1000
Promotion hold time	Time that the passive device (in active/passive mode) or the active-secondary device (in active/active mode) will wait before taking over as the active or active-primary device after communications with the HA peer have been lost. This hold time will begin only after the peer failure declaration has been made.	2000/500	2000/500	2000/500

Timers	Description	PA-7050 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 device in active/passive mode and to the active-primary device in active/active mode. This timer is recommended to avoid a failover when both devices experience the same link/path monitor failure simultaneously.	500/500	500/500	7000/5000
Hello interval	The time interval in milliseconds between the hello packets that are sent to verify that the HA functionality on the other firewall is operational. The range is 8000-60000 ms with a default of 8000 ms for all platforms.	8000/8000	8000/8000	8000/8000
Maximum no. of flaps	A flap is counted when the firewall leaves the active state within 15 minutes after it last left the active state. This value indicates the maximum number of flaps that are permitted before the firewall is determined to be suspended and the passive firewall takes over (range 0-16, default 3).	3/3	3/3	Not Applicable

Use the High Availability Timer Profiles	
<p>Select the timer.</p>	<ol style="list-style-type: none"> <li data-bbox="781 275 1398 331">1. Select Device > High Availability, and edit the Election Settings section.  <ol style="list-style-type: none"> <li data-bbox="781 537 1446 594">2. Select the HA Timer Settings drop down. You can select Recommended, Aggressive or Advanced.
<p>View the preset values for each profile.</p> <p>See the table on the previous page for a description for each timer and the current preset values for your hardware platform.</p>	<ul style="list-style-type: none"> <li data-bbox="781 621 1463 741">• To view the preset values, select Advanced in the HA Timer Settings drop-down. The list of all the timers and their corresponding values is displayed. You can modify the values to meet your needs.  <ul style="list-style-type: none"> <li data-bbox="781 1207 1446 1264">• To load and view the preset values for the Recommended profile, click the Load Recommended link. <li data-bbox="781 1289 1479 1346">• To view the preset values for the Aggressive profile, click the Load Aggressive link.

On upgrade, the current/existing HA settings are saved to the **Advanced** profile. If you prefer to load the preset values for the **Recommended** or **Aggressive** profiles, use the instructions provide above.



Panorama Features

The following sections describe the new Panorama features and provide instructions for setting them up:

- ▲ [Panorama Log Forwarding](#)
- ▲ [Scheduled Dynamic Updates](#)
- ▲ [Support for Dual URL Filtering Databases](#)

Panorama Log Forwarding

All Palo Alto Networks next-generation firewalls can generate logs that provide an audit trail of the activities and events on the firewall. To centrally monitor the logs and to generate reports, you must forward the logs generated on the managed firewalls to Panorama. With this release, you can configure Panorama to [aggregate the logs and forward them](#) to a remote logging destination such as a syslog server.

In addition to logs, emails and SNMP traps can also be aggregated and forwarded from Panorama to a remote destination. Forwarding logs from Panorama reduces the load on the firewalls and provides a reliable and streamlined approach to combine and forward syslogs/SNMP traps/email notifications to remote destinations.

Table: Panorama Log Forwarding to an External Destination Per Platform

Platform/Deployment	Panorama Logs	Device Logs
Panorama Virtual Appliance	To forward Panorama logs: Panorama > Log Settings > System Panorama > Log Settings > Config	To forward device logs: Panorama > Log Settings Select the subtab for each log type: System, Config, Traffic, Threat, HIP Match and WildFire
Distributed Log Collection Architecture with: <ul style="list-style-type: none"> • Panorama M-100 with default Collector and/or Managed Collectors or <ul style="list-style-type: none"> • Panorama Virtual Appliance with Managed Collectors 	To forward both Panorama local logs and managed collector logs: Panorama > Log Settings > System Panorama > Log Settings > Config	To forward device logs that are being collected on the Log Collector Group: Panorama > Collector Groups > Collector Log Forwarding Select the subtab for each log type: System, Config, Traffic, Threat, HIP Match and WildFire



In order to forward logs from Panorama, you must have first configured the firewalls to forward logs to Panorama. For instructions, refer to the [Panorama Administrator's Guide](#).

Enable Log forwarding to External Destinations	
<p>Step 1 Set up server profiles for each external destination to which you want to forward logs.</p>	<ol style="list-style-type: none"> 1. Set up one or more of the following server profiles: <ol style="list-style-type: none"> a. SNMP: Panorama > Server Profiles > SNMP Trap b. Email: Panorama > Server Profiles > Email c. Syslog: Panorama > Server Profiles > Syslog <p>To forward logs to a syslog server, you can configure the transport medium to use UDP, TCP or SSL. See Support for Syslog Over TCP and SSL for setting up the syslog server profile.</p> <p>By default, each syslog entry is appended with the FQDN (hostname and domain name if configured) of the appliance that forwards the logs—Panorama or Managed Collector; the unique identifier of the firewall that generated the log entry is included in the data. To change what is listed in the syslog header, see Enhancement in the Syslog Header.</p>
<p>Step 2 Configure Panorama to forward logs.</p>	<ol style="list-style-type: none"> 1. Refer to Table: Panorama Log Forwarding to an External Destination Per Platform for instructions for details on forwarding logs for your platform/deployment. 2. Commit the changes on Panorama.

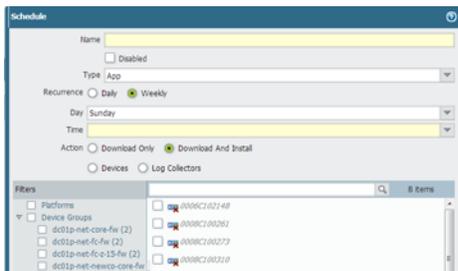
Scheduled Dynamic Updates

As a Panorama administrator, you can now schedule [dynamic content updates](#)—Apps and Threats, WildFire, Antivirus, and the BrightCloud URL Database updates—on the managed devices and managed collectors. You can install Apps, Threats, Antivirus, WildFire, and BrightCloud URL Filtering updates on the managed devices; only Apps and Threats updates can be installed on Log Collectors.

Schedule Dynamic Content Updates using Panorama

Step 1 Schedule dynamic updates.

You can define up to 100 scheduled content updates on Panorama.



1. Select **Panorama > Device Deployment > Dynamic Updates**.
2. Click **Schedules > Add** to set the schedule of each update type.
3. Add a **Name** to describe the schedule.
4. Select the **Type** of update, and specify how often you want the updates to occur by selecting a **Recurrence** value. 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).
5. 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. The timezone on Panorama is used to perform the download/installation.
6. Specify whether you want the system to **Download And Install** the update on the managed devices and collectors (best practice) or **Download Only**, where the content is downloaded to Panorama.
7. (For the **Download and Install** option) Select the **Devices/ Log Collectors** on which the update will be installed.
 -  Before the updates is installed on the device, a license check is performed. For a installation to succeed, the managed device must have a valid license.
8. Click **OK** to save the schedule settings.
9. Click **Commit** to save the settings to the running configuration.
10. Repeat the steps for each additional update you would like to schedule.



As a best practice, be sure to stagger the updates that you schedule because Panorama 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.

Schedule Dynamic Content Updates using Panorama (Continued)

Step 2 Verify the software and content update version running on each managed device and/or managed collector.



If the scheduled update fails, a system log is generated; if the scheduled update succeeds, a configuration log is generated.

Verify the Version for a Managed Device:

Select **Panorama > Managed Devices** and then locate the device(s) and review the content and software versions on the table.

Device Group	Device Name	Status		Software Version	Apps and Threat	Antivirus	URL Filtering	GlobalProtect Client	WildFire
		Conn...	Template						
▼ Branch (1/1 Devices Connected)									
Branch	SupportFW-07	<input checked="" type="checkbox"/>	In sync	5.0.0	347-1647	862-1186	4061	1.1.3	15901-23121

Verify the Version for a Managed Collector:

Enter the following command from the CLI:

```
show system info
```

The following details display after a successful update:

```
sw-version: 5.1.0-b10
app-version: 366-1738
app-release-date: 2013/03/29 15:46:03
av-version: 1168-1550
av-release-date: 2013/04/21 14:31:27
threat-version: 366-1738
threat-release-date: 2013/03/29 15:46:03
```

Support for Dual URL Filtering Databases

In order to create security, QoS, captive portal and decryption policies on Panorama that reference URL categories and/or URL profiles, you must enable a [URL Filtering vendor](#) on Panorama. Enabling a URL Filtering vendor on Panorama allows Panorama to obtain the content files that reference the URL categories that the vendor supports; the database is not downloaded on Panorama. While you can select only one URL filtering vendor—BrightCloud or PAN-DB—you can use the same shared policies to manage devices (running PAN-6.0) that are enabled for either vendor. This is possible because when you push policies from Panorama, managed devices running PAN-OS 6.0 can detect a vendor mismatch and automatically migrate URL categories and/or URL profiles to (one or more) categories that align with that of the vendor enabled on it.

Enable URL Filtering Vendor	
<p>Step 1 Select the URL filtering vendor to enable on Panorama.</p> <p> Enable the same URL Filtering vendor on Panorama as that of managed devices running PAN-OS versions earlier than 6.0. Then, when you push policies to managed devices running PAN-OS 6.0 that are enabled for the other vendor, 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.</p>	<ol style="list-style-type: none"> 1. Select Panorama > Setup > Management, and click the Edit button in the General Settings section. 2. Select the vendor from the URL Filtering Database drop-down.
<p>Step 2 Verify that the categories are available for referencing in policies.</p> <p>Because the URL database is not downloaded on Panorama, unlike the firewalls (Device > Licenses), you cannot view the download status of the database.</p>	<ol style="list-style-type: none"> 1. Select Objects > Security Profiles > URL Filtering. 2. Click Add and verify that the categories are displayed on the right pane of the URL Filtering Profile window. 



User-ID Features

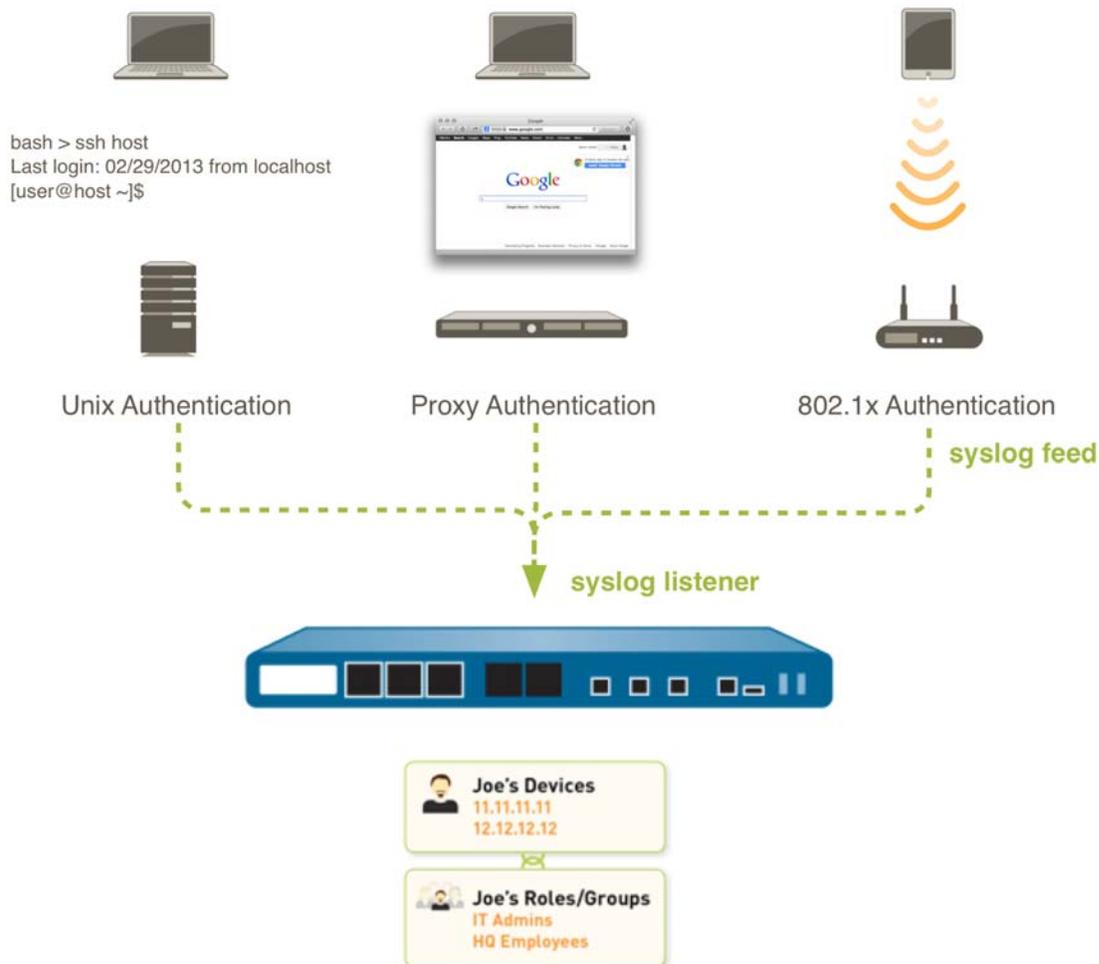
The following sections describe the new User-ID features and provide instructions for setting them up:

- ▲ [User-ID Integration With Syslog](#)
- ▲ [Support for Custom Terminal Service Solutions](#)

User-ID Integration With Syslog

In environments with existing network services that authenticate users, such as wireless controllers, 802.1x devices, Apple Open Directory servers, proxy servers, or other Network Access Control (NAC) mechanisms, the firewall **User-ID** agent (either the Windows agent or the agentless user mapping feature on the firewall) can now listen for authentication syslog messages from those services. Syslog filters, which you define, allow the User-ID agent to parse and extract usernames and IP addresses from authentication syslog events generated by the external service, and add the information to the User-ID IP address to username mappings maintained by the firewall. Previously this could only be done using the XML API interface.

Figure: Syslog Integration with User-ID



You can configure both the integrated PAN-OS User-ID agent and the Windows-based User-ID agent for syslog as described in the following sections:

- ▲ [Configure the PAN-OS Integrated User-ID Agent as a Syslog Listener](#)
- ▲ [Configure the Windows User-ID Agent as a Syslog Listener](#)

Configure the PAN-OS Integrated User-ID Agent as a Syslog Listener

The following procedure shows how to configure the User-ID agent on the firewall as a Syslog listener.



The PAN-OS integrated User-ID agent accepts syslogs over SSL and UDP only. As a best practice, always use SSL to listen for syslog messages. However, if you must use UDP, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending UDP traffic to the firewall.

Collect User Mappings from Syslog Senders

Step 1 Determine whether there is a pre-defined syslog filter for your particular syslog sender(s).

Palo Alto Networks provides several pre-defined syslog filters, which are delivered as Application content updates and are therefore updated dynamically as new filters are developed. The pre-defined filters are global to the firewall, whereas manually defined filters apply to a single virtual system only.



Any new syslog filters in a given content update will be documented in the corresponding release note along with the specific regex used to define the filter.

1. Verify that your Application or Application and Threat database is up to date:
 - a. Select **Device > Dynamic Updates**.
 - b. Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates.
 - c. If a new update is available, **Download** and **Install** it.
2. Check to see what pre-defined filters are available:
 - a. Select **Device > User Identification > User Mapping**.
 - b. In the **Server Monitoring** section of the screen, click **Add**.
 - c. Select **Syslog Sender** as the server **Type**.
 - d. Select the **Filter** drop-down and check to see if there is a filter for the manufacturer and product you plan to forward syslogs from. If the filter you need is available, skip to [Step 5](#) for instructions on defining the servers. If the filter you need is not available, continue to [Step 2](#).

The screenshot shows the configuration interface for a 'User Identification Monitored Server'. The 'Name' and 'Description' fields are empty. The 'Enabled' checkbox is checked. The 'Type' is set to 'Syslog Sender'. The 'Network Address' field is empty. The 'Connection Type' is set to 'SSL'. The 'Filter' dropdown menu is open, displaying a list of pre-defined filters. The 'Default Domain Name' is set to 'Aerohive AP v1.0.0'. The list of filters includes: Aerohive AP v1.0.0, BlueCoat Log Main Format Proxy Authentication, BlueCoat Proxy SG Proxy Log, BlueCoat Squid Web Proxy Authentication, Cisco ASA Any Connect v1.0.0, Cisco ASA IPsec v1.0.0, Citrix Access Gateway v1.0.0, Juniper IC v1.0.0, Juniper SA Net Connect v1.0.0, Squid Web Proxy Authentication, SSH Authentication, and Unix PAM Authentication.

Collect User Mappings from Syslog Senders (Continued)

Step 2 If the filter you need is not on the list of Filters, you must manually define syslog filter(s) for extracting the User-ID IP address to username mapping information from syslog messages.

In order to be parsed by the User-ID agent, syslog messages must meet the following criteria:

- Each syslog message must be a single line text string. Line breaks are delimited by a carriage return and a new line (\r\n) or a new line (\n).
- The maximum allowed size of an individual syslog message is 2048 bytes.
- Syslog messages sent over UDP must be contained in a single packet; messages sent over SSL can span multiple packets.
- A single packet may contain multiple syslog messages.

1. Review the syslogs generated by the authenticating service to identify the syntax of the login events. This enables you to create the matching patterns that will allow the firewall to identify and extract the authentication events from the syslogs.



While reviewing the syslogs also determine whether the domain name is included in the log entries. If the authentication logs do not contain domain information, consider defining a default domain name when adding the syslog sender to the monitored servers list in [Step 5](#).

2. Select **Device > User Identification > User Mapping** and edit the Palo Alto Networks User-ID Agent Setup section.

3. On the **Syslog Filters** tab, **Add** a new syslog parse profile.

4. Enter a name for the **Syslog Parse Profile**.

5. Specify the **Type** of parsing to use to filter out the user mapping information by selecting one of the following options:

- **Regex Identifier**—With this type of parsing, you specify regular expressions to describe search patterns for identifying and extracting user mapping information from syslog messages. Continue to [Step 3](#) for instructions on creating the regex identifiers.
- **Field Identifier** —With this type of parsing, you specify a string to match the authentication event, and prefix and suffix strings to identify the user mapping information in the syslogs. Continue to [Step 4](#) for instructions on creating the field identifiers.

Collect User Mappings from Syslog Senders (Continued)

Step 3 If you selected **Regex Identifier** as the parsing **Type**, create the regex matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a regex configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User: johndoe1
Source:192.168.3.212
```



If the syslog contains a standalone space and/or tab as a delimiter you must use an `\s` (for a space) and/or `\t` (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event Regex** field. For example, when matched against the example syslog message, the following regex instructs the firewall to extract the first {1} instance of the string `authentication success`. The backslash before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character: `(authentication\success){1}`.
2. Enter the regex for identifying the beginning of the username in the authentication success messages in the **Username Regex** field. For example, the regex `User:([a-zA-Z0-9\\._]+)` would match the string `User: johndoe1` in the example message and extract `acme\johndoe1` as the User-ID.
If the syslogs do not contain domain information and you require domain names in your user mappings, be sure to enter the **Default Domain Name** when defining the monitored server entry in [Step 5](#).
3. Enter the regex for identifying the IP address portion of the authentication success messages in the **Address Regex** field. For example, the following regular expression `Source:([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})` would match an IPv4 address (`Source:192.168.0.212` in the example syslog).

Collect User Mappings from Syslog Senders (Continued)

Step 4 If you selected **Field Identifier** as the parsing **Type**, define the string matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a field identifier configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoel
Source:192.168.3.212
```



If the syslog contains a standalone space and/or tab as a delimiter you must use an `\s` (for a space) and/or `\t` (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event String** field. For example, when matched against the sample syslog message, you would enter the string `authentication success` to identify authentication events in the syslog.
2. Enter the matching string for identifying the beginning of the username field within the authentication syslog message in the **Username Prefix** field. For example, the string `User :` identifies the beginning of the username field in the sample syslog.
3. Enter the **Username Delimiter** to mark the end of the username field within an authentication syslog message. For example, if the username is followed by a space, you would enter `\s` to indicate that the username field is delimited by a space in the sample log.
4. Enter the matching string for identifying the beginning of the IP address field within the authentication event log in the **Address Prefix** field. For example, the string `Source:` identifies the beginning of the address field in the example log.
5. Enter the **Address Delimiter** to mark the end of the IP address field within the authentication success message within the field. For example, if the address is followed by a line break, you would enter `\n` to indicate that the address field is delimited by a new line.

Step 5 Define the servers that will be sending syslog messages to the firewall for user mapping purposes.

You can define entries for up to 50 syslog senders per virtual system and up to a total of 100 monitored servers, including syslog senders, Microsoft Active Directory, Microsoft Exchange, or Novell eDirectory servers. The User-ID agent/firewall will discard any syslog messages received from servers that are not on this list.



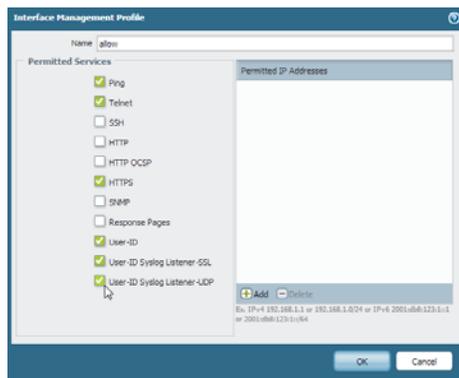
A Syslog sender using SSL to connect will only show a **Status** of **Connected** when there is an active SSL connection. Syslog senders using UDP will not show a **Status** value.

1. Select **Device > User Identification > User Mapping**.
2. In the Server Monitor section of the screen, click **Add**.
3. Enter a **Name** and **Network Address** for the server.
4. Select **Syslog Sender** as the server **Type**.
5. Make sure the **Enabled** check box is selected.
6. (Optional) If the syslogs that the authenticating device sends do not include domain information in the login event logs, enter the **Default Domain Name** to append to the user mappings.
7. Click **OK** to save the settings.

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.

 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.



Step 7 Save the configuration.

1. Select **Network > 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.



On the Windows User-ID agent, the default listening port for syslog over UDP or TCP is 514, but the port value is configurable. For the agentless User Mapping feature on the firewall only syslog over UDP and SSL are supported and the listening ports (514 for UDP and 6514 for SSL) are not configurable; they are enabled through the management service only.

3. Click **OK** to save the interface management profile.



Even after enabling the User-ID Syslog Listener service on the interface, the interface will only accept syslog connections from servers that have a corresponding entry in the User-ID monitored servers configuration. Connections or messages from servers that are not on the list will be discarded.

Click **Commit** to save the configuration.

Collect User Mappings from Syslog Senders (Continued)

Step 8 Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:

To see the status of a particular syslog sender:

```
admin@PA-5050> show user server-monitor state Syslog2
UDP Syslog Listener Service is enabled
SSL Syslog Listener Service is enabled

Proxy: Syslog2(vsys: vsys1)      Host: Syslog2(10.5.204.41)
number of log messages          : 1000
number of auth. success messages : 1000
number of active connections    : 0
total connections made          : 4
```

To see how many log messages came in from syslog senders and how many entries were successfully mapped:

```
admin@PA-5050> show user server-monitor statistics

Directory Servers:
Name                TYPE      Host           Vsys    Status
-----
AD                  AD        10.2.204.43   vsys1   Connected

Syslog Servers:
Name                Connection Host           Vsys    Status
-----
Syslog1             UDP           10.5.204.40   vsys1   N/A
Syslog2             SSL           10.5.204.41   vsys1   Not connected
```

To see how many user mappings were discovered through syslog senders:

```
admin@PA-5050> show user ip-user-mapping all type SYSLOG

IP                Vsys  From      User                               IdleTimeout(s) M
axTimeout(s)
-----
192.168.3.8      vsys1  SYSLOG   acme\jreddick                     2476           2
476
192.168.5.39    vsys1  SYSLOG   acme\jdonaldson                   2480           2
480
192.168.2.147   vsys1  SYSLOG   acme\ccrisp                       2476           2
476
192.168.2.175   vsys1  SYSLOG   acme\jjaso                        2476           2
476
192.168.4.196   vsys1  SYSLOG   acme\jblevins                     2480           2
480
192.168.4.103   vsys1  SYSLOG   acme\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

The following workflow describes how to configure a Windows-based User-ID agent to listen for syslogs from authenticating services.

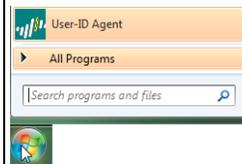


The Windows User-ID agent accepts syslogs over TCP and UDP only. However, you must use caution when using UDP to receive syslog messages because it is an unreliable protocol and as such there is no way to verify that a message was sent from a trusted syslog server. Although you can restrict syslog messages to specific source IP addresses, an attacker can still spoof the IP address, potentially allowing the injection of unauthorized syslog messages into the firewall. As a best practice, use TCP instead of UDP. In either case, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending syslogs to the User-ID agent.

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

Step 1 Launch the User-ID Agent application.

Click Start and select **User-ID Agent**.



Step 2 Manually define syslog filter(s) for extracting the User-ID IP address to username mapping information from syslog messages.

In order to be parsed by the User-ID agent, syslog messages must meet the following criteria:

- Each syslog message must be a single line text string. Line breaks are delimited by a carriage return and a new line (\r\n) or a new line (\n).
- The maximum allowed size of an individual syslog message is 2048 bytes.
- Syslog messages sent over UDP must be contained in a single packet; messages sent over SSL can span multiple packets.
- A single packet may contain multiple syslog messages.

1. Review the syslogs generated by the authenticating service to identify the syntax of the login events. This enables you to create the matching patterns that will allow the firewall to identify and extract the authentication events from the syslogs.



While reviewing the syslogs also determine whether the domain name is included in the log entries. If the authentication logs do not contain domain information, consider defining a default domain name when adding the syslog sender to the monitored servers list in [Step 5](#).

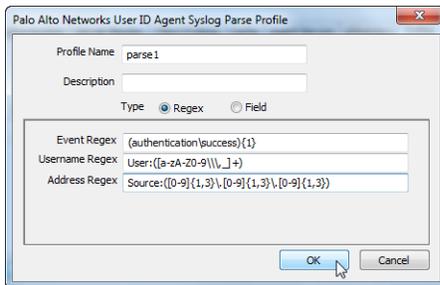
2. Select **User Identification > Setup** and click **Edit** in the Setup section of the dialog.
3. On the **Syslog** tab, **Add** a new syslog parse profile.
4. Enter a **Profile Name** and **Description**.
5. Specify the **Type** of parsing to use to filter out the user mapping information by selecting one of the following options:
 - **Regex**—With this type of parsing, you specify regular expressions to describe search patterns for identifying and extracting user mapping information from syslog messages. Continue to [Step 3](#) for instructions on creating the regex identifiers.
 - **Field**—With this type of parsing, you specify a string to match the authentication event, and prefix and suffix strings to identify the user mapping information in the syslogs. Continue to [Step 4](#) for instructions on creating the field identifiers.

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

Step 3 If you selected **Regex** as the parsing **Type**, create the regex matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a regex configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoe1
Source:192.168.3.212
```



 If the syslog contains a standalone space and/or tab as a delimiter you must use an `\s` (for a space) and/or `\t` (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event Regex** field. For example, when matched against the example syslog message, the following regex instructs the firewall to extract the first `{1}` instance of the string `authentication success`. The backslash before the space is a standard regex escape character that instructs the regex engine not to treat the space as a special character: `(authentication\success){1}`.
2. Enter the regex for identifying the beginning of the username in the authentication success messages in the **Username Regex** field. For example, the regex `User:([a-zA-Z0-9\\._]+)` would match the string `User:johndoe1` in the example message and extract `acme\johndoe1` as the User-ID.



If the syslogs do not contain domain information and you require domain names in your user mappings, be sure to enter the **Default Domain Name** when defining the monitored server entry in [Step 5](#).

3. Enter the regex for identifying the IP address portion of the authentication success messages in the **Address Regex** field. For example, the following regular expression `Source:([0-9]{1,3}\.){0-9}[0-9]{1,3}` would match an IPv4 address (`Source:192.168.0.212` in the example syslog).
4. Click **OK** to save the profile.

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

Step 4 If you selected **Field Identifier** as the parsing **Type**, define the string matching patterns for identifying the authentication events and extracting the user mapping information.

The example below shows a field identifier configuration for matching syslog messages with the following format:

```
[Tue Jul 5 13:15:04 2005 CDT] Administrator
authentication success User:johndoel
Source:192.168.3.212
```



If the syslog contains a standalone space and/or tab as a delimiter you must use an \s (for a space) and/or \t (for a tab) in order for the agent to parse the syslog.

1. Specify how to match successful authentication events in the syslogs by entering a matching pattern in the **Event String** field. For example, when matched against the sample syslog message, you would enter the string `authentication success` to identify authentication events in the syslog.
2. Enter the matching string for identifying the beginning of the username field within the authentication syslog message in the **Username Prefix** field. For example, the string `User :` identifies the beginning of the username field in the sample syslog.
3. Enter the **Username Delimiter** to mark the end of the username field within an authentication syslog message. For example, if the username is followed by a space, you would enter `\s` to indicate that the username field is delimited by a standalone space in the sample log.
4. Enter the matching string for identifying the beginning of the IP address field within the authentication event log in the **Address Prefix** field. For example, the string `Source:` identifies the beginning of the address field in the example log.
5. Enter the **Address Delimiter** to mark the end of the IP address field within the authentication success message within the field. For example, if the address is followed by a line break, you would enter `\n` to indicate that the address field is delimited by a new line.
6. Click **OK** to save the profile.

Step 5 Enable the syslog listening service on the agent.



As a best practice, make sure that the syslog server and client are both on a dedicated, secure VLAN to prevent untrusted hosts from sending syslogs to the User-ID agent.

1. Select the **Enable Syslog Service** check box.
2. (Optional) Modify the **Syslog Service Port** number to match the port number used by the syslog sender (Default=514).
3. To save the agent syslog configuration, click **OK**.

Step 6 Define the servers that will be sending syslog messages to the User-ID agent.

You can define entries for up to 100 syslog senders. The User-ID agent will discard any syslog messages received from servers that are not on this list.

1. Select **User Identification > Discovery**.
2. In the **Servers** section of the screen, click **Add**.
3. Enter a **Name** and **Server Address** for the server that will be sending syslogs to the agent.
4. Select **Syslog Sender** as the **Server Type**.
5. Select a **Filter** you defined in [Step 2](#).
6. (Optional) If the syslogs that the authenticating device sends do not include domain information in the login event logs, enter the **Default Domain Name** to append to the user mappings.
7. Click **OK** to save the settings.

Step 7 Save the configuration.

Click **Commit** to save the configuration.

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

Step 8 Verify the configuration by opening an SSH connection to the firewall and then running the following CLI commands:

To see the status of a particular syslog sender:

```
admin@PA-5050> show user server-monitor state Syslog2
UDP Syslog Listener Service is enabled
SSL Syslog Listener Service is enabled

Proxy: Syslog2(vsys: vsys1)   Host: Syslog2(10.5.204.41)
number of log messages       : 1000
number of auth. success messages : 1000
number of active connections  : 0
total connections made       : 4
```

To see how many log messages came in from syslog senders and how many entries were successfully mapped:

```
admin@PA-5050> show user server-monitor statistics

Directory Servers:
Name                TYPE      Host           Vsys    Status
-----
AD                  AD        10.2.204.43   vsys1   Connected

Syslog Servers:
Name                Connection Host           Vsys    Status
-----
Syslog1             UDP           10.5.204.40   vsys1   N/A
Syslog2             SSL           10.5.204.41   vsys1   Not connected
```

To see how many user mappings were discovered through syslog senders:

```
admin@PA-5050> show user ip-user-mapping all type SYSLOG

IP                Vsys  From      User                               IdleTimeout(s) M
axTimeout(s)
-----
192.168.3.8      vsys1  SYSLOG    acme\jreddick                     2476           2
476
192.168.5.39    vsys1  SYSLOG    acme\jdonaldson                   2480           2
480
192.168.2.147   vsys1  SYSLOG    acme\ccrisp                       2476           2
476
192.168.2.175   vsys1  SYSLOG    acme\jjaso                        2476           2
476
192.168.4.196   vsys1  SYSLOG    acme\jblevins                     2480           2
480
192.168.4.103   vsys1  SYSLOG    acme\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
```

Support for Custom Terminal Service Solutions

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. The User-ID XML API has been extended to provide similar functionality for non-Windows terminal servers as described in the following sections:

- ▲ [About the User-ID XML API Terminal Services Extensions](#)
- ▲ [Construct API Scripts to Send User Mapping Information to the Firewall](#)

About the User-ID XML API Terminal Services Extensions

The User-ID API is an XML 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 mappings from multi-user systems such as terminal servers, the User-ID XML API has been extended as follows:

- **<multiusersystem>**—A new User-ID XML API message format has been added in order to set 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>**—This new parameter has been added to 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 in its IP address-port-user mappings for enforcement of user- and group-based security policy rules. When a user exhausts all of 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.

Construct API Scripts to Send User Mapping Information to the Firewall

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 using the firewall's API key for secure communication.



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.

Use the User-ID XML API to Map Non-Windows Terminal Services Users

Step 1 Generate the API key that will be used to authenticate the API communication between the firewall and the Terminal server. To generate the key you must provide login credentials for an administrative account; the API is available to all administrators (including role based administrators with XML API privileges enabled).



Any special characters in the password must be URL/percent-encoded.

From a browser, log in to the firewall. Then, to generate the API key for the firewall, open a new browser window and enter the following URL:

```
https://<Firewall-IPaddress>/api/?type=keygen&user=<username>&password=<password>
```

Where <Firewall-IPaddress> is the IP address or FQDN of the firewall and <username> and <password> are the credentials for the administrative user account on the firewall. For example:

```
https://10.1.2.5/api/?type=keygen&user=admin&password=admin
```

The firewall responds with a message containing the key, for example:

```
<response status="success">
  <result>
    <key>k7J335J6hI7nBxIqyfa62sZugWx7ot%2BgzEA9UonlZRg=
  </key>
</result>
</response>
```

Use the User-ID XML API to Map Non-Windows Terminal Services Users (Continued)

<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: 1000 	<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 blocksize is 4000 and each multi-user system can allocate a maximum of 1000 blocks.</p> <p>If you define a custom blocksize and or port range, keep in mind that you must configure the values such that every port in the range gets allocated and that there are no gaps or unused ports. For example, if you set the port range to 1000-1499, you could set the block size to 100, but not to 200. This is because if you set it to 200, there would be unused ports at the end of the range.</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 blockstart values would be 1000, 1200, 1400, 1600, or 1800. Blockstart values of 1001, 1300, or 1850 would be unacceptable because some of the port numbers in the range would be left unused.</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 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 1000 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> </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%2BgzEA9UOnlZRg</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%2BgzEA9UOnlZRg&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%2BgzEA9UOnlZRg</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%2BgzEA9UOnlZRg&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
```



Virtualization Features

The following section describes the new features that support your virtual network needs and provides instructions for setting them up:

- ▲ VM Monitoring Agent
- ▲ Dynamic Address Groups
- ▲ Support for the VM-Series Firewall on the Citrix SDX Server
- ▲ Support for the VM-Series NSX Edition

VM Monitoring Agent

The firewall can monitor the VMware vCenter server and/or an ESX(i) server version 4.1 or 5.0, and poll for information on IP address and tags on newly provisioned Virtual Machines (VM), or on VMs that have been updated or moved on the network. In lieu of using external scripts and the XML API on the firewall, you can configure the **VM Information Sources** on the firewall to automate the mechanism for registering VM IP addresses and the tags associated with the VMs. This feature in conjunction with [VM Monitoring Agent](#), allows you to automatically synchronize the firewall with changes in the virtual environment and easily adapt and scale existing policies with changes on the virtual network.

You can configure up to 10 VM information sources on the firewall. This capability is also supported on the [Windows User-ID agent](#) and can monitor up to 100 sources. By default, the traffic between the firewall and the monitored sources uses the management (MGT) port on the firewall.

Set up the VM Monitoring Agent	
<p>Step 1 Enable the VM Monitoring Agent.</p> <p> Up to 10 sources can be configured for each firewall, or for each virtual system on a multiple virtual systems capable firewall. If you use the Windows User-ID agent, support for up to 100 sources is available.</p> <p>If your firewalls are configured in a high availability configuration:</p> <ul style="list-style-type: none"> • 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. 	<ol style="list-style-type: none"> 1. Select Device > VM Information Sources. 2. Click Add and enter the following information: <ol style="list-style-type: none"> a. A Name to identify the VMware ESX(i) or vCenter server that you want to monitor. b. Enter the Host information for the server—hostname or IP address and the Port on which it is listening. c. Select the Type to indicate whether the source is a VMware ESX(i) server or a VMware vCenter server. d. Add the credentials (Username and Password) to authenticate to the server specified above. e. (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. f. (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. g. Click OK, and Commit the changes. h. Verify that the connection Status displays as  connected.

Set up the VM Monitoring Agent (Continued)

Step 2 Verify the connection status.

1. Verify that the connection **Status** displays as  connected. If the connection status is pending or disconnected, verify that the source is operational and that the firewall is able to access the source. If you use a port other than the MGT port for communicating with the monitored source, you must change the service route (**Device > Setup > Services**, click the **Service Route Configuration** link and modify the **Source Interface** for the **VM Monitor** service).



Set up VM Monitoring on the User-ID Agent for Windows

Set up the VM Monitoring on the User-ID Agent

Step 1 Launch the User-ID Agent application.

1. Select **User-ID Agent** from the Windows **Start** menu.

Step 2 Add a VM information source.



The User-ID agent supports up to 100 VM Information sources.

1. Select **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.
 - Select the **Type** to indicate whether the source is a **VMware-ESX(i)** server or a **VMware-vCenter** server.
 - Enter the hostname or IP address of the VM **Host**.
 - Enter the **Port** number on which the VM source is listening.
 - Add the credentials (**Username** and **Password**) to authenticate to the server.
 - (Optional) Modify the **Update Interval (sec)** to a value in the range of 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.
 - Make sure the **Enable** check box is selected.
 - (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 **Timeout (hours)** 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**.

Set up the VM Monitoring on the User-ID Agent	
<p>Step 3 Verify the connection status.</p>	<p>Verify that the connection Status displays as Connected.</p> <p>If the connection status is pending or disconnected, verify that the source is operational and that the User-ID agent is able to access the source.</p>
<p>Step 4 Configure the firewall 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 information on IP addresses and tags on newly provisioned Virtual Machines (VM), or on VMs that have been updated or moved on the network:</p> <ol style="list-style-type: none"> 1. From the firewall web interface, 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.

Dynamic Address Groups

Dynamic address groups 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 its role on the network or the different kinds of traffic it processes.

A dynamic address group uses *Tags* as a filtering criteria to determine its members. A tag is a string or attribute that the firewall uses to match on and determine its group members. Tags use logical *and* and *or* operators for defining the filtering criteria.

Tags can be defined statically on the firewall and/or registered (dynamically) to the firewall. All entities that have the tags and match the defined criteria become members of the dynamic group. 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 in policy and the policy must be committed on the device.

The IP address and associated tags for an entity can be dynamically *registered* on the firewall using the XML API or the [VM Monitoring Agent](#) on the firewall; each registered IP address can have up to 32 tags. 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). Because the members of a dynamic address group are automatically updated, using dynamic address groups in lieu of static address objects, allows you to adapt to changes in your environment without relying on a system administrator to make policy changes and committing them on the firewall.

Use the following table to verify the maximum number of IP addresses that can be registered for each model of firewall:

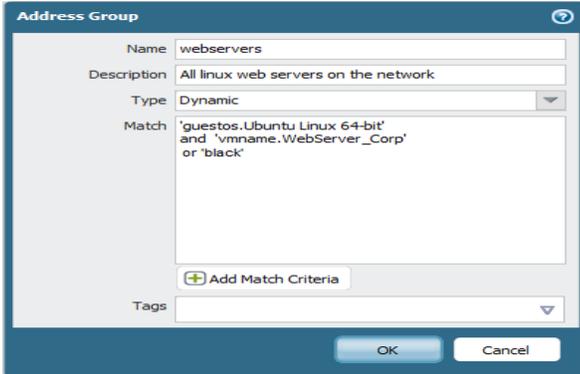
Platform	Maximum number of dynamically registered IP addresses
PA-7050, PA-5060, VM-1000	100,000
PA-5050	50,000
PA-5020	25,000
PA-4000 Series, PA-3000 Series	5000
PA-2000 Series, PA-500, PA-200, VM-300, VM-200, VM-100	1000

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

<p>Step 1 Enable VM Source Monitoring.</p>	<p>See Support for the VM-Series Firewall on the Citrix SDX Server.</p>
<p>Step 2 Create dynamic address groups on the firewall.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>View the tutorial to see a big picture view of the feature.</p> </div>	<ol style="list-style-type: none"> 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. <ol style="list-style-type: none"> a. 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. <div style="text-align: center; margin: 10px 0;">  </div> <ol style="list-style-type: none"> <li style="padding-left: 20px;">b. Click OK. <ol style="list-style-type: none"> 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 tftp	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.

Support for the VM-Series Firewall on the Citrix SDX Server

The Palo Alto Networks VM-Series firewall is the virtualized form of the Palo Alto Networks firewall. It is designed for use in a virtualized datacenter environment where it can protect and secure traffic within server networks; it is particularly well suited for private and public cloud deployments. To reduce your carbon footprint and consolidate key functions on a single server, you can deploy one or more instances of the VM-Series firewall on the [Citrix SDX](#) server and [VMware ESXi](#) server.

The following sections describe how to deploy the VM-Series firewall on a Citrix SDX server. In addition it details how to deploy the VM-Series firewall in conjunction with the NetScaler VPX, which is a virtual NetScaler appliance hosted on the SDX server. When deployed together, this solution enables secure application delivery along with network security, availability, performance, and visibility.

- ▲ [About the VM-Series Firewall on the Citrix SDX Server](#)
- ▲ [Integrate the VM-Series on to the SDX Server](#)
- ▲ [Supported Deployments](#)
- ▲ [Install the VM-Series Firewall](#)

About the VM-Series Firewall on the Citrix SDX Server

One or more instances of the VM-Series firewall can be deployed to secure east-west and/or north-south traffic on the network; virtual wire interfaces, Layer 2 interfaces, and Layer 3 interfaces are supported. To deploy the firewall, see [Support for the VM-Series NSX Edition](#).

Once deployed the VM-Series firewall works harmoniously with the NetScaler VPX (if needed), which is a virtual NetScaler appliance deployed on the SDX server. The NetScaler VPX provides load balancing and traffic management functionality and is typically deployed in front of a server farm to facilitate efficient access to the servers. For a complete overview of NetScaler feature/functionality, refer to the [Citrix NetScaler web site](#). When the VM-Series is paired to work with the NetScaler VPX, the complementary capabilities enhance your traffic management, load balancing, and application/network security needs.

This document assumes that you are familiar with the networking and configuration on the NetScaler VPX. In order to provide context for the terms used in this section, here is a brief refresher on the NetScaler owned IP addresses that are referred to in this document:

- NetScaler IP address (NSIP): The NSIP is the IP address for management and general system access to the NetScaler itself, and for HA communication.
- Mapped IP address (MIP): A MIP is used for server-side connections. It is not the IP address of the NetScaler. In most cases, when the NetScaler receives a packet, it replaces the source IP address with a MIP before sending the packet to the server. With the servers abstracted from the clients, the NetScaler manages connections more efficiently.
- Virtual server IP address (VIP): A VIP is the IP address associated with a vsriver. It is the public IP address to which clients connect. A NetScaler managing a wide range of traffic may have many VIPs configured.
- Subnet IP address (SNIP): When the NetScaler is attached to multiple subnets, SNIPs can be configured for use as MIPs providing access to those subnets. SNIPs may be bound to specific VLANs and interfaces.

For examples on deploying the VM-Series firewall and the NetScaler VPX together, see [Supported Deployments](#).

Requirements

You can deploy multiple instances of the VM-Series firewall on the Citrix SDX server. Because each instance of the firewall requires a minimum resource allocation—number of CPUs, memory and disk space—on the SDX server, make sure to conform to the specifications below to ensure optimal performance.

Requirement	Detail
SDX platforms	<ul style="list-style-type: none"> • 11500, 13500, 14500, 16500, 18500, 20500; • 17550, 19550, 20550, 21550
SDX version	10.1+ 10.1 is not supported; a software version higher than 10.1. is required.
Citrix XenServer version	6.0.2 or later
Minimum System Resources Plan and allocate the total number of data interfaces that you might require on the VM-Series firewall. This task is essential during initial deployment, because adding or removing interfaces to the VM-Series firewall after initial deployment will cause the data interfaces (Eth 1/1 and Eth 1/2) on the VM-Series firewall to re-map to the adapters on the SDX server. Each data interface sequentially maps to the adapter with the lowest numerical value, and this remapping can cause a configuration mismatch on the firewall.	<ul style="list-style-type: none"> • Two vCPUs per VM-Series firewall; one for the management plane and one for the dataplane. You can assign 2 or 6 additional vCPUs to allocate a total of 2, 4 or 8 vCPUs to the firewall; the management plane only uses one vCPU and any additional vCPUs are assigned to the dataplane. • Two network interfaces: one dedicated for management traffic and one for data traffic. For management traffic, you can use the 0/x interfaces on the management plane or the 10/x interfaces on the dataplane. Assign additional network interfaces for data traffic, as required for your network topology. • 4GB of memory. If you allocate additional memory, it will be used by the management plane only. • 40GB of virtual disk space. You can add disk space of up to 2TB; disk space in excess of the minimum 40GB requirement is used for logging purposes only.

Limitations

The VM-Series firewall deployed on the Citrix SDX server has the following limitations:

- Up to 24 total ports can be configured. One port will be used for management traffic and up to 23 can be used for data traffic.
- Jumbo frames are not supported.
- Link aggregation is not supported.

Integrate the VM-Series on to the SDX Server

On the SDX server, the VM-Series firewall can be deployed as a standalone virtual appliance that secures east-west and/or north-south traffic on the network. You can deploy the VM-Series firewall in virtual wire mode, Layer 2 mode or Layer 3 mode; one or more instances of the VM-Series can be installed on the SDX server. To deploy the firewall, see [Install the VM-Series Firewall](#).

In addition, the VM-Series firewall can be deployed to work harmoniously with the NetScaler VPX, which is a virtual NetScaler appliance deployed on the SDX server. The NetScaler VPX provides load balancing and traffic management functionality and is typically deployed in front of a server farm to facilitate efficient access to the servers. For a complete overview of NetScaler features/functionality, see <http://www.citrix.com/netscaler>. When the VM-Series is paired to work with the NetScaler VPX, the complementary capabilities enhance your traffic management, load balancing, and application/network security.

This article assumes that you are familiar with the networking and configuration on the NetScaler VPX. In order to provide context for the terms used in this section, here is a brief refresher on the NetScaler owned IP addresses that are referenced in this document:

- NetScaler IP address (NSIP): The NSIP is the IP address for management and general system access to the NetScaler itself, and for HA communication.
- Mapped IP address (MIP): A MIP is used for server-side connections. It is not the IP address of the NetScaler. In most cases, when the NetScaler receives a packet, it replaces the source IP address with a MIP before sending the packet to the server. With the servers abstracted from the clients, the NetScaler manages connections more efficiently.
- Virtual server IP address (VIP): A VIP is the IP address associated with a vserver. It is the public IP address to which clients connect. A NetScaler managing a wide range of traffic may have many VIPs configured.
- Subnet IP address (SNIP): When the NetScaler is attached to multiple subnets, SNIPs can be configured for use as MIPs providing access to those subnets. SNIPs may be bound to specific VLANs and interfaces.

For examples and instructions on deploying the VM-Series firewall and the NetScaler VPX together, see [Supported Deployments](#).

Supported Deployments

You can deploy one or more instances of the VM-Series firewall on the SDX server. In the following scenarios, the VM-Series firewall secures traffic destined to the servers on the network. It works in conjunction with the NetScaler VPX to manage traffic before or after it reaches the NetScaler VPX.

- ▲ [Scenario 1—Secure North-South Traffic](#)
- ▲ [Scenario 2—Secure East-West Traffic](#)

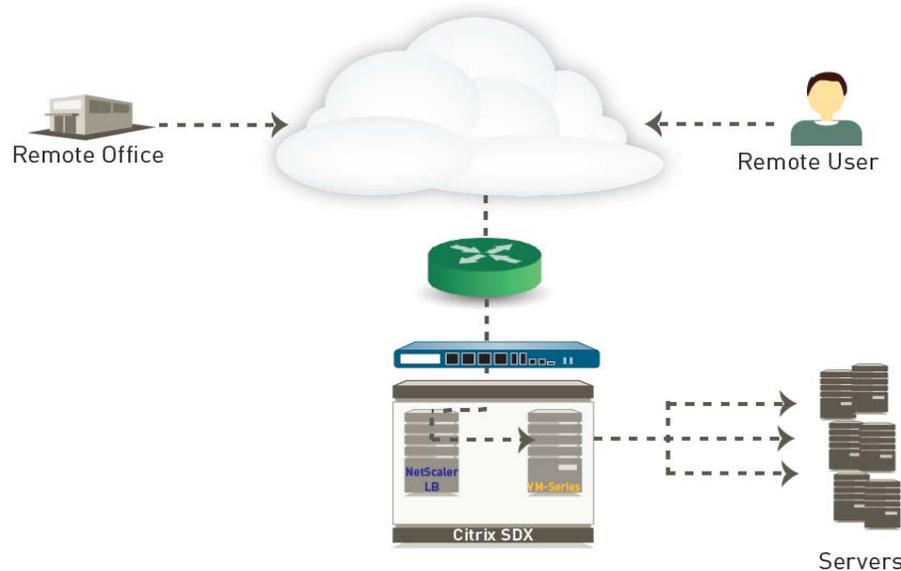
Scenario 1—Secure North-South Traffic

To secure north-south traffic, you have the following options:

- ▲ VM-Series Firewall Between the NetScaler VPX and the Servers
- ▲ VM-Series Firewall Before the NetScaler VPX

VM-Series Firewall Between the NetScaler VPX and the Servers

The perimeter firewall gates all traffic in to the network. All traffic permitted into the network flows through the NetScaler VPX and then through the VM-Series firewall before the request is forwarded to the servers.



In this scenario, the VM-Series firewall secures north-south traffic and can be deployed using virtual wire, L2, or L3 interfaces.

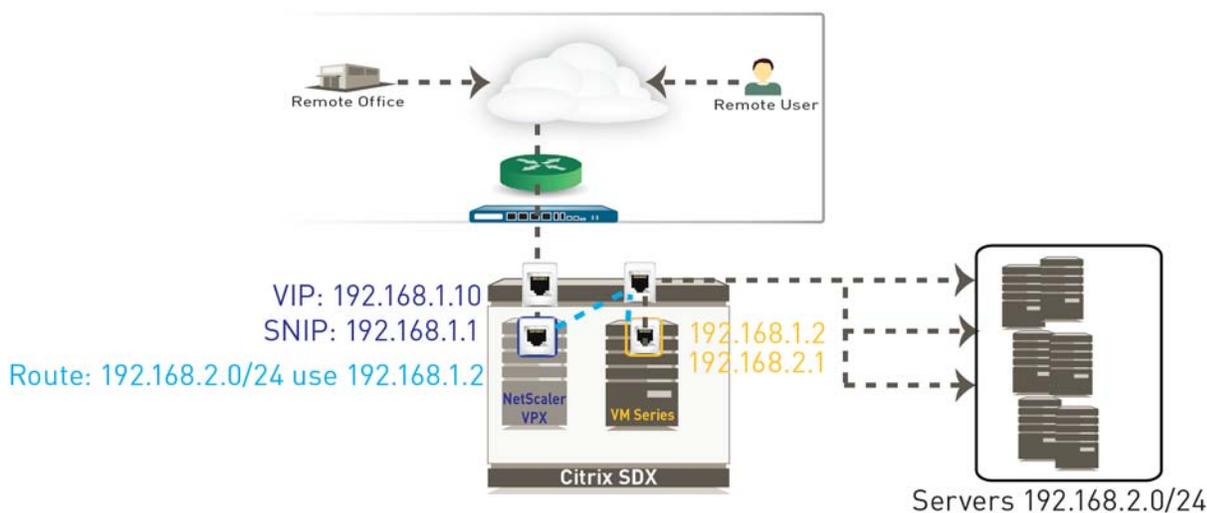
- ▲ VM-Series Firewall with L3 Interfaces
- ▲ VM-Series Firewall with L2 or Virtual Wire Interfaces

VM-Series Firewall with L3 Interfaces

Deploying the firewall with L3 interfaces allows you to scale more easily as you deploy new servers and new subnets. You can deploy multiple instances of the firewall to manage traffic to each new subnet and then configure the firewalls as a high availability pair, if needed.

Using an L3 interface allows you make minimal changes to the SDX server/network configuration because the SNIP to reach the servers is removed from the NetScaler VPX and is configured on the VM-Series firewall. With this approach, only one data interface is used on the VM-Series firewall, hence only one zone can be defined. As a result, when defining the policy rules you must specify the source and destination IP address/subnets across which to enforce security rules. For details, see [Deploy the VM-Series Firewall Using L3 Interfaces](#) in the [VM-Series Deployment Guide](#).

Topology After Adding the VM-Series Firewall with L3 Interfaces



In this example, the public IP address that the clients connect to (VIP on the NetScaler VPX), is 192.168.1.10. For providing access to the servers on subnet 192.168.2.x, the configuration on the VPX references the subnets (SNIP) 192.168.1.1 and 192.168.2.1. Based on your network configuration and default routes, the routing on servers might need to be changed.

When you set up the VM-Series firewall, you must add a data interface (for example eth1/1), and assign two IP addresses to the interface. One IP address must be on the same subnet as the VIP and the other must be on the same subnet as the servers. In this example, the IP addresses assigned to the data interfaces are 192.168.1.2 and 192.168.2.1. Because only one data interface is used on the VM-Series firewall, all traffic belongs to a single zone, and all intra zone traffic is implicitly allowed in policy. Therefore, when defining the policy rules you must specify the source and destination IP address/subnets across which to enforce security rules.

Even after you add the VM-Series firewall on the SDX server, the IP address that the clients continue to connect to is the VIP of the NetScaler VPX (192.168.1.10). However, to route all traffic through the firewall, on the NetScaler VPX you must define a route to the subnet 192.168.2.x. In this example, to access the servers this route must reference the IP address 192.168.1.2 assigned to the data interface on the VM-Series firewall. Now all traffic destined for the servers is routed from the NetScaler VPX to the firewall and then on to the servers. The return traffic uses the interface 192.168.2.1 on the VM-Series and uses the SNIP 192.168.1.1 as its next hop.



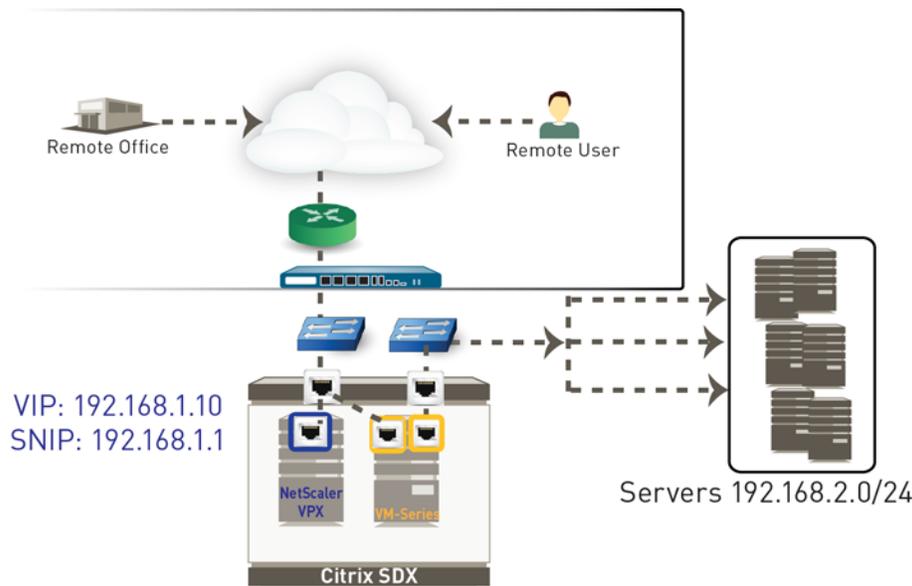
For security compliance, if USIP (Use client Source IP) is enabled on the NetScaler VPX, then the VM-Series firewall requires a default route that points to the SNIP 192.168.1.1, in this example. If a default NAT (mapped/SNIP) IP address is used, then you do not need to define a default route on the VM-Series firewall.

For instructions, see [Deploy the VM-Series Firewall Using L3 Interfaces](#) in the [VM-Series Deployment Guide](#).

VM-Series Firewall with L2 or Virtual Wire Interfaces

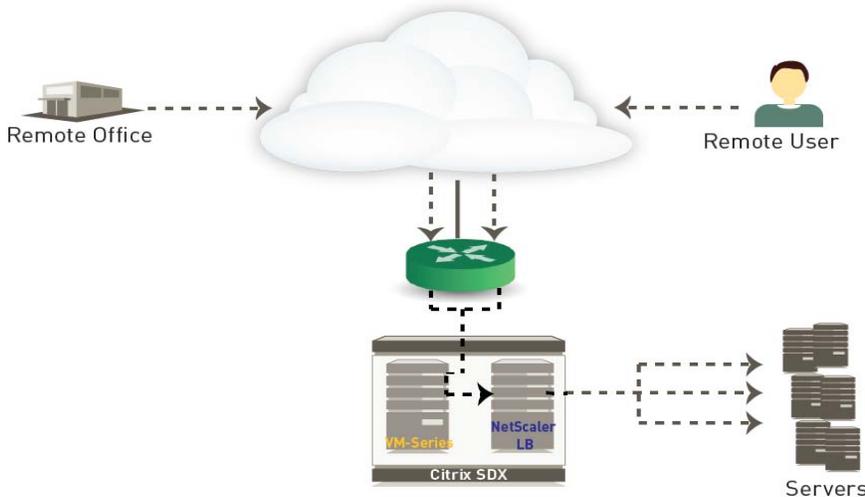
Deploying the VM-Series firewall using L2 interfaces or virtual wire interfaces requires reconfiguration on the NetScaler VPX to remove direct connection to the servers. The VM-Series firewall can then be cabled and configured to transparently intercept and enforce policy on traffic destined to the servers. In this approach two data interfaces are created on the firewall and each belongs to a distinct zone. The security policy is defined to allow traffic between the source and destination zones. For details, see [Deploy the VM-Series Firewall Using Layer 2 \(L2\) or Virtual Wire Interfaces](#) in the [VM-Series Deployment Guide](#).

Topology After Adding the VM-Series Firewall with L2 or Virtual Wire Interfaces



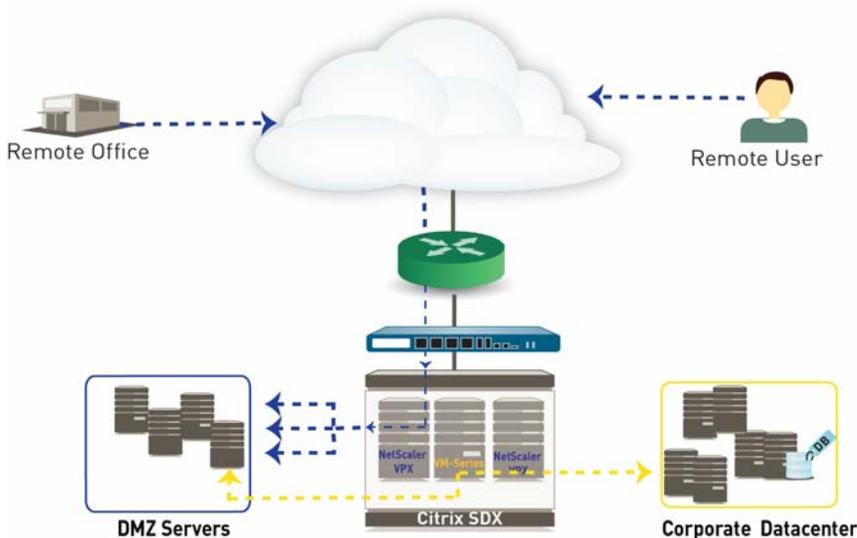
VM-Series Firewall Before the NetScaler VPX

In this scenario, the perimeter firewall is replaced with the VM-Series firewall that can be deployed using L3, L2, or virtual wire interfaces. All traffic on your network is secured by the VM-Series firewall before the request reaches the NetScaler VPX and is forwarded to the servers. For details, see [Deploy the VM-Series Firewall Before the NetScaler VPX](#) in the [VM-Series Deployment Guide](#).



Scenario 2—Secure East-West Traffic

The VM-Series firewall is deployed along with two NetScaler VPX systems that service different server segments on your network or operate as termination points for SSL tunnels. In this scenario, the perimeter firewall secures incoming traffic. Then, the traffic destined to the DMZ servers flows to a NetScaler VPX that load balances the request. To add an extra layer of security to the internal network, all east-west traffic between the DMZ and the corporate network are routed through the VM-Series firewall. The firewall can enforce network security and validate access for that traffic. For details, see [Secure East-West Traffic with the VM-Series Firewall](#) in the [VM-Series Deployment Guide](#).



Install the VM-Series Firewall

A support account and a valid VM-Series license are required to obtain the .xva base image file that is required to install the VM-Series firewall on the SDX server. If you have not already registered the capacity auth-code that you received with the order fulfillment email, with your support account, see [Register the VM-Series Firewall in the VM-Series Deployment Guide](#). After registration is completed, continue to the following tasks:

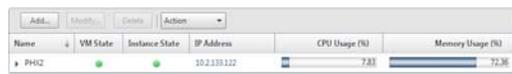
- ▲ [Upload the Image to the SDX Server](#)
- ▲ [Provision the VM-Series Firewall](#)

Upload the Image to the SDX Server

To provision the VM-Series firewall, you need to obtain the .xva image file and upload it to the SDX server.

Upload the XVA Image to the SDX Server	
<p>Step 1 Download and extract the base image zip file to a local computer.</p>	<ol style="list-style-type: none"> Go to https://support.paloaltonetworks.com/ and download the VM-Series Citrix SDX Base Image zip file. Unzip the base image zip file, and extract the .xva file. This .xva file is required for installing the VM-Series firewall.
<p>Step 2 Upload the image from the local computer onto the Citrix SDX server.</p>	<ol style="list-style-type: none"> Launch the web browser and log in to the SDX server. Select Configuration > Palo Alto VM-Series > Software Images In the Action drop-down, select Upload... and Browse to the location of the saved .xva image file. Select the image and click Open. Upload the image to the SDX server. 

Provision the VM-Series Firewall

Provision the VM-Series Firewall on the SDX Server													
<p>Step 1 Access the SDX server.</p>	<p>Launch the web browser and connect to the SDX server.</p>												
<p>Step 2 Create the VM-Series firewall.</p>  <p>Allocate the total number of data interfaces that you might require on the VM-Series firewall during initial deployment. Adding or removing interfaces to the VM-Series firewall after initial deployment will cause the data interfaces (Eth 1/1 and Eth 1/2) on the VM-Series firewall to re-map to the adapters on the SDX server. Each data interface sequentially maps to the adapter with the lowest numerical value, and can therefore cause a configuration mismatch on the firewall.</p>	<ol style="list-style-type: none"> 1. Select Configuration > Palo Alto VM-Series > Instances. 2. Click Add. 3. Enter a name for the VM-Series firewall. 4. Select the .xva image that you uploaded earlier. This image is required to provision the firewall. 5. Allocate the memory, additional disk space, and the virtual CPUs for the VM-Series firewall. To verify resource allocation recommendations, see Requirements. 6. Select the network interfaces. <ol style="list-style-type: none"> a. Use the management interfaces 0/1 or 0/2 and assign an IP address, netmask, and gateway IP address.  <p>If needed, you can use a data interface on the SDX server for managing the firewall.</p> <ol style="list-style-type: none"> b. Select the data interfaces that will be used for handling traffic to and from the firewall. b. If you plan to deploy the interfaces as Layer 2 or virtual wire interfaces, select the Allow L2 Mode option so that the firewall can receive and forward packets for MAC addresses other than its own MAC address.  7. Review the summary and click Finish to begin the installation process. It takes 5-8 minutes to provision the firewall. When completed, use the management IP address to launch the web interface of the firewall.  <table border="1"> <thead> <tr> <th>Name</th> <th>VM State</th> <th>Instance State</th> <th>IP Address</th> <th>CPU Usage (%)</th> <th>Memory Usage (%)</th> </tr> </thead> <tbody> <tr> <td>PHQ2</td> <td>Running</td> <td>Running</td> <td>10.2.10.102</td> <td>78</td> <td>72</td> </tr> </tbody> </table>	Name	VM State	Instance State	IP Address	CPU Usage (%)	Memory Usage (%)	PHQ2	Running	Running	10.2.10.102	78	72
Name	VM State	Instance State	IP Address	CPU Usage (%)	Memory Usage (%)								
PHQ2	Running	Running	10.2.10.102	78	72								

Support for the VM-Series NSX Edition

The [VM-Series NSX edition firewall](#) is jointly developed by Palo Alto Networks and VMware. This solution uses the NetX API to integrate the Palo Alto Networks next-generation firewalls and Panorama with VMware ESXi servers to provide comprehensive visibility and safe application enablement of all datacenter traffic including intra-host virtual machine communications.

The following topics provide information about the VM-Series NSX edition firewall:

- ▲ [VM-Series NSX Edition Firewall Overview](#)
- ▲ [Deploy the VM-Series NSX Edition Firewall](#)

VM-Series NSX Edition Firewall Overview

NSX, the VMware Networking and Security platform designed for the software-defined data center (SDDC), offers the ability to deploy the Palo Alto Networks firewall as a service on ESXi servers. The term *software-defined data center (SDDC)* is a VMware term that refers to a data center where infrastructure—compute resources, network, and storage—is virtualized using VMware NSX.

To keep pace with the changes in the agile SDDC, the NSX edition of the VM-Series firewall simplifies the process of deploying a next-generation firewall and continually enforcing security and compliance for the east-west traffic in the SDDC. For details on the VM-Series NSX edition, see the following topics:

- ▲ [What are the Components of the Solution?](#)
- ▲ [How Do the Components Work Together?](#)
- ▲ [What are the Benefits of the Solution?](#)

What are the Components of the Solution?

The components of this joint Palo Alto Networks and VMware solution are:

Provider	Component	Version	Description
VMware	vCenter Server	5.5	The vCenter server is the centralized management tool for the vSphere suite.
	NSX Manager	6.0	VMware's Networking and Security platform must be installed and registered with the vCenter server. The NSX Manager is required to deploy the VM-Series NSX edition firewall on the ESXi hosts within a ESXi cluster.
	ESXi Server	5.5	ESXi is a hypervisor that enables compute virtualization.

Provider	Component	Version	Description
Palo Alto Networks	PAN-OS	6.0	<p>The VM-Series base image (PA-VM-NSX-6.0.0.zip) used for deploying the VM-Series NSX edition firewall is PAN-OS version 6.0.</p> <p>The minimum system requirements for deploying the VM-Series NSX edition firewall on the ESXi server are as follows:</p> <ul style="list-style-type: none"> Two vCPUs. One for the management plane and one for the dataplane. You can assign 2 or 6 additional vCPUs to allocate a total of 2, 4 or 8 vCPUs to the firewall; the management plane only uses one vCPU and any additional vCPUs are assigned to the dataplane. 5GB of memory. Any additional memory will be used by the management plane only. 40GB of virtual disk space.
	Panorama	6.0	<p>Panorama is the centralized management tool for the Palo Alto Networks next-generation firewalls. In this solution, Panorama works with the NSX Manager to deploy, license, and centrally administer—configuration and policies—on the VM-Series NSX edition firewall.</p> <p>Panorama must be able to connect to the NSX Manager, the vCenter server, the VM-Series firewalls and the Palo Alto Networks update server.</p> <p>The minimum system requirement for Panorama is as follows:</p> <ul style="list-style-type: none"> Two 8-Core vCPUs (2.2GHz); use 3GHz if you have 10 or more firewalls 4GB RAM; 16GB recommended if have 10 or more firewalls <p>40GB disk space; To expand log capacity, you must add a virtual disk or set up access to an NFS datastore. For details, refer to the Panorama Administrator's Guide.</p>
	VM-Series NSX Edition	6.0	The only VM-Series license available in this solution is the VM-1000 in hypervisor mode (VM-1000-HV).

vCenter Server

The vCenter server is required to manage the NSX Manager and the ESXi hosts in your datacenter. This joint solution requires that the ESXi hosts be organized into one or more clusters on the vCenter server and must be connected to a distributed virtual switch.

For information on clusters, distributed virtual switch, DRS, and the vCenter server, refer to your VMware documentation: <https://www.vmware.com/support/vcenter-server.html>.

NSX Manager

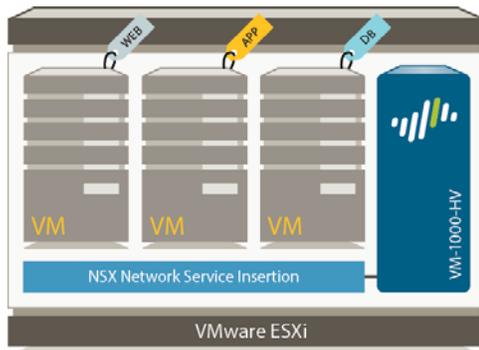
NSX is VMware's network virtualization platform that is completely integrated with vSphere. The NSX Firewall and the Service Composer are key features of the NSX Manager. The NSX firewall is a logical firewall that allows you to attach network and security services to the virtual machines, and the Service Composer allows you to group virtual machines and create policy to redirect traffic to the VM-Series firewall (Palo Alto Networks NGFW service).

Panorama

Panorama is used to register the NSX edition of the VM-Series firewall as the Palo Alto Networks next-generation firewall (NGFW) service on the NSX Manager. Registering the Palo Alto Networks NGFW service on the NSX Manager allows the NSX Manager to deploy the NSX edition of the VM-Series firewall on each ESXi host in the ESXi cluster.

Panorama serves as the central point of administration of the VM-Series NSX edition firewalls. When a new VM-Series NSX edition firewall is deployed, it communicates with Panorama to obtain the license and receives its configuration/policies from Panorama. All configuration elements, policies, and Dynamic Address Groups on the VM-Series NSX edition firewalls can be centrally managed on Panorama using Device Groups and Templates. The REST-based XML API integration in this solution, enables Panorama to synchronize with the NSX Manager and the VM-Series NSX edition firewalls to allow the use of Dynamic Address Groups and share context between the virtualized environment and security enforcement. For more information, see [Policy Enforcement using Dynamic Address Groups](#).

VM-Series NSX Edition



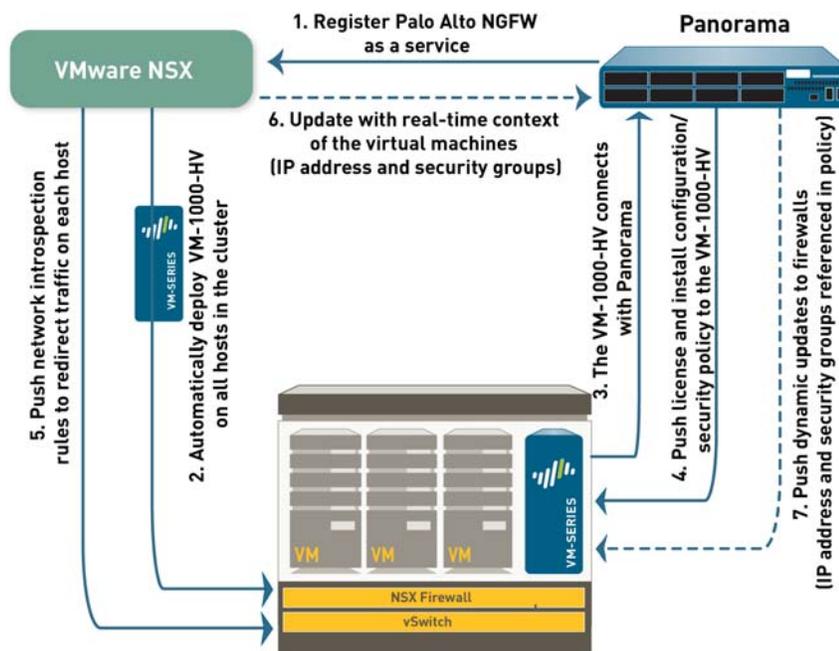
The VM-Series NSX edition is the VM-Series firewall that is deployed on the ESXi hypervisor. The integration with the NetX API makes it possible to automate the process of installing the VM-Series firewall directly on the ESXi hypervisor, and allows the hypervisor to forward traffic to the VM-Series firewall without using the vSwitch configuration; it therefore, requires no change to the virtual network topology.

The VM-Series NSX edition only supports virtual wire interfaces. In this edition, ethernet 1/1 and ethernet 1/2 are bound together through a virtual wire and use the NetX dataplane API to communicate with the hypervisor. Layer 2 or Layer 3 interfaces are neither required nor supported on the VM-Series NSX edition, and therefore no switching or routing actions can be performed by the firewall.

The only license available for this version of the VM-Series firewall is the VM-1000-HV. For complete information on the maximum capacities supported on the VM-1000-HV license refer to the VM-Series datasheet.

How Do the Components Work Together?

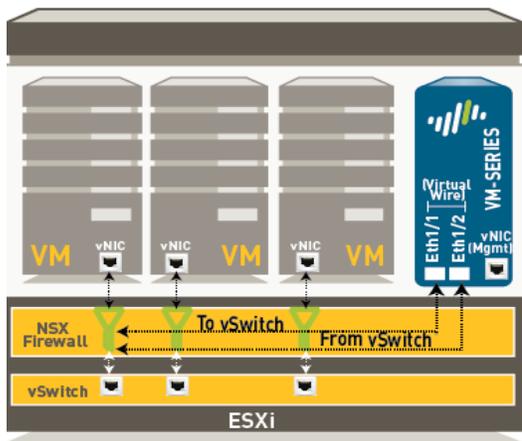
To meet the security challenges in the software-defined datacenter, the NSX Manager, ESXi servers and Panorama work harmoniously to automate the deployment of the VM-Series firewall.



1. Register the Palo Alto Networks NGFW service—The first step is to register the Palo Alto Networks NGFW as a service on the NSX Manager. The registration process uses the NetX management plane API to enable bi-directional communication between Panorama and the NSX Manager. Panorama is configured with the IP address and access credentials to initiate a connection and register the Palo Alto Networks NGFW service on the NSX Manager. The configuration includes the URL for accessing the VM-Series base image that is required to deploy the VM-Series NSX edition firewall, the authorization code for retrieving the license and the device group to which the VM-Series firewalls will belong. The NSX manager uses this management plane connection to share updates on the changes in the virtual environment with Panorama.

2. Deploy the VM-Series automatically from NSX—The NSX Manager collects the VM-Series base image from the URL specified during registration and installs an instance of the VM-Series firewall on each ESXi host in the ESXi cluster. From a static management IP pool (that you define on the NSX Manager), a management IP address is assigned to the VM-Series firewall and the Panorama IP address is provided to the firewall. When the firewall boots up, the NetX dataplane integration API connects the VM-Series firewall to the hypervisor so that it can receive traffic from the vSwitch.

Traffic Flow on the VM-Series NSX Edition



3. Establish communication between the VM-Series firewall and Panorama: The VM-Series firewall then initiates a connection to Panorama to obtain its license. Panorama retrieves the license from the update server and pushes it to the firewall. The VM-Series firewall receives the license (VM-1000-HV) and reboots with a valid serial number.

4. Install configuration/policy from Panorama to the VM-Series firewall: The VM-Series firewall reconnects with Panorama and provides its serial number. Panorama now adds the firewall to the device group that was defined in the registration process and pushes the default policy to the firewall. The VM-Series firewall is now available as a security virtual machine

that can be further configured to safely enable applications on the network.

5. Push traffic redirection rules from NSX Firewall: On the Service Composer on the NSX Firewall, create security groups and define network introspection rules that specify traffic from which guests are steered to the VM-Series firewall. See [Integrated Policy Rules](#) for details.

6. Receive real-time updates from NSX Manager: The NSX Manager sends real-time updates on the changes in the virtual environment to Panorama. These updates include information on the security groups and IP addresses of guests that are part of the security group from which traffic is redirected to the VM-Series firewall. See [Integrated Policy Rules](#) for details.

7. Use Dynamic Address Groups in policy and push dynamic updates from Panorama to the VM-Series firewalls: On Panorama, use the real-time updates on security groups to create Dynamic Address Groups, bind them to security policies and then push these policies to the VM-Series firewalls. Every VM-Series firewall in the device group will have the same set of policies and is now completely marshaled to secure the SDDC. See [Policy Enforcement using Dynamic Address Groups](#) for details.

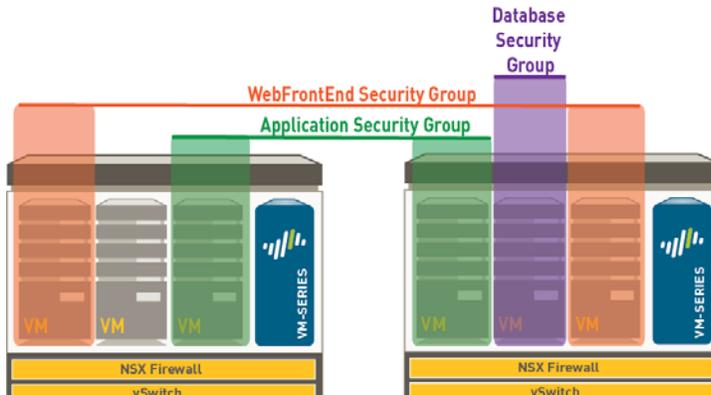
Integrated Policy Rules

The NSX Firewall and the VM-Series firewall work in concert to enforce security; each provides a set of traffic management rules that are applied to the traffic on each ESXi host. The first set of rules is defined on the NSX Firewall; these rules determine traffic from which guests in the cluster are steered to the VM-Series firewall. The second set of rules (next-generation firewall rules) is defined on Panorama and pushed to the VM-Series firewalls. These are security enforcement rules for the traffic that is steered to the Palo Alto Networks NGFW service. These rules determine how the VM-Series firewall must process, that is allow, deny, inspect, and constrain, the application for enabling it safely on your network.

Rules defined on the NSX Firewall—The rules for directing traffic from the guests on each ESXi host are configured on the NSX Manager. The Service Composer on the NSX Manager allows you to define what kind of security protection, such as firewall rules are to be applied to the guests in the ESXi cluster. To define the rules on the NSX Firewall, you must first aggregate the guests into security groups, and then create NSX service composer policies to redirect the traffic from these security groups to the Palo Alto Networks NGFW service on the NSX Firewall.

- The following diagram illustrates how security groups can be composed of guests across different ESXi hosts within a cluster.

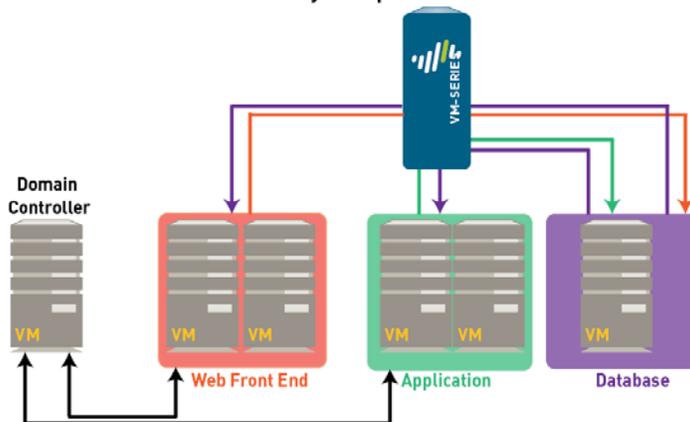
Grouping Guests into Security Groups Within a Cluster



For traffic that needs to be inspected and secured by the VM-Series firewall, the NSX service composer policies redirect the traffic to the Palo Alto Networks NGFW service. This traffic is then steered to the VM-Series firewall and is first processed by the VM-Series firewall before it goes to the virtual switch.

Traffic that does not need to be inspected by the VM-Series firewall, for example network data backup or traffic to an internal domain controller, does not need to be redirected to the VM-Series firewall and can be sent to the virtual switch for onward processing.

Traffic Between Security Groups Redirected to the Palo Alto NGFW



Traffic that is not Redirected to the Palo Alto NGFW Service

- Rules centrally managed on Panorama and applied by the VM-Series firewall**—The next-generation firewall rules are applied by the VM-Series firewall. These rules are centrally defined and managed on Panorama using templates and device groups and pushed to the VM-Series firewalls. The VM-Series firewall then enforces security policy by matching on source or destination IP address—the use of Dynamic Address Groups allows the firewall to populate the members of the Dynamic Address Groups in real time—and forwards the traffic to the filters on the NSX Firewall.

To understand how the NSX Manager and Panorama stay synchronized with the changes in the SDDC and ensure that the VM-Series firewall consistently enforces policy, see [Policy Enforcement using Dynamic Address Groups](#).

Policy Enforcement using Dynamic Address Groups

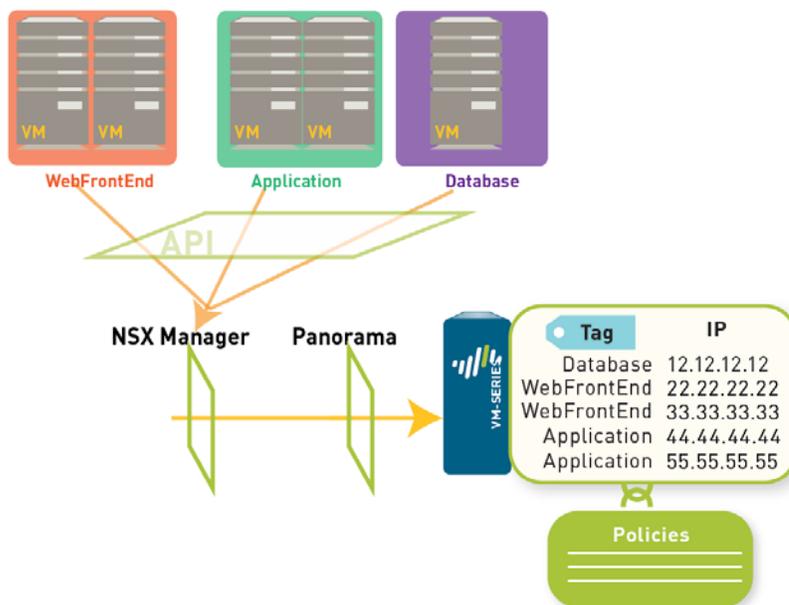
Unlike the other versions of the VM-Series firewall, the NSX edition does not use security zones as the primary traffic segmentation mechanism because both virtual wire interfaces belong to the same zone. Instead, the NSX edition uses Dynamic Address Groups to segment traffic.

A Dynamic Address Group is used as a source or destination object in security policy. Because IP addresses are constantly changing in a datacenter environment, Dynamic Address Groups offer a way to automate the process of referencing source and/or destination addresses within security policies. Unlike static address objects that must be manually updated in configuration and committed whenever there is an address change (addition, deletion, or move), Dynamic Address Groups automatically adapt to changes.

All security groups defined on the NSX Manager are automatically provided as updates to Panorama using the NetX API management plane integration and can be used as filter criteria to create Dynamic Address Groups; the firewall filters for the name of the security group, which is a tag, to find all the members that belong to a security group.

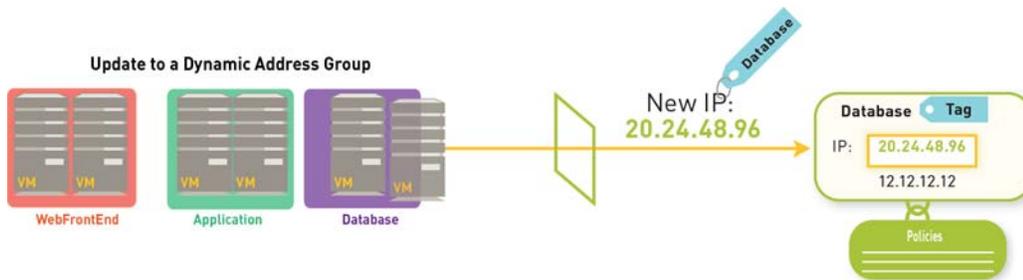
If, for example, you have a multi-tier architecture for web applications, on the NSX Manager you create three security groups for the WebFrontEnd servers, Application servers and the Database servers. The NSX Manager updates Panorama with the name of the security groups and the IP address of the guests that are included in each security group.

Dynamic Address Groups Match on Security Group Names



On Panorama, you can then create three Dynamic Address Groups to match for objects that are tagged as Database, Application and WebFrontEnd. Then, in security policy you can use the Dynamic Address Groups as source or destination objects, define the applications that are permitted to traverse across these servers, and push the rules to the VM-Series firewalls.

Each time a guest is added or modified in the ESXi cluster or a security group is updated or created, the NSX Manager uses the REST-based XML API to update Panorama with the IP address, and the security group to which a guest belongs.



To ensure the name of each security group is unique, the vCenter server assigns a Managed Object Reference (MOB) ID to the name you define for the security group. The syntax to display the name of a security group on Panorama is `specified_name-securitygroup-number`; for example, `WebFrontEnd-securitygroup-47`.

When Panorama receives the API notification, it verifies/updates the IP address of each guest and the security group to which that guest belongs. Then, Panorama pushes these real-time updates to all the firewalls that are included in the device group and notifies device groups in the service manager configuration on Panorama.

On each firewall, all policy rules that reference these Dynamic Address Groups are updated at runtime. Because the firewall matches on the security group tag to determine the members of a Dynamic Address Group, you do not need to modify or update the policy when you make changes in the virtual environment. The firewall matches the tags to find the current members of each Dynamic Address Group and applies the security policy to the source/destination IP address that are included in the group.

What are the Benefits of the Solution?

The NSX edition of the VM-Series firewall is focused on securing east-west communication in the software-defined datacenter. Deploying the firewall has the following benefits:

- Simpler Deployment**—The NSX Manager automates the process of delivering next-generation firewall security services, and the VM-Series firewall allows for transparent security enforcement. When a new ESXi host is added to a cluster, a new VM-Series firewall is automatically deployed, provisioned and available for immediate policy enforcement without any manual intervention. The automated workflow allows you to keep pace with the virtual machine deployments in your datacenter. The hypervisor mode on the firewall removes the need to reconfigure the ports/ vswitches/ network topology; because each ESXi host has an instance of the firewall, the traffic does not need to traverse the network or be back hauled for inspection and consistent enforcement of policies.
- Tighter Integration between Virtual Environment and Security Enforcement for Dynamic Security**—Dynamic Address Groups maintain awareness of changes in the virtual machines/applications and ensure that security policy stays in tandem with the changes in the network. The awareness provides visibility and protection of applications in an agile environment.
- Sturdier Centralized Management**—The firewalls deployed using this solution are licensed and managed by Panorama, the Palo Alto Networks central management tool. Using Panorama to manage both the perimeter and datacenter firewalls (the hardware-based and virtual firewalls) allows you to centralize policy management and maintain agility and consistency in policy enforcement throughout the network.

In summary, this solution ensures that the dynamic nature of the virtual network is secured with minimal administrative overhead. You can successfully deploy applications with greater speed, efficiency, and security.

Deploy the VM-Series NSX Edition Firewall

To deploy the NSX edition of the VM-Series firewall, use the following workflow:

- ❑ **Step 1: Set up the Components**—To deploy the VM-Series NSX edition, set up the following components:
 - Set up the vCenter server, install and register the NSX Manager with the vCenter server.
If you have not already set up the virtual switch(es) and grouped the ESXi hosts in to clusters, refer to the VMware documentation for instructions on setting up the vSphere environment. This document does not take you through the process of setting up the VMware components of this solution.
 - Upgrade Panorama (or install, if needed) to version 6.0. Create a Device Group and Template on Panorama. If you are new to Panorama, refer to the [Panorama Administrator's Guide](#) for instructions on setting up Panorama.
 - Download and save the ovf template for the NSX edition of the VM-Series firewall on a web server. The NSX Manager must have network access to this web server so that it can deploy the VM-Series firewall as needed. You cannot host the ovf template on Panorama.
- ❑ **Step 2: Register**—Configure Panorama to register the VM-Series firewall as a service on the NSX Manager. When registered, the VM-Series firewall is added to the list of network services that can be transparently deployed as a service by the NSX Manager.
The connection between Panorama and the NSX Manager is also required for licensing and configuring the firewall.
- ❑ **Step 3: Deploy the Firewall and Create Policies** —Install the VM-Series firewall and create policies to redirect traffic to the VM-Series firewall and to secure the traffic that is redirected to the firewall.
 - (On the NSX Manager) Define the IP address pool. An IP address from the defined range is assigned to the management interface of each instance of the VM-Series firewall.
 - (On the NSX Manager) Deploy the VM-Series firewall. The NSX Manager automatically deploys an instance of the VM-300-HV on each ESXi host in the cluster.
 - (On the NSX Manager) Set up the service composer and create security groups. A security group assembles the specified guests/applications so that you can apply policy to the group.
 - (On Panorama) Apply policies to the VM-Series firewall. From Panorama, you define, push, and administer policies centrally on all the VM-Series firewalls. On Panorama, create Dynamic Address Groups for each security group and reference the Dynamic Address Groups in policy, and then push the policies to the managed firewalls.
This centralized administration mechanism allows you to secure guests/applications with minimal administrative intervention.
 - (On the NSX Manager) Define the network introspection rules that redirect traffic to the VM-Series firewall.
- ❑ **Step 4: Monitor and Maintain Network Security**—Panorama provides a comprehensive, graphical view of network traffic. Using the visibility tools on Panorama—the Application Command Center (ACC), logs, and the report generation capabilities—you can centrally analyze, investigate and report on all network activity, identify areas with potential security impact, and translate them into secure application enablement policies. Refer to the [Panorama Administrator's Guide](#) for more information.

- **Step 5: Upgrade the software version**— When upgrading the VM-Series NSX edition firewalls, you must first upgrade Panorama before upgrading the firewalls. To upgrade the firewalls, see [Upgrade the PAN-OS Software Version \(NSX Edition\)](#).

For step-by-step instructions for the tasks in this workflow, see [Set Up a VM-Series NSX Edition Firewall](#).