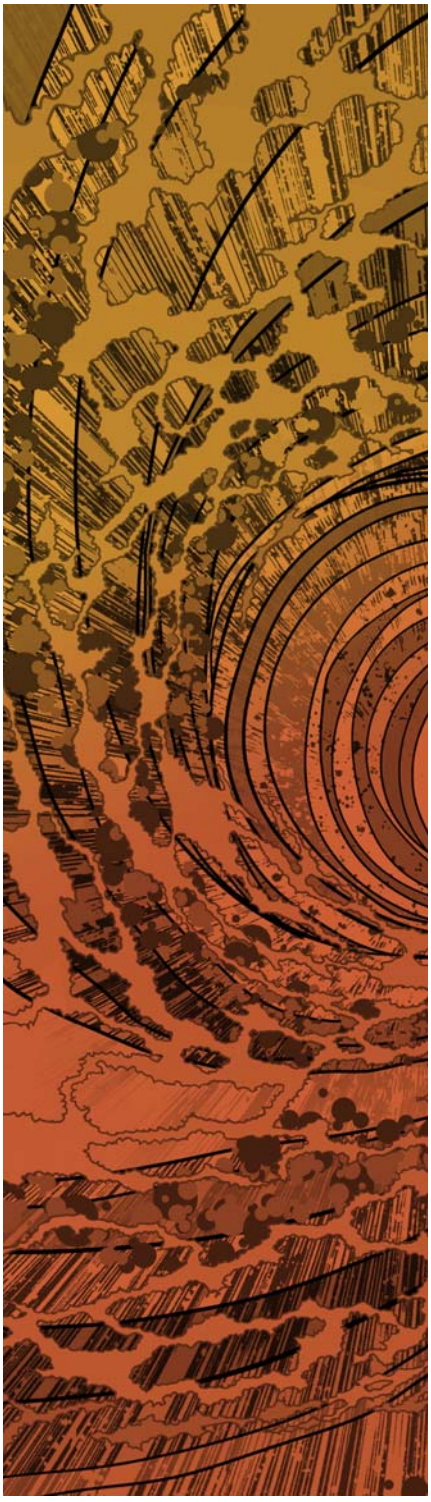

paloalto NETWORKS®

PAN-OS® New Features Guide

Version 6.1

Contact Information

Corporate Headquarters:

Palo Alto Networks

4401 Great America Parkway

Santa Clara, CA 95054

https://www.paloaltonetworks.com/company/contact-us

About this Guide

This guide describes how to use the new features introduced in PAN-OS 6.1. For additional information, refer to the following resources:

- For information on the additional capabilities and for instructions on configuring the features on the firewall, refer to https://www.paloaltonetworks.com/documentation.
- For access to the knowledge base and community forums, refer to https://live.paloaltonetworks.com.
- For contacting support, for information on support programs, to manage your account or devices, or to open a support case, refer to https://www.paloaltonetworks.com/support/tabs/overview.html.
- For the most current PAN-OS and Panorama 6.1 release notes, go to https://www.paloaltonetworks.com/documentation/61/pan-os/pan-os-release-notes.html.

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

Palo Alto Networks, Inc.
www.paloaltonetworks.com

Revision Date: July 11, 2016

Table of Contents

Upgrade Your Firewalls to PAN-OS 6.1

Upgrade/Downgrade Considerations

Table: PAN-OS 6.1 Upgrade/Downgrade Considerations lists the new features that have upgrade and/or downgrade impact. Make sure you understand the changes that will occur in the configuration prior to upgrading to or downgrading from PAN-OS 6.1. For additional information about this release, refer to the Release Notes.

Table: PAN-OS 6.1 Upgrade/Downgrade Considerations

Feature	Upgrade Considerations	Downgrade Considerations
Configurable Key Size for SSL Forward Proxy Server Certificates	• The default key size for SSL/TLS Forward Proxy Server certificates changes from 1024-bit RSA to **Defined by destination host**.	• The default key size for the SSL/TLS Forward Proxy Server certificates changes from **Defined by destination host** to 1024-bit RSA.
LACP		• Before downgrading, you must disable LACP for any aggregate group that uses it. PAN-OS retains all other aggregate group and interface settings.
Security Policy Rulebase Enhancements	• A new **Rule Type** classification indicates whether a security rule matches intrazone traffic, interzone traffic, or both (called *universal*). • All existing rules in the rulebase are converted to universal rules. • Default rules are displayed at the end of the security rulebase. By default, the treatment of traffic that does not match any rule in the rulebase is unchanged: intrazone traffic is allowed and interzone traffic is denied. However, you can now override this default behavior.	• The **Rule Type** is removed from all rules and all **intrazone** and **interzone** rules you have defined will be removed; all rules of type **universal** will be preserved. If you do not want a particular **intrazone** or **interzone** rule to be removed, you must change the **Rule Type** to **universal** before downgrading. • Default rules are no longer displayed at the end of the security rulebase. If you modified the default rules prior to downgrade, traffic that does not match any other security rule will revert to the default: intrazone traffic is allowed and interzone traffic is denied.
Session End Reason Logging	• Logs of subtype end, generated in an older PAN-OS release, will have a session end reason of unknown.	• PAN-OS ignores the session end reasons for logs generated in PAN-OS 6.1 or later. Logs generated in earlier releases do not have session end reasons.
TCP Session Closing Timers	• The TCP-Wait timer setting becomes the setting of the TCP Half Closed timer.	• The TCP-Wait timer takes the setting of the TCP Half Closed timer. The TCP Time Wait and TCP Unverified RST timers go away.

Feature	Upgrade Considerations	Downgrade Considerations
Upgrade the WF-500 Appliance and Enable Windows 7 64-bit Support		• When you upgrade the WF-500 WildFire appliance to version 6.1, you must first install the WFWin7_64 base image and the WFWin7_64 add-on image before upgrading the operating system. • If you downgrade a WF-500 appliance with operating system 6.1 installed and the appliance is configured to use the Windows 7 64-bit virtual machine (VM) sandbox, you must configure the appliance to use a different VM after the downgrade. This is because the Windows 7 64-bit VM is only available in version 6.1 or above.

Upgrade to PAN-OS 6.1

How you upgrade to PAN-OS 6.1 depends on whether you have standalone firewalls or firewalls in a high availability (HA) configuration and whether your firewalls are managed by Panorama. Review the Release Notes and then follow the procedure that matches your configuration:

- Upgrade Firewalls Using Panorama
- Upgrade the Firewall to PAN-OS 6.1
- Upgrade an HA Firewall Pair to PAN-OS 6.1

When upgrading firewalls that you manage with Panorama or firewalls that are configured to forward content to a WF-500 appliance, you must first upgrade Panorama and/or the WF-500 appliance first, before upgrading the firewalls.

Upgrade Firewalls Using Panorama

Review the Release Notes and then use the following procedure to upgrade firewalls that Panorama manages. This procedure applies to standalone firewalls and firewalls configured in a high availability (HA) configuration.

Upgrade Firewalls Using Panorama	
Step 1 Save a backup of the current configuration file on each managed firewall that you plan to upgrade. 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.	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.
Step 2 Install the content updates. Make sure the firewalls you plan to upgrade are running content release version 454 or later.	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.  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)		
Step 3	Determine the upgrade path. 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 6.1.1, you must: • Download and install PAN-OS 5.0.0 and reboot. • Download and install PAN-OS 6.0.0 and reboot. • Download the PAN-OS 6.1.0 image (you do not need to install it). • Download and install PAN-OS 6.1.1 and reboot.	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: • If PAN-OS 6.0.0 or later is currently installed, continue to Step 4. • If a version of PAN-OS prior to 6.0.0 is currently installed, follow the upgrade path to 6.0.0 before upgrading to 6.1. Refer to the Release Notes for your currently installed PAN-OS version for upgrade instructions.
Step 4	Download the software updates.	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 6.1.1, download the images that have File Name PanOS_3000-6.1.1 and PanOS_5000-6.1.1. 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)

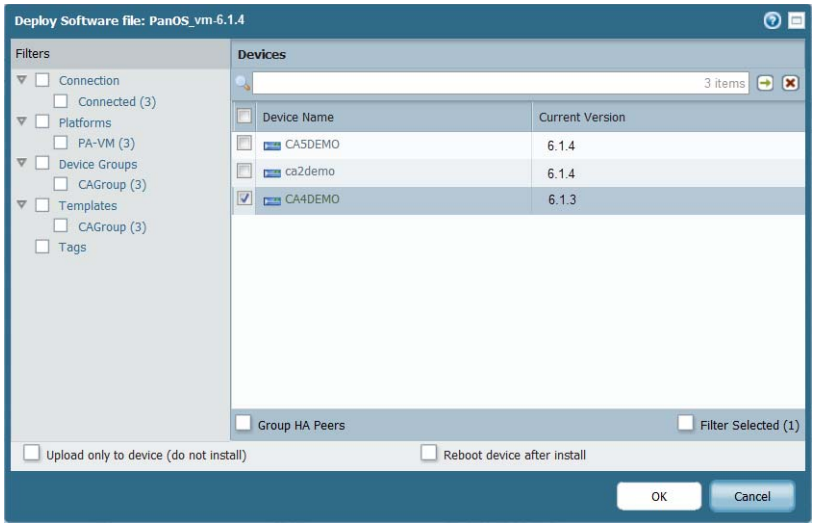

Step 6	Verify the software and content release version running on each managed device.	1. On Panorama, select **Panorama > Managed Devices**. 2. Locate the firewalls and review the content and software versions in the table. 

Upgrade the Firewall to PAN-OS 6.1

Review the Release Notes and then use the following procedure to upgrade a firewall that is not in an HA configuration to PAN-OS 6.1.

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

Upgrade PAN-OS

Step 1	Save a backup of the current configuration file. 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.	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.

Upgrade PAN-OS (Continued)		
Step 2	Make sure the firewall is running Content Release version 454 or later.	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**.
Step 3	Determine the upgrade path. 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. For example, if you want to upgrade from PAN-OS 4.1.8 to PAN-OS 6.1.1, you must: • Download and install PAN-OS 5.0.0 and reboot. • Download and install PAN-OS 6.0.0 and reboot. • Download the PAN-OS 6.1.0 image (you do not need to install it). • Download and install PAN-OS 6.1.1 and reboot.	1. Select **Device > Software**. 2. Check which version has a check mark in the Currently Installed column and proceed as follows: • If PAN-OS 6.0.0 or later is currently installed, continue to Step 4. • If a version of PAN-OS prior to 6.0.0 is currently installed (as shown here), follow the upgrade path to 6.0.0 before you can upgrade to 6.1. Refer to the release notes for your currently installed PAN-OS version for upgrade instructions. 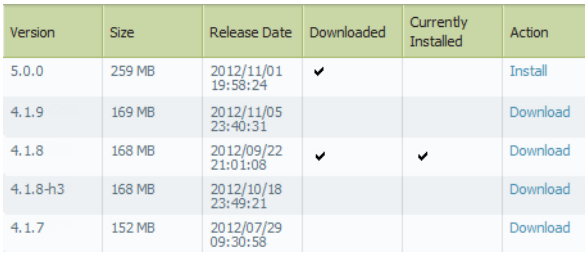
Step 4	Install PAN-OS 6.1. If your firewall does not have Internet access from the management port, you can download the software update from the Palo Alto Networks Support Site (https://support.paloaltonetworks.com). You can then manually **Upload** it to your firewall.	1. Click **Check Now** to check 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: • 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.
Step 5	Verify that the firewall is passing traffic.	Select **Monitor > Session Browser**.

Upgrade an HA Firewall Pair to PAN-OS 6.1

Review the Release Notes and then use the following procedure to upgrade a pair of firewalls in an HA configuration. This procedure applies to both Active/Passive and Active/Active configurations.

Upgrading PAN-OS on firewalls in a high availability (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. If you have session synchronization enabled, this will continue to function during the upgrade process as long as you are upgrading from PAN-OS 6.0.x to PAN-OS 6.1. If you are upgrading the pair from an older feature release of PAN-OS, session syncing between the firewalls will not work and, if a failover occurs before both firewalls are running the same version of PAN-OS, session forwarding could be impacted. In this case, if session continuity is required, you must temporarily permit non-syn-tcp while the session table is rebuilt.

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		
Step 1	Save a backup of the current configuration file.  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.	Perform these steps on each firewall in the pair: 1. Select **Device > Setup > Operations** and click **Export named configuration snapshot**. 2. Select the XML file that contains your running configuration (for example, **running-config.xml**) and click **OK** to export the configuration file. 3. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the upgrade.
Step 2	Make sure each device running Content Release version 454 or later.	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)

Step 3 Determine the upgrade path.

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.

For example, if you want to upgrade from PAN-OS 4.1.8 to PAN-OS 6.1.1, you must:

- Download and install PAN-OS 5.0.0 and reboot.
- Download and install PAN-OS 6.0.0 and reboot.
- Download the PAN-OS 6.1.0 image (you do not need to install it).
- Download and install PAN-OS 6.1.1 and reboot.

1. Select **Device > Software**.
2. Check which version has a check mark in the Currently Installed column and proceed as follows:
   - If PAN-OS 6.0.0 or later is currently installed, continue to Step 4.
   - If a version of PAN-OS prior to 6.0.0 is currently installed (as shown here), follow the upgrade path to 6.0.0 before you can upgrade to 6.1. Refer to the release notes for your currently installed PAN-OS version for upgrade instructions.

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

Step 4 Install PAN-OS 6.1 on the passive device (active/passive) or on the active-secondary device (active/active).

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.

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:
   - 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.

Step 5 Suspend the active/active-primary firewall.

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. (Optional) If you have session synchronization enabled and you did not upgrade from PAN-OS 6.0.x, run the operational command **`set session tcp-reject-non-syn no`**. This will rebuild the session table so that sessions that started prior to the upgrade will continue.

Upgrade PAN-OS (Continued)		
Step 6	Install PAN-OS 6.1 on the other device in the pair.  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.	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: • 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. 5. (Optional) 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 4, revert back by running the **`set session tcp-reject-non-syn yes`** command.  If the preemptive option is configured, the current passive device will revert to active when state synchronization is complete.
Step 7	Verify that the devices are passing traffic as expected. In an active/passive deployment, the active device should be passing traffic and in an active/active deployment both devices should be passing traffic.	Run the following CLI commands to confirm that the upgrade succeeded: • (Active device(s) only) To verify that the active devices are passing traffic, run **`show session all`**. • To verify session synchronization, run **`show high-availability interface ha2`** and make sure that the Hardware Interface counters on the CPU table are increasing as follows: • In an active/passive configuration, only the active device will show packets transmitted and the passive device will only show packets received.  If you have enabled HA2 keep-alive, the hardware interface counters on the passive peer will show both transmit and receive packets. This occurs because HA2 keep-alive is bidirectional which means that both peers transmit HA2 keep-alive packets. • In the active/active configuration, you will see packets received and packets transmitted on both devices.

Downgrade from PAN-OS 6.1

The way you downgrade from PAN-OS 6.1 depends on whether you are downgrading to a previous feature release (where the first or second digit in the PAN-OS version changes, for example 6.1 to 6.0 or 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 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	
Step 1 Save a backup of the current configuration file.  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.	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.

Downgrade to a Previous Maintenance Release	
Step 2 Install the previous maintenance release image.  If your firewall does not have Internet access from the management port, you can download the software update from the Palo Alto Networks Support Portal. You can then manually **Upload** it to your firewall.	1. 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: • 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 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 6.1 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, on downgrade 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, thereby resulting in a loss of logs.

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

Downgrade to a Previous Feature Release	
Step 1 Save a backup of the current configuration file.  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.	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.

Downgrade to a Previous Feature Release	
Step 2 Install the previous feature release image.  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.	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 6.1.1 and want to downgrade to PAN-OS 6.0.3, select autosave-6.0.3. 5. After the install completes, reboot using one of the following methods: • If you are prompted to reboot, click **Yes**. • If you are not prompted to reboot, select **Device > Setup > Operations** and click **Reboot Device** in the Device Operations section.

Management Features

Authenticated NTP

The firewall can use a Network Time Protocol (NTP) server to synchronize the firewall's local clock. You can now configure the firewall to authenticate time updates from the configured NTP server. Authenticated NTP allows you to prevent tampering with the firewall's clock and resulting disruptions to logging and schedule-based policies and services. To enable Authenticated NTP, first designate an NTP server and then select a type of authentication for the firewall to use to authenticate the time updates.

Configure Authenticated NTP		
Step 1	Add an NTP server that the firewall will use to synchronize its local clock.	1. Select **Device > Setup > Services** and edit the Services. 2. Select the **NTP** tab. 3. In the **NTP Server Address** field, enter the IP address or hostname of an NTP server that you want to use to synchronize the local clock of the firewall. Optionally enter the IP address or hostname of a secondary NTP server with which to synchronize the clock. The firewall does not prioritize the primary or secondary; it can synchronize with either NTP server.

Step 2	Choose the type of authentication to use to authenticate time updates from the NTP server.	Continuing on the **NTP** tab, select the **Authentication Type** for the firewall to use to authenticate time updates from each NTP server:
		• **None**—(Default) Select this option to disable NTP Authentication.
		• **Symmetric Key**—Select this option for the firewall to use symmetric key exchange (shared secrets) to authenticate the NTP server's time updates. If you select Symmetric Key, continue by entering the following fields:
		  • **Key ID**—Enter the Key ID (1- 65534).
		  • **Algorithm**—Select the Algorithm to use in NTP authentication (MD5 or SHA1).
		  • **Authentication Key/Confirm Authentication Key**—Enter and confirm the authentication algorithm's authentication key.
		• **Autokey**—Select this option for the firewall to use autokey (public key cryptography) to authenticate the NTP server's time updates.

Step 3	Save your settings.	Click **OK** and **Commit**.

App Scope Enhancements

The App Scope uses charts and maps to graphically represent network activity. It presents a summary of the change in traffic patterns over a period of time and allows you to detect anomalies. For the rendering of charts and maps in App Scope, the firewall now uses HTML5 and JavaScript instead of Adobe Flash. With this change, the following enhancements are available in **Monitor > App Scope**:

- Ability to toggle the attributes in the legend to only view chart details that you want to review. The ability to include or exclude data from the chart allows you to change the scale and review details more closely.

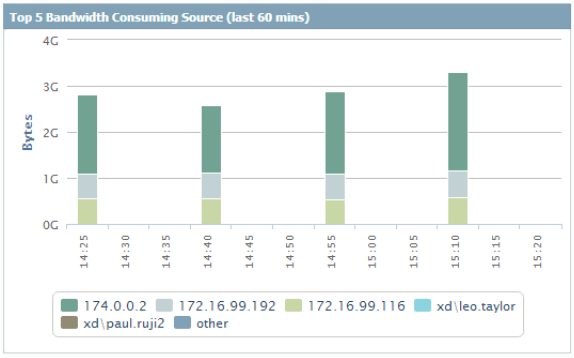

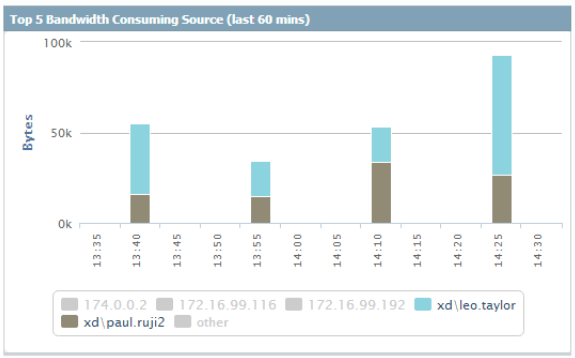

- Ability to click into an attribute in a bar chart and drill down to the related sessions in the ACC. Click into an Application name, Application Category, Threat Name, Threat Category, Source IP address or Destination IP address on any bar chart to filter on the attribute and view the related sessions in the ACC.
- Ability to export a chart or map to PDF or as an image. For portability and offline viewing, you can **Export** charts and maps as PDFs or PNG images.

The firewall uses geolocation for creating traffic maps and threat maps. Now, the firewall is placed at the bottom of the threat and traffic maps screen if the geolocation coordinates are not specified on the firewall.

Security Policy Rulebase Enhancements

The Security policy rulebase has been enhanced to enable more streamlined control over intrazone (within a zone) and interzone (between zones) traffic, allowing you to create rules that enable visibility and control over intrazone or interzone traffic for multiple zone pairs in a single rule rather than having to create separate rules for each pair. To enable this flexibility, a new *Rule Type* classification indicates whether the rule matches intrazone traffic, interzone traffic, or both (called *universal*).

In addition, the rules that the firewall uses for handling intrazone and interzone traffic that doesn't match any other rule have now been exposed, allowing you override select settings on these default rules, such as enabling logging so that you can see what traffic isn't matching any of the rules you have defined.

The following sections describe these security rulebase enhancements and how to use them to simplify your security policy rulebase:

- Use the New Rule Types in Policy
- Modify the Default Rules

Use the New Rule Types in Policy

Each security policy rule now includes a *Rule Type* that specifies whether the rule matches intrazone traffic, interzone traffic, or both intrazone and interzone traffic (called *universal* traffic). Because you can now create a single rule that matches either intrazone traffic or interzone traffic within one or more source zones (intrazone) or source-destination pairs (interzone), the total number of rules you now need is reduced. For example, suppose you want to log all traffic between zones A, B, and C, but not traffic within the zones. Previously you would have had to create a separate rule for each source-destination zone pair (in both directions) and enable logging on each individual rule, but now you can instead create a single *interzone* rule matching multiple zone pairs and enable logging on that one rule. Table: Security Policy Rule Types describes the three new rule types and how they match traffic.

Table: Security Policy Rule Types

Rule Type	Description
universal (default)	Matches both intrazone traffic and interzone traffic in the specified source and destination zones. For example, a universal rule with source zones A and B and destination zones A and B would match traffic: • Within zone A • Within zone B • From zone A to zone B • From zone B to zone A This is the default Rule Type when creating a new security rule. Upon upgrade to PAN-OS 6.1, all existing rules in your security rulebase are converted to universal rules.
intrazone	Matches traffic within the specified source zones, but not between them (you cannot specify a destination zone for intrazone rules). For example, an intrazone rule with source zones A and B would match traffic: • Within zone A • Within zone B It would NOT match traffic between zones A and B.

Rule Type	Description
interzone	Matches traffic between the specified source and destination zones and each rule can match on multiple source and destination zone pairs. For example, an interzone rule with source zones A, B, and C and destination zones A and B would match traffic: • From zone A to zone B • From zone B to zone A • From zone C to zone A • From zone C to zone B This rule would NOT match traffic within zone A, within zone B, or within zone C.

Upon upgrade to PAN-OS 6.1, the firewall converts all existing security rules to **universal** rules. You can add new rules or modify the existing rules as follows:

Set the Rule Type in a Security Policy Rule		
Step 1	Select or add a security rule.	1. Select **Policies > Security** on the firewall or **Policies > Security > Default Rules** on Panorama. 2. Select the location where you want to make the default rule modifications: • On a firewall enabled for multiple virtual systems, select the **Virtual System** on which to make the changes (this step is not required on firewalls that do not have multiple virtual systems). • On Panorama, select a specific **Device Group** on which to make the changes, or select **Shared**. 3. Select a rule **Name** to modify or **Add** a new rule.
Step 2	Define the rule type and make sure zone settings are appropriate. Upon upgrade to PAN-OS 6.1, the firewall converts all existing security rules to **universal** rules. If you subsequently downgrade to an earlier PAN-OS version, the **Rule Type** is removed from all rules and all **intrazone** and **interzone** rules you have defined will be removed.	1. On the **General** tab, select a value from the **Rule Type** drop-down. 2. On the **Source** and **Destination** tabs, set the **Source Zone** and **Destination Zone** as appropriate for the selected rule type as detailed in Table: Security Policy Rule Types. Keep in mind that security rules of type **intrazone** only allow selection of one or more Source Zones; you cannot select a Destination Zone. Selecting multiple Source Zones in an **intrazone** rule controls matching traffic within each of the specified zones, but not between them. 3. Complete any additional security policy rule settings as necessary.
Step 3	Commit the changes.	Click **Commit**.

Modify the Default Rules

The security rulebase includes default rules that instruct the firewall how to handle traffic that does not match any other rule. Per these default rules, the firewall allows all intrazone traffic and denies all interzone traffic that doesn't match another rule. These default rules— named *interzone-default* and *intrazone-default*— are now exposed within the Security policy rulebase. Because these rules are intended to only match traffic

that does not match any other rule, they appear at the end of the rulebase (and are numbered accordingly). On Panorama, the default rules come after all Pre Rules and Post Rules, ensuring that the evaluation order for rules on Panorama is top down in the following order: Shared Pre Rules, Device Group Pre Rules, Local rules defined on the firewall, Device Group Post Rules, Shared Post Rules, Default Rules.

Although the rules are initially read-only (either because they are part of the predefined configuration settings or because they were pushed from Panorama), you can override some of the rule settings, enabling more granular control over the handling of traffic that doesn't match any rule. For example, you can enable logging on the default rules for visibility into the traffic that is not matching any of the security policy rules you have created. Or, you can attach security profiles to scan the traffic for threats.

Because the default rules are part of the predefined configuration on the firewall and on Panorama, you cannot see them when you run the `show running security-policy` CLI command until you override them.

Only the following settings on the **General** and **Action** tabs of a default rule are available for modification:

Field	Description
Tags	Attach tags to help visually identify the default rule.
Action	Change whether traffic matching the default rules is blocked or allowed.
Profile Setting	Attach security profiles (or a security profile group) to apply to traffic matching the rules.
Log Setting	Select a logging option to identify when (or if) to log of matching traffic, or select a Log Forwarding profile to trigger when there is a traffic match.

The default rules are not included in the Security policy rulebase filtering and therefore are always displayed on the rulebase even when you apply filters that do not match them.

To modify the default rules, you must first override them. The process is similar to how you would override a template setting that was pushed from Panorama to a firewall. On the firewall, you can modify default rules that are part of the predefined configuration, or pushed from a Panorama Shared or Device Group context; the default rules are virtual system (VSYS) specific. On Panorama, you can override the default rules that are part of the predefined configuration; you can modify them in a Device Group or Shared context.

Modify the Default Security Rules

Step 1 Override the predefined rule you want to modify so that it is no longer read-only.

If you want to see where the current default rule settings are coming from before overriding them, hover over the icon.

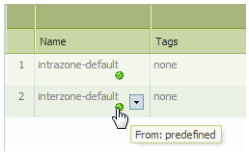

1. Select **Policies > Security** on the firewall or **Policies > Security > Default Rules** on Panorama.
2. Select the location where you want to make the default rule modifications:
   - On a firewall enabled for multiple virtual systems, select the **Virtual System** on which to make the changes (this step is not required on firewalls that do not have multiple virtual systems).
   - On Panorama, select a specific **Device Group** on which to make the changes, or select **Shared**.
3. Select the default rule you want to modify and click **Override**.

Step 2 Edit the default rule.

If you override a predefined rule and then want to set it back to its default, select the rule and click **Revert**.

Upon downgrade from PAN-OS 6.1 to an earlier version, the firewall reverts any changes you have made to the default rules to their predefined settings and removes them from the security rulebase. In this state, the firewall will allow all intrazone traffic and deny all interzone traffic that doesn't match any other rule.

1. Modify any of the following fields:
   - On the **General** tab, select any **Tags** you want to attach to the rule.
   - On the **Actions** tab, modify any of the following:
     - Change the **Action** (**Deny** or **Allow**)
     - Add a security profile by setting the **Profile Type** field to **Profiles** and then selecting the appropriate profiles from the corresponding drop-downs.
     - Add a security profile group by setting the **Profile Type** field to **Group** and selecting the **Group Profile**.
     - Change when to trigger a log for traffic matching the rule (**Log at Session Start** or **Log at Session End**).
     - Enable log forwarding for logs triggered by this rule by selecting or creating a **Log Forwarding** profile.
2. Click **OK** to save the changes. Notice that the rule now shows the override icon:

   Hover over the icon to see the location you are overriding from (it will show **Predefined** if you are overriding the default firewall or Panorama configuration or **Panorama** if you are overriding a default rule pushed to a firewall from a Panorama Device Group or Shared context:

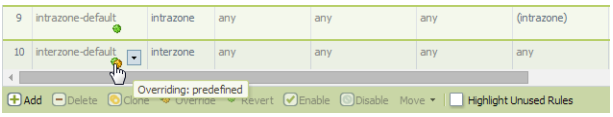

Step 3 Commit the changes.

Click **Commit**.

Multiple M-100 Appliance Interfaces

The Panorama M-100 appliance now supports separate interfaces for configuration (of firewalls, log collectors, and Panorama itself), log collection, and communication within collector groups. For the configuration function, you must assign the MGT (Eth0) interface. For log collection and collector group communication, you can assign the Eth1 or Eth2 interface to perform either or both functions. For example, you can configure Eth1 for log collection and configure Eth2 for collector group communication, or vice versa. However, you cannot assign multiple interfaces to a single function. For example, you cannot configure both Eth1 and Eth2 to perform log collection. By default, the M-100 appliance uses one interface (MGT) for all three functions but configuring separate interfaces is a best practice to improve:

- Security—For the management interface, define a separate subnet that is more private than the subnets used for log collection and collector group communication.
- Control traffic prioritization—Reserve the management interface for control traffic so it does not compete for bandwidth with log collection traffic.
- Performance and scalability—Avoid imposing the traffic load of all three functions on a single interface.

If the Eth1 or Eth2 interface goes down, Panorama does not fail back to MGT for the functions that used those interfaces. The system logs record interface failures.

Before configuring M-100 appliances to use multiple interfaces, verify the following:

- ❒ The M-100 appliances (Panorama or log collector mode) have PAN-OS 6.1 or later installed.
- ❒ The firewalls have PAN-OS 5.0 or later installed.
- ❒ The initial configuration of the M-100 appliances and firewalls is complete. This includes defining log collectors and collector groups.

To complete the configuration of an interface, you must specify the IP address, netmask (for IPv4) or prefix length (for IPv6), and default gateway. The Eth1 and Eth2 interfaces require complete configurations or else the commit operation will fail. You can commit a partial configuration for the MGT interface (for example, you might omit the default gateway) but then you can only access the M-100 appliance using the console port for future configuration changes. It is recommended that you commit a complete configuration for MGT.

Configure an M-100 Appliance to Use Multiple Interfaces		
Step 1	Configure the Management (MGT), Eth1, and Eth2 interfaces on the M-100 appliance.	1. Select **Panorama > Setup > Management**. 2. Edit the Management, Eth1 and Eth2 Interface Settings to define the configuration settings fore each interface and then click **OK**. 3. Click **Commit**, select **Panorama** as the **Commit Type**, then click **OK**.

Configure an M-100 Appliance to Use Multiple Interfaces (Continued)		
Step 2	Assign functions to the Eth1 and Eth2 interfaces of each log collector in the collector group.	1. Select **Panorama > Managed Collectors** and select the log collector. 2. In the **General** tab, select the interfaces that the collector will use for **Device Log Collection** (**Eth1** in this example) and **Collector Group Communication** (**Eth2** in this example). For an M-100 appliance that is in log collector mode, you must configure the additional interface settings in the **Eth1** and **Eth2** tabs. For the local (default) log collector on the M-100 appliance that is in Panorama mode, you already configured these settings on the **Panorama > Setup > Management** page. 3. Click **OK** and **Commit**, select **Panorama** as the **Commit Type**, then click **OK**. 4. Click **OK** and **Commit**, select **Collector Group** as the **Commit Type**, then click **OK**.
Step 3	If the collector group contains M-100 appliances in log collector mode, verify they are synchronized and connected.	Select **Panorama > Managed Collectors**. The Connected column displays a check mark ☑ icon to indicate that a log collector is connected to Panorama. The Configuration Status column indicates whether the configurations you committed to Panorama and the managed collectors are or are not synchronized with each other. A configuration mismatch can occur if: • After configuring the interfaces you performed a Panorama commit but not a collector group commit. • One or more log collectors has a PAN-OS version older than 6.1 and has an existing Eth1/Eth2 configuration. This usually occurs after a downgrade from 6.1 to 6.0. The Run Time Status column indicates whether the log collectors are connected to each other. To see more status information for a collector, hover over the corresponding Configuration or Run Time **Detail** cell and click the ellipsis **...** icon: a popup displays the information (if any).

Extended SNMP Support

PAN-OS support for Simple Network Management Protocol (SNMP) now includes the features described in the following topics:

- SNMP Support for LACP
- SNMP Support for M-100 Appliance Eth1 and Eth2 Interface Statistics

SNMP Support for LACP

PAN-OS support for Simple Network Management Protocol (SNMP) now includes the ability to monitor the status of aggregate groups that have Link Aggregation Control Protocol (LACP) enabled. When the firewall logs LACP events, it also generates traps. The following topics list the SNMP tables and traps that PAN-OS implements for LACP:

- SNMP Tables for LACP
- SNMP Traps for LACP

SNMP Tables for LACP

Download the IEEE 802.3 LAG MIB definition file to your trap receiver to enable monitoring of Link Aggregation Control Protocol (LACP) aggregate groups. PAN-OS implements the following SNMP tables for LACP:

Table	Description
Aggregator Configuration Table (dot3adAggTable)	This table contains information about every aggregate group that is associated with a firewall. Each aggregate group has one entry. Some table objects have restrictions, which the *dot3adAggIndex* object describes. This index is the unique identifier that the local system assigns to the aggregate group. It identifies an aggregate group instance among the subordinate managed objects of the containing object. The identifier is read-only. The ifTable MIB (a list of interface entries) does not support logical interfaces and therefore does not have an entry for the aggregate group.
Aggregation Port List Table (dot3adAggPortListTable)	This table lists the ports associated with each aggregate group in a firewall. Each aggregate group has one entry. The *dot3adAggPortListPorts* attribute lists the complete set of ports associated with an aggregate group. Each bit set in the list represents a port member. For non-chassis platforms, this is a 64-bit value. For chassis platforms, the value is an array of eight 64-bit entries.
Aggregation Port Table (dot3adAggPortTable)	This table contains LACP configuration information about every port associated with an aggregate group in a firewall. Each port has one entry. The table has no entries for ports that are not associated with an aggregate group.
LACP Statistics Table (dot3adAggPortStatsTable)	This table contains link aggregation information about every port associated with an aggregate group in a firewall. Each port has one row. The table has no entries for ports that are not associated with an aggregate group.

The *dot3adTablesLastChanged* object indicates the time of the most recent change to dot3adAggTable, dot3adAggPortListTable, and dot3adAggPortTable.

SNMP Traps for LACP

To enable PAN-OS to generate traps related to Link Aggregation Control Protocol (LACP), update your SNMP trap receiver with the latest version of the PAN-OS MIB. The MIB includes the following LACP-related traps:

Trap Name	Description
panLACPLostConnectivityTrap	The peer lost connectivity to the firewall.
panLACPUnresponsiveTrap	The peer does not respond to the firewall.
panLACPNegoFailTrap	LACP negotiation with the peer failed.
panLACPSpeedDuplexTrap	The link speed and duplex settings on the firewall and peer do not match.
panLACPLinkDownTrap	An interface in the aggregate group is down.
panLACPLacpDownTrap	An interface was removed from the aggregate group.
panLACPLacpUpTrap	An interface was added to the aggregate group.

SNMP Support for M-100 Appliance Eth1 and Eth2 Interface Statistics

For the M-100 appliance, Panorama support for Simple Network Management Protocol (SNMP) now includes the ability to monitor statistics for the Eth1 and Eth2 interfaces, in addition to the Eth0 (MGT) interface. The statistics are only available if you configure Multiple M-100 Appliance Interfaces (otherwise Panorama does not use Eth1 or Eth2). The ifTable MIB provides the statistics, which are the same as those provided for Eth0. In a MIB browser, you can see the interface entries under **mib-2 > interfaces > ifTable**.

Configurable Key Size for SSL Forward Proxy Server Certificates

When responding to a client in an SSL Forward Proxy session, the firewall creates a copy of the certificate presented to it by the destination server and uses it to establish its connection with the client. By default, the firewall generates certificates with the same key size as the certificate presented by the destination server. However, PAN-OS now provides the following key options for the SSL/TLS Forward Proxy Server feature. PAN-OS generates a certificate with the specified key size for its connection with the client.

Key Option	Description
Defined by destination host	PAN-OS determines the key size to use for the certificate based on the key size of the destination server certificate. If the destination server uses a 1024-bit RSA key, PAN-OS generates a certificate with that key size and an SHA-1 hashing algorithm. If the destination server uses a key size larger than 1024 bits (for example, 2048 bits or 4096 bits), PAN-OS generates a certificate that uses a 2048-bit RSA key and SHA-256 algorithm. This is the default setting.
1024-bit RSA	PAN-OS generates certificates that use a **1024-bit RSA** key and SHA-1 hashing algorithm regardless of the key size of the destination server certificates. As of December 31, 2013, public certificate authorities (CAs) and popular browsers have limited support for X.509 certificates that use keys of fewer than 2048 bits. In the future, depending on security settings, when presented with such keys the browser might warn the user or block the SSL/TLS session entirely.
2048-bit RSA	PAN-OS generates certificates that use a **2048-bit RSA** key and SHA-256 hashing algorithm regardless of the key size of the destination server certificates. Public CAs and popular browsers support 2048-bit keys, which provide better security than the 1024-bit keys.

The following procedure describes how to select a key option:

Configure a Key Size for SSL Forward Proxy Server Certificates

1. In a firewall, select **Device > Setup > Session**.

   In Panorama, select **Device > Setup > Session** and select a **Template**.
2. In the Decryption Settings section, click **Forward Proxy Server Certificate Settings**.
3. Select a **Key Size**:
   - **Defined by destination host**
   - **1024-bit RSA**
   - **2048-bit RSA**

   
   Changing the key size clears the current certificate cache.
4. Click **OK** and **Commit**.

Default Profile Group and Log Forwarding Settings

You can now use the web interface to allow new security policies and new security zones to use your organization's preferred settings for security profile groups or log forwarding by default. This allows you to quickly and easily create new security policies or add a new security zone with having to manually select your preferred settings for security profiles or log forwarding each time. Setting up a default security profile group or default log forwarding profile also helps you to enforce consistency for administrators creating new policy rules or security zones, by including your organization's preferred profile group and log forwarding options in new policies or zones automatically. Use the following topics to set up a default security profile group to be used in new security policies or a default log forwarding profile to be used in new security policies and new security zones:

- Set Up a Default Security Profile Group
- Set Up a Default Log Forwarding Profile

This feature is only applicable when using the web interface and does not apply when using the CLI or XML API.

Set Up a Default Security Profile Group

A security profile group is a set of security profiles that can be treated as a unit and then easily added to security policies to provide additional protection from threats, vulnerabilities, and data leaks. You can allow a default security profile group to used as the default profile setting for new security policies. Name a security profile group *default* to allow the profile group to be used as the default profile setting for a new security rule. See the following steps for details on how to use the web interface to add or override a default security profile group:

- If no *default* security profile exists, the profile settings for a new security policy are set to **None** by default.
- After upgrading to PAN-OS 6.1.0, this setting does not affect any currently configured security rules; if a security profile group is named *default*, only new security policies will use that group as the policy's default profile settings.

Set Up or Override a Default Security Profile Group	
• Set up a default security profile group.	1. Select **Objects > Security Profile Groups** and add a new security profile group or modify an existing security profile group. 2. **Name** the security profile group *default*: 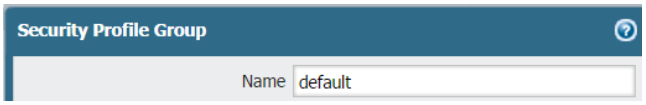 3. Click **OK** and **Commit**. 4. Confirm that the *default* security profile group is included in new security policies by default: a. Select **Policies > Security** and **Add** a new security policy. b. Select the **Actions** tab and view the **Profile Setting** fields: 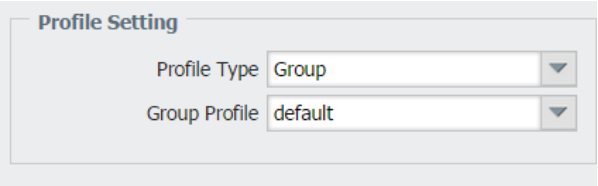 By default, the new security policy correctly shows the **Profile Type** set to Group and the *default* **Group Profile** is selected.
• Override a default security profile group.	If you have an existing default security profile group, and you do not want that set of profiles to be attached to a new security policy, you can continue to modify the Profile Setting fields according to your preference. Begin by selecting a different Profile Type for your policy (**Policies > Security > Security Policy Rule > Actions**).

Set Up a Default Log Forwarding Profile

Log forwarding profiles allow you to forward traffic logs to Panorama or an external system. A log forwarding profile can be added to a security zone to forward zone protection logs or to a security policy to forward logs for traffic that matches that policy. You can now allow a default log forwarding profile to be used as the default log forwarding settings for new security policies and new security zones. Name a log forwarding profile *default* to allow that profile to be used as the default log forwarding settings for new security rules and new security zones. See the following steps for details on how to use the web interface to add or override a default log forwarding profile:

- If no *default* log forwarding profile is configured, the log forwarding settings for a new security policy or new security zone will be set to **None** by default.
- After upgrading to PAN-OS 6.1.0, this setting does not affect any currently configured security policies or security zones; if a log forwarding profile is named *default*, only new security policies and security zones will use that profile as the policies' or zones' default log forwarding settings.

Set up or override a default log forwarding profile	
• Set up a default log forwarding profile.	1. Select **Objects > Log Forwarding** and add a new log forwarding profile or modify an existing profile. 2. **Name** the security profile group *default*: 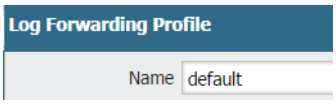 3. Click **OK** and **Commit**. 4. Confirm that the *default* log forwarding profile is included in new security policies by default: a. Select **Policies > Security** and **Add** a new security policy. b. Select the **Actions** tab and ensure the **Log Forwarding** field shows the *default* profile selected: 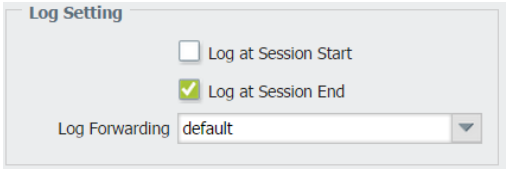 5. Confirm the *default* log forwarding profile is included in new security zones by default: a. Select **Network > Zones** and **Add** a new security zone. b. Ensure the **Log Setting** field shows the *default* log forwarding profile selected: 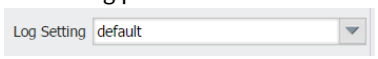
• Override a default log forwarding profile.	If you have an existing default log forwarding profile, and you do not want the log forwarding settings defined in that profile to be applied to a new security policy or a new security zone, continue to modify the **Log Setting** field in the policy or zone according to your preference.

WildFire Features

- Upgrade the WF-500 Appliance and Enable Windows 7 64-bit Support
- Signature/URL Generation on the WildFire Appliance
- Content Updates on the WF-500 WildFire Appliance
- WildFire Email Link Analysis
- Email Header Information in WildFire Logs
- Flash and Office Open XML File Type Support
- WildFire Analysis Report Enhancements
- WildFire XML API Support on the WF-500 Appliance

Upgrade the WF-500 Appliance and Enable Windows 7 64-bit Support

To upgrade the WF-500 WildFire appliance operating system to version 6.1, you must first download and install the Windows 7 64-bit image. The VM images can be as large as 4GB, so they must be downloaded from the Palo Alto Networks update servers and then hosted on an SCP-enabled server that you provide. You will then use the SCP client on the appliance to download the images from the SCP-enabled server prior to upgrading the appliance.

The appliance can only use one environment at a time to analyze samples, so after upgrading the appliance, you should review the list of available VM images and then choose the image that best fits your environment. In the case of Windows 7, if your environment has a mix of Windows 7 32-bit and Windows 7 64-bit systems, it is recommended that you choose the Windows 7 64-bit image to ensure that both 32-bit and 64-bit PE files will be analyzed in the sandbox environment.

Upgrade the WF-500 appliance before upgrading the firewalls that are configured to forward samples to it.

The following workflow describes how to upgrade the WF-500 appliance and enable the Windows 7 64-bit environment:

WF-500 WildFire Appliance Upgrade	

Step 1 Determine the upgrade path and download a base image file if needed.
  You cannot upgrade directly to the WildFire appliance operating system version 6.1 from version 5.1. Although you do not have to install version 6.0.0, you must first download the image and then download and install version 6.1.0. All releases have the requirement to download the base image files to skip a feature release.
	1. Log in to the WF-500 appliance and run the following command:
	`admin@WF-500> show system info`
	2. Check the `sw-version:` field to determine the installed version and proceed as follows:
	• If version 6.0.0 or later is installed, continue to step Step 2.
	• If a version prior to 6.0.0 is installed, continue the steps in this section.
	3. To download the 6.0.0 base image, run the following command:
	`admin@WF-500> request system software download version 6.0.0`
	4. To check the status of the download, run the following command:
	`admin@WF-500> show jobs all`
	5. After the download completes, continue to Step 2.

WF-500 WildFire Appliance Upgrade (Continued)		
Step 2	Download the required WildFire files to prepare for the 6.1.0 upgrade. In this case, you will need the WildFire operating system 6.1.0 image file, the Windows 7 64-bit base image, and the Windows 7 64-bit add-on image.	1. Download the WildFire operating system 6.1.0 base image file by running the following command: `admin@WF-500> request system software download version 6.1.0` 2. On the Palo Alto Networks Support site, click Software Updates and in the WF-500 Guest VM Images section locate and download the latest Windows 7 64-bit base image and the Windows 7 64-bit Add-on image.  The VM files can be as large as 4GB, so ensure that your Secure Copy (SCP) enabled server software supports file transfers over 4GB and verify that there is enough free space to temporarily store the files. The file names will be similar to the following: • **Base Image**—WFWin7_64Base_m-1.0.0_64base • **Add-on Image**—WFWin7_64Addon1_m-1.0.0_64addon 3. Move the files to your SCP-enabled server and note the file name and directory path.

Step 3	Download the VM images to the WF-500 appliance.	1. Download the base image file from the SCP-enabled server by running the following operational command:
		`admin@WF-500> scp import wildfire-vm-image from username@host:path`
		`For example:`
		`admin@WF-500> scp import wildfire-vm-image from bart@10.43.15.41:c:/scp/WFWin7_64Base_m-1.0.0_64base`
		  The SCP path following the IP or hostname may vary depending on the SCP software that you are using. For Windows, it will be `c:/folder/filename` or `//folder/filename`; for Unix/Mac systems, it will be `/folder/filename` or `//folder/filename`.
		2. Download the add-on image by running the following command:
		`admin@WF-500> scp import wildfire-vm-image from username@host:path`
		For example:
		`admin@WF-500> scp import wildfire-vm-image from bart@10.43.15.41:c:/scp/WFWin7_64Base_m-1.0.0_64addon`

Step 4	Install the Windows 7 64-bit VM images.	1. To install the Windows 7 64-bit base image, run the following command:
		`admin@WF-500> request system wildfire-vm-image upgrade install WFWin7_64Base_m-1.0.0_64base`
		2. To install the Windows 7 64-bit Add-on image, run the following command:
		`admin@WF-500> request system wildfire-vm-image upgrade install WFWin7_64Base_m-1.0.0_64addon`

Step 5	Install the 6.1 operating system image file.	Install the WF-500 appliance operating system image that you downloaded previously by running the following command: `admin@WF-500> request system software install file WildFire_m-6.1.0`

WF-500 WildFire Appliance Upgrade (Continued)		

Step 6	Restart the appliance and confirm that the installation was successful.	1. Confirm that the upgrade has completed by running the following command and look for the job type `Install` and status `FIN`:
		`admin@WF-500> **show jobs all**`
		`Enqueued ID Type Status`
		`Result Completed`
		`2014/07/30 10:38:48 2 Downld **FIN**`
		`OK 10:39:08`
		2. After the upgrade is complete, restart the appliance using the following command:
		`admin@WF-500> **request restart system**`
		3. Verify that the `sw-version` field shows 6.1 by running the following command:
		`admin@WF-500> **show system info | match sw-version**`

Step 7	(Optional) Enable the Windows 7 64-bit sandbox environment.	1. To view the active virtual machine image, run the following command and view the Selected VM field:
		`admin@WF-500> **show wildfire status**`
		2. To view a list of available virtual machines images, run the following command:
		`admin@WF-500> **show wildfire vm-images**`
		The following output shows that vm-5 is the Windows 7 64-bit image:
		`vm-5`
		`Windows 7 64bit, Adobe Reader 11, Flash 11, Office 2010. Support PE, PDF, Office 2010 and earlier`
		3. To select the image that will be used, enter configuration mode (type configure) and enter the following command:
		`admin@WF-500# **set deviceconfig setting wildfire active-vm <vm-image-number>**`
		For example, to use vm-5, run the following command:
		`admin@WF-500# **set deviceconfig setting wildfire active-vm vm-5**`
		4. `Commit` the configuration.

Signature/URL Generation on the WildFire Appliance

The WF-500 appliance can now generate signatures locally, eliminating the need to send any data to the public cloud in order to block malicious content. The WF-500 appliance can now analyze files forwarded to it from Palo Alto Networks firewalls or from the WildFire API and generate the following types of signatures that block both the malicious files as well as associated command and control traffic:

- **Antivirus signatures**—Detect and block malicious files. These signatures are added to WildFire and Antivirus content updates.
- **DNS signatures**—Detect and block callback domains for command and control traffic associated with malware. These signatures are added to WildFire and Antivirus content updates.
- **URL Categorization**—Categorizes callback domains as malware and updates the URL category in PAN-DB.

The WF-500 appliance will analyze Java content, but will not generate signatures for malicious samples. You must download the sample from the WildFire Submission log and upload it to the WildFire cloud for signature generation.

Firewalls must be running PAN-OS 6.1 or later to enable dynamic updates from a WF-500 appliance. In addition, you must configure the firewalls to receive content updates from the WF-500 appliance, which can occur as frequently as every five minutes. You can optionally send the malware sample file analysis data (or just the XML report if you don't want to send the sample) to the WildFire cloud to enable signature generation for distribution through Palo Alto Networks content releases.

When the local storage on the appliance is full, new signatures/URL categorizations will overwrite existing ones, beginning with the oldest ones first.

The following topics describe how to enable signature/URL generation on the WF-500 appliance and how to configure firewalls to retrieve content updates from the appliance:

▲ Enable Signature/URL Generation on the WF-500 Appliance

▲ Configure a Firewall to Retrieve Updates From a WF-500 Appliance

Enable Signature/URL Generation on the WF-500 Appliance

Use the following procedure to enable a WildFire appliance to generate signatures locally.

Enable Signature/URL Generation on the WildFire Appliance	

Step 1 Enable Signature/URL Generation.	1. Log in to the appliance and type `configure` to enter configuration mode.
	2. To enable all threat prevention options, enter the following command:
	`admin@WF-500# set deviceconfig setting wildfire signature-generation av yes dns yes url yes`
	3. Commit the configuration:
	`admin@WF-500# commit`
	4. Configure a Firewall to Retrieve Updates From a WF-500 Appliance.

Enable Signature/URL Generation on the WildFire Appliance		

Step 2	(Optional) Configure the WF-500 appliance to forward analysis reports or malicious samples (all files/URLs that are determined to be malicious) to the Palo Alto Networks WildFire cloud. If Packet Captures (PCAPS) are enabled, the PCAP will also be forwarded with the sample file.	1. To auto submit analysis reports, enter the following command:
		`admin@WF-500# set deviceconfig setting wildfire cloud-intelligence submit-report yes`
		If submit-sample is enabled as described in the following step, there is no need to enable submit-report because the sample will be re-analyzed in the cloud and a new report will be generated.
		2. To auto submit file samples, enter the following command:
		`admin@WF-500# set deviceconfig setting wildfire cloud-intelligence submit-sample yes`
		3. Commit the configuration:
		`admin@WF-500# commit`

Configure a Firewall to Retrieve Updates From a WF-500 Appliance

If you Enable Signature/URL Generation on the WF-500 Appliance, you can configure your firewalls to retrieve regular content updates from the appliance. This ensures that your network is protected from threats WildFire detects in your local environment. As a best practice, you should configure your firewalls to retrieve content updates from the Palo Alto Networks Update Servers and from the WildFire cloud. This will ensure that your firewalls receive signatures based on threats that are detected world wide.

The following workflow describes how to configure a Palo Alto Networks firewall to retrieve content updates from a WildFire appliance.

Configure the Firewall to Retrieve Updates from the WildFire Appliance		
Step 1	Launch the web interface and go to the Dynamic Updates page.	Select **Device > Dynamic Updates**.
Step 2	Check for the latest updates. To check the status of an action, click **Tasks** (on the lower right-hand corner of the window).	1. Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates. The link in the **Action** column indicates whether an update is available: • **Download**—Indicates that a new update file is available. Click the link to begin downloading the file directly to the firewall. After successful download, the link in the **Action** column changes from **Download** to **Install**. The following screen capture shows the new WF-Private section in Dynamic Updates. This is where you will download updates from the WF-500 appliance. 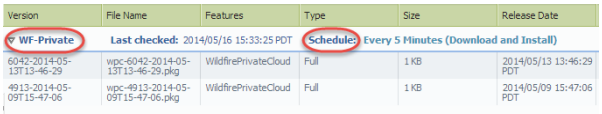 • **Revert**—Indicates that the corresponding update has been downloaded previously. You can choose to revert to the previously installed version of the update.
Step 3	Install the updates.	1. Click the **Install** link in the **Action** column. When the installation completes, a check mark displays in the **Currently Installed** column.

Configure the Firewall to Retrieve Updates from the WildFire Appliance (Continued)

Step 4 Schedule the update.

To receive updates at the minimal interval, configure the firewall to retrieve updates every five minutes.

1. Set the schedule of each update type by clicking the **None** link.

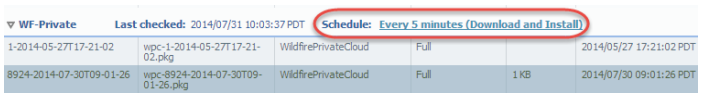

2. Specify how often you want the updates to occur by selecting a value from the **Recurrence** drop-down. The WF-500 appliance updates are available **Every 5 minutes** (best practice), **Every 15 minutes**, **Every 30 minutes**, or **Every Hour**.
3. Specify the **Time** and (or, minutes past the hour in the case of WildFire), if applicable depending on the **Recurrence** value you selected, **Day** of the week that you want the updates to occur.
4. Specify whether you want the system to **Download And Install** the update (best practice) or **Download Only**.
5. Specify how long after a content release to wait before performing a content update by entering the number of hours to wait in the **Threshold (Hours)** field. This provides added protection in the event that there are errors in a content release.
6. Click **OK** to save the schedule settings.
7. Click **Commit** to save the settings to the running configuration.

Content Updates on the WF-500 WildFire Appliance

To support the ability to generate signatures on the local WF-500 WildFire appliance, daily content updates are now available for the appliance. These content updates equip the appliance with the most up-to-date threat information for accurate malware detection and improve the appliance's ability to differentiate the malicious from the benign.

- Install Content Updates Directly from the Update Server
- Install Content Updates from an SCP-Enabled Server

Install Content Updates Directly from the Update Server

The following procedure describes how to install content updates on a WildFire appliance directly from the Palo Alto Networks Update Server.

Install Content Updates from the Update Server	

Step 1 Verify connectivity from the appliance to the update server and identify the content update to install.	1. Log in to the WildFire appliance and view the current system version information by running the following command:
	`admin@wf-500> show system info`
	The `wfm-content-version` field will show the current version.
	2. Confirm that the appliance can communicate with the Palo Alto Networks Update Server and view available updates by running the following command:
	`admin@wf-500> request wf-content upgrade check`
	The command queries the Palo Alto Networks Update Server and provides information about available updates and identifies the version that is currently installed on the appliance.
	`admin@wf-500> request wf-content upgrade check`
	`Version Size Released on Downloaded Installed`
	`---`
	`2-115 57MB 2014/04/20 20:00:08 PDT no no`
	`2-39 44MB 2014/02/12 14:04:27 PST yes current`
	If the appliance is not able to connect to the update server, you will need to allow connectivity from the appliance to the Palo Alto Networks Update Server, or download and install the update using SCP as described in Install Content Updates from an SCP-Enabled Server.

Install Content Updates from the Update Server (Continued)		

Step 2	Download and install the latest content update.	1. Download the latest content update:
		`admin@wf-500> request wf-content upgrade download latest`
		2. View the status of the download:
		`admin@wf-500> show jobs all`
		Run `show jobs pending` to view pending jobs. The following output shows that the download (job id 5) has finished downloading (Status FIN):
		`admin@wf-500> request wf-content upgrade download latest`
		`Download job enqueued with jobid 5`
		`admin@wf-500> show jobs all`
		`Enqueued ID Type Status Result Completed`
		`2014/04/22 03:42:20 5 Downld FIN OK 03:42:23`
		3. After the download is complete, install the update you just downloaded:
		`admin@wf-500> request wf-content upgrade install version latest`
		Run the `show jobs all` command again to view the status of the install.

Step 3	Verify the content update.	To verify the content release version, run the following command and refer to the `wfm-content-version` field:
		`admin@wf-500> show system info`
		The following shows an example output with content update version 2-115 installed:
		`admin@wf-500> show system info`
		`hostname: wf-500`
		`ip-address: 10.5.164.245`
		`netmask: 255.255.255.0`
		`default-gateway: 10.5.164.1`
		`mac-address: 00:25:90:c3:ed:56`
		`vm-interface-ip-address: 192.168.2.2`
		`vm-interface-netmask: 255.255.255.0`
		`vm-interface-default-gateway: 192.168.2.1`
		`vm-interface-dns-server: 192.168.2.1`
		`time: Mon Apr 21 09:59:07 2014`
		`uptime: 17 days, 23:19:16`
		`family: m`
		`model: WF-500`
		`serial: abcd3333`
		`sw-version: 6.1.0`
		`wfm-content-version: 2-115`
		`wfm-release-date: 2014/08/20 20:00:08`
		`logdb-version: 6.1.2`
		`platform-family: m`

Step 4	(Optional) Schedule content updates to install the latest updates on the firewall at a set interval.
	You can configure the appliance to install daily or weekly and either download only or download and install the updates.
		1. To schedule the appliance to download and install content updates, run the following command:
		`admin@WF-500# set deviceconfig system update-schedule wf-content recurring [daily | weekly] action [download-and-install | download-only]`
		For example, to download and install updates daily at 8:00 am, run the following command:
		`admin@WF-500# set deviceconfig system update-schedule wf-content recurring daily action download-and-install at 08:00`
		2. Commit the configuration
		`admin@WF-500# commit`

Install Content Updates from an SCP-Enabled Server

The following procedure describes how to install content updates on a WildFire appliance that does not have direct connectivity to the Palo Alto Networks Update Server. You will need a Secure Copy (SCP)-enabled server for this procedure.

Install Updates from an SCP-Enabled Server		
Step 1	Retrieve the content update file from the update server.	1. Log in to the Palo Alto Networks support site and click Dynamic Updates. 2. In the WildFire Appliance section, locate the latest WF-500 content update and download it. 3. Copy the content update file to an SCP-enabled server and note the file name and directory path.

Step 2	Install the content update on the WildFire appliance.	1. Log in to the WildFire appliance.
		2. To download the update file from the SCP server, run the following command:
		`admin@WF-500> scp import wf-content from username@host:path`
		For example:
		`admin@WF-500> scp import wf-content from bart@10.10.10.5:c:/updates/panup-all-wfmeta-2-182.tgz`
		If your SCP server is running on a non-standard port or if you need to specify the source IP, you can also define those options in the **scp import** command.
		3. To install the update, enter the following command:
		`admin@WF-500> request wf-content upgrade install file panup-all-wfmeta-2-182.tgz`
		To view status of the install, run the following command:
		`admin@WF-500> show jobs all`

Step 3	Verify the content update.	To view the current content version, run the following command and refer to the `wfm-content-version` field:
		`admin@wf-500> show system info | match wf-content-version`
		The following output shows an example output with content update version 2-115 installed:
		`admin@wf-500> show system info | match wf-content-version`
		`wfm-content-version: 2-115`

WildFire Email Link Analysis

The firewall can now extract HTTP/HTTPS links contained in SMTP and POP3 email messages and forward the links to the WildFire cloud for analysis (this feature is not supported on the WF-500 WildFire appliance). You enable this functionality by configuring the firewall to forward the **email-link** file type. Note that the firewall only extracts links and associated session information (sender, recipient, and subject) from the email messages that traverse the firewall; it does not receive, store, forward, or view the email message.

After receiving an email link from a firewall, WildFire visits the links to determine if the corresponding web page hosts any exploits. If it determines that the page itself is benign, no log entry will be sent to the firewall. However, if it detects malicious behavior on the page, it returns a malicious verdict and:

- Generates a detailed analysis report and logs it to the WildFire Submissions log on the firewall that forwarded the links. This log now includes the email header information—email sender, recipient, and subject—so that you can identify the message and delete it from the mail server and/or track down the recipient and mitigate the threat if the email has already been delivered and/or opened.
- Adds the URL to PAN-DB and categorizes the URL as malware.

Note that if the link corresponds to a file download, WildFire does not analyze the file. However, the firewall will forward the corresponding file to WildFire for analysis if the end user clicks the link to download it as long as the corresponding file type is enabled for forwarding.

To ensure that you gain the full benefits of this feature, confirm the following on each firewall that will forward samples to WildFire.

- ❒ A valid WildFire subscription is installed.
- ❒ WildFire content updates are configured to download-and-install frequently (every 15 minutes at minimum).
- ❒ PAN-DB is the active URL filtering vendor.

Configure Email Link Analysis

The following workflow describes how to enable the email link analysis feature:

Configure WildFire Email Link Analysis		
Step 1	Verify that the firewall is configured to forward files to the WildFire cloud.	1. Log in to the firewall and select **Device > Setup > WildFire**. 2. Click the Edit icon in the General Settings section. 3. Confirm that the **WildFire Server** field is set to `wildfire-public-cloud`.

Configure WildFire Email Link Analysis (Continued)		
Step 2	Configure a file blocking profile to define the applications and file types that will trigger forwarding to WildFire. In this example, you will forward the file type **email-link**.  In the following steps, you can select the file type **any** to forward all supported file types.	1. On the firewall that will forward files to WildFire, select **Objects > Security Profiles > File Blocking**. 2. Click **Add** to add a new profile and enter a **Name** and **Description**. 3. Click **Add** in the file blocking profile window and then click **Add** again. Click in the **Names** field and enter a rule name. 4. Select the **Applications** that will match this profile. For example, select **smtp** and **pop3**, which are the supported email applications for email link forwarding. 5. Select the **File Types** that will trigger the forwarding action. Choose **email-link** to forward http/https links in SMTP/POP3 traffic that matches the profile. 6. Select the **Direction** of the traffic that will be analyzed. For example, select **download** to detect and forward incoming SMTP/POP3 emails that contain an http/https link. 7. Select **Action** and choose **forward**. This will enable the firewall to forward the link to WildFire for analysis. 8. Click **OK** to save.
Step 3	Attach the file blocking profile to a security rule.	1. Select **Policies > Security**. 2. Click **Add** to create a new rule for the zones to which to apply WildFire forwarding, or select an existing security rule. 3. On the **Actions** tab, select the **File Blocking** profile from the drop-down.
Step 4	Configure the firewall to include email header information in the WildFire analysis reports. You can use the email header information to track down malicious emails and identify the recipient quickly. For more information, see Email Header Information in WildFire Logs.	1. Select **Device > Setup > WildFire** and then click the Edit icon in the Session Information Settings section. 2. Select the following options: **Email Sender**, **Email recipient**, and **Email subject**. 3. Click **OK** to save the changes.

Configure WildFire Email Link Analysis (Continued)		
Step 5	Verify that the firewall is forwarding email links.  WildFire does not generate logs for benign links, so to test this feature you must have access to an email that contains a malicious link. Also, the wildfire-upload-skip log will not be generated for skipped links.	1. Send an email containing a malicious HTTP or HTTPS link through the firewall. The email must be transferred over SMTP or POP3 and must match a policy with a file blocking profile that will forward the URL link. 2. Select **Monitor > Logs > Data Filtering**. 3. View the **Action** column for status. For each email link that matches a file blocking profile, the following actions are possible: • **forward**—The file blocking profile triggered the forwarding of the email link from the dataplane to the management plane. At this point, the link has not been forwarded to the WildFire cloud. • **wildfire-upload-success**—The firewall forwarded the email link to the WildFire cloud for analysis. For each malicious link that is successfully uploaded to WildFire, a subsequent log will be generated in the **WildFire Submissions** log that will contain a detailed analysis report. • **wildfire-upload-fail**—The link could not be uploaded to WildFire. This is typically caused by network communication issues between the firewall and the WildFire cloud. Verify connectivity and check DNS.
Step 6	View the analysis results for a malicious link.	1. Select **Monitor > Logs > WildFire Submissions** and locate the log that corresponds to the link and then click the detailed log view icon. You can filter by the category malicious to make it easier to find the log. 2. Select the **Log Info** tab to view session and email header information related to the malicious email link. 3. Select the **WildFire Analysis Report** tab to view full details of the email link analysis. For information on configuring reports and alerts for malicious verdicts, refer to the Palo Alto Networks WildFire Administrator's Guide.

Email Header Information in WildFire Logs

The firewall now captures email header information—email sender, recipient(s), and subject—and sends it along with the corresponding email attachments and email links that it forwards to WildFire. If WildFire determines that the email attachment or link is malicious, it includes the email header information in the WildFire Submissions log that it returns to the firewall. This information can help you to quickly track down and remediate threats that are detected in emails received by your users. Note that neither the firewall nor WildFire receive, store, or view the actual email contents. If the email header option is enabled, the firewall will automatically capture email header information associated with all email attachments and email links it forwards to WildFire.

The following workflow describes how to enable the email header options, how to set the User-ID attribute, and how to locate log information to help you identify recipients who have downloaded malicious email links or attachments.

Configure the Email Header Option for WildFire Logs		
Step 1	Change the WildFire Server setting on the firewall to point to the beta WildFire cloud server.  After you have completed beta testing, change the WildFire server back to `wildfire-public-cloud`.	1. On the firewall running PAN-OS 6.1 that is configured to forward samples to WildFire, select **Device > Setup > WildFire**. 2. Click the Edit icon and in the **WildFire Server** field, enter `beta.wildfire.paloaltonetworks.com`. 3. **Commit** the configuration.
Step 2	Enable the email header option.	1. On the firewall that will forward files to WildFire, select **Device > Setup > WildFire**. 2. Edit the Session Information Settings section and enable one or more of the options (**Email sender**, **Email recipient**, and **Email subject**). 3. Click **OK** to save.
Step 3	(Optional) Configure the User-ID option to enable the firewall to match User-ID information with email header information identified in email links and email attachments forwarded to WildFire. When a match occurs, the user name in the WildFire log email header section will contain a link that when clicked, will bring up the ACC filtered by the user or group of users.	1. Select **Device > User Identification > Group Mapping Settings**. 2. Select the desired group mapping profile to modify it. 3. In the **Server Profile** tab in the Mail Domains section, populate the **Domain List** field: • **Mail Attributes**—This field is automatically populated after you fill in the **Domain List** field and click **OK**. The attributes are based on your LDAP server type (Sun/RFC, Active Directory, and Novell). • **Domain List**—Enter the list of email domains in your organization using a comma separated list up to 256 characters.

Configure the Email Header Option for WildFire Logs (Continued)	
Step 4 Confirm that email header information is appearing in the WildFire reports. Within approximately 15 minutes after the file or link is forwarded, a WildFire log will be generated. Benign email links are not logged.	1. Select **Monitor > Logs > Data Filtering** from the firewall that forwarded the email link or attachment. 2. View the log and analysis report by selecting **Monitor > Logs > WildFire Submissions** and locate the corresponding log for the malicious link or file attachment. 3. Click the **log details** icon in the first column. In the **Log Info** tab, you will see the new email information in the Email Headers section. 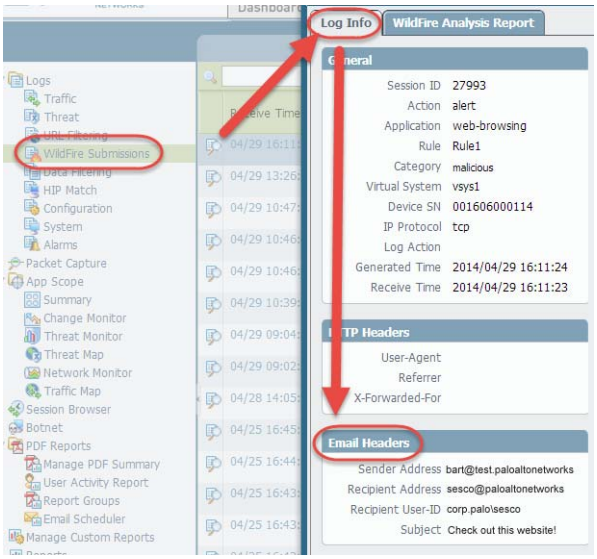 If User-ID is configured on the firewall, the domain and user name collected by User-ID are displayed in the **Recipient User-ID** field. Use the email header and User-ID information to track down the message on the mail server to delete it or use the information to locate the recipient to remove the threat if the email has already been opened.

Flash and Office Open XML File Type Support

Palo Alto Networks firewalls can now forward Flash content embedded in web pages to WildFire cloud and/or the WildFire appliance for analysis. In addition, WildFire now also creates antivirus signatures for Flash applets and Office Open XML (OOXML) 2007+ documents that it determines to be malicious and delivers the signatures through antivirus updates, enabling you to alert or block malicious content in these types of files. To support this capability, the firewall must have a WildFire subscription and be running Content Release version 450 or later.

- **Forward Flash Content**—Select the **flash** file type in a file blocking profile, select the forward action, and then attach the profile to a security rule. The **flash** file type includes the swf and swc (compressed flash components) file types.
- **Forward Office Open XML Files**—Select the **msoffice** file type in a file blocking profile, select the forward action, and then attach the profile to a security rule.

For details on configuring file blocking profiles and security policies, refer to the PAN-OS Administrator's Guide.

WildFire Analysis Report Enhancements

The WildFire detailed analysis report generated by the WildFire cloud and the WildFire appliance (WF-500) has been enhanced to improve efficiency in identifying and tracking down threats. For details on all report features, refer to the Palo Alto Networks WildFire Administrator's Guide.

The new/enhanced report features include:

- **Behavioral Summary and Severity**—This section of the report now describes details on each behavior that the sample file exhibits and a new column named Severity indicates the severity of each behavior. The severity gauge will show one bar for low severity and additional bars for higher severity levels. This information is also added to the dynamic and static analysis sections.

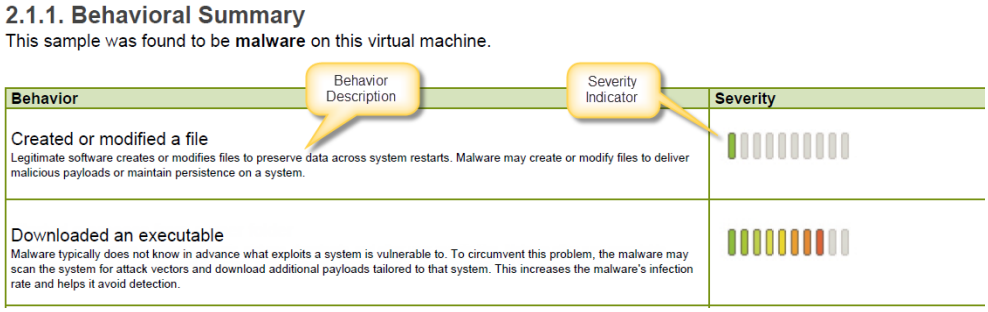

- **SHA-1**—The WildFire cloud and the WildFire appliance now generate a SHA-1 file hash for each sample that is analyzed and the hash value is added to the WildFire analysis report. The SHA-1 hash provides support for systems that perform API functions on WildFire that do not support the existing SHA-256 and MD5 algorithms. In the WildFire report, the hash values are located in the File Information section and in the File Activity section.

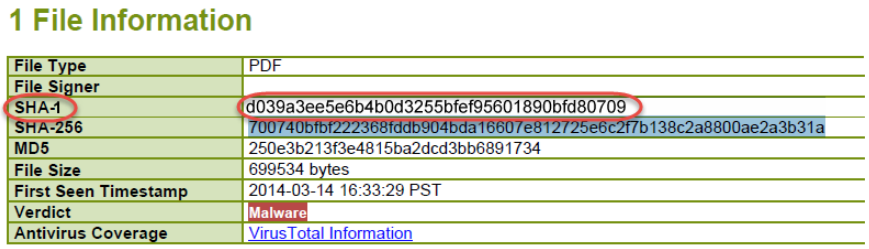

- **Coverage Status**—A new coverage status section dynamically updates when the report is rendered on the firewall with up-to-date information about what signature and URL filtering coverage that Palo Alto Networks currently provides to protect against the threat.

  The following coverage information is provided for active signatures:

  - **Coverage Type**—The type of protection provided by Palo Alto Networks (virus, dns, wildfire, or malware url). The type `wildfire` indicates that the malware was discovered by WildFire and protection was initially delivered as a WildFire signature.
  - **Signature ID**—A unique ID number assigned to each signature that Palo Alto Networks provides. This option is not applicable to malware URLs.

- **Detail**—The well-known name of the virus.
- **Date Released**—The date that Palo Alto Networks released coverage to protect against the malware.
- **Content Version**—The version number for the content release that provides protection against the malware.

The coverage status is dynamic, so the information is updated each time you view a WildFire detailed report for malware. Because this section is dynamic, it will only appear when viewing the log on the firewall and will not appear when viewing a PDF version of the report.

If the firewall is configured to forward files to a WildFire appliance, the firewall will query the appliance and the WildFire cloud to determine if coverage information is available. If `Coverage Status` is available for both systems (Cloud/Appliance), a separate table will appear for each system.

- **File Activity**—Three new fields now provide more details on files that were analyzed by Wildfire:
  - **Size(B)**—Indicates the size of the file in bytes.
  - **File Type**—Indicates the type of file, such as Office document, PDF, and PE.
  - **Hash**—Displays the MD5, SHA-1, and SHA-256 hash values. When viewing the report on the firewall, click a hash value to bring up a pop-up with the hash highlighted, so it can be copied and used in a search.

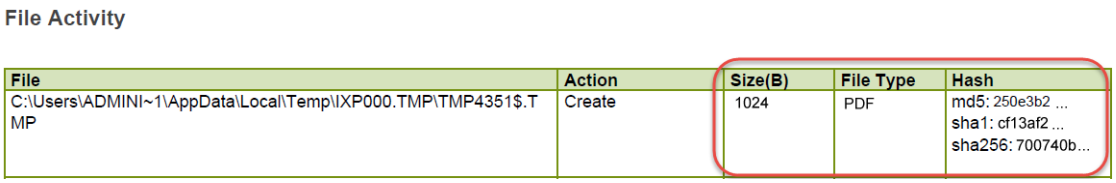

- **Miscellaneous Updates**
  - In the File Information section, the Sample File download link now uses the name of the file as the link name. If the file name cannot be determined, the link name is Download File.
  - Minor updates were also made to the WildFire reports on the WildFire Portal. The field hostname/IP is renamed to Firewall hostname/IP to make it clear that the IP address is for the firewall that forwarded the sample; source and destination have been renamed to file source and file destination to make it clear that the IP address is related to the source/destination of the sample file; and File Name is renamed to File Name/URL because the threat may be a malicious file or a URL.

WildFire XML API Support on the WF-500 Appliance

The WF-500 appliance now supports the WildFire XML API. To use the WildFire XML API with the appliance, you must generate the API keys on the appliance. The WF-500 appliance supports up to 100 API keys.

The following topics describe how to manage API keys on the appliance and provide an example on using the WildFire API to submit file samples to the appliance.

- Generate API Keys on the WildFire Appliance
- Manage API Keys on the WildFire Appliance
- Use the WildFire API on a WildFire Appliance

Generate API Keys on the WildFire Appliance

Generate an API Key	

Step 1 Generate a new API key on the WildFire appliance.
 	
As a best practice, leave out the key-value option in this step and the firewall will generate a key automatically. If you manually enter a key, the key-value must be 64 alpha characters (a-z) or numbers (0-9) that you randomly choose.
	1. Log in to the WildFire appliance CLI.
	2. Generate the API key using one of the following methods:
	• To generate a key automatically, run the following command:
	`admin@WF-500> create wildfire api-key name key-name`
	For example, to create a key with the name my-api-key, run the following:
	`admin@WF-500> create wildfire api-key name my-api-key`
	• To generate a key manually, run the following command (where key-value is a 64-bit key):
	`admin@WF-500> create wildfire api-key name my-api-key key <key-value>`

Step 2 View the API keys that you generated.	Run the following command to view all API keys: `admin@WF-500> show wildfire api-key all` This command also shows the date the key was generated and the last time the key was used. In this example, the following key was generated with the name my-api-key: **0377785F3F1A3D2DC6BCF2342730700747FBF4A23BD69F455F142494BC43D4A1**

Manage API Keys on the WildFire Appliance

This section describes some useful commands that you can use to manage WildFire API keys on the appliance and describes how to export and import the keys. For example, you may want to export all of your keys for backup purposes or to make it easier to access the keys from the systems that will use the API to perform various functions on the appliance.

Manage API Keys	
• Use the following commands to disable API keys temporarily, enable keys, or delete keys that are no longer used.	• To disable or enable an API key, run the following command: `admin@WF-500>` **`edit wildfire api-key status [disable | enable] key <api-key>`** For example, to disable the API key used in this example, run the following command: `admin@WF-500>` **`edit wildfire api-key status disable key 0377785F3F1A3D2DC6BCF2342730700747FBF4A23BD69F455F142494BC43D4A1`**  In the above command, you can type the first few unique digits of the key and then hit tab to fill in the remaining digits. • To delete an API key, run the following command: `admin@WF-500>` **`delete wildfire api-key key`** ***`api-key`*** For example: `admin@WF-500>` **`delete wildfire api-key key 377785F3F1A3D2DC6BCF2342730700747FBF4A23BD69F455F142494BC43D4A1`**
• Use the following commands to import or export API keys from the appliance using Secure Copy (SCP).	• Save all API keys to a file to prepare the keys for export by running the following command from configuration mode: `admin@WF-500#` **`save wildfire api-key to <filename>`** For example: `admin@WF-500>` **`save wildfire api-key to my-api-keys`** To SCP the API key file to an SCP-enabled server, run the following operational command: `admin@WF-500>` **`scp export wildfire-api-keys to <username@host:path>`** For example: `admin@WF-500>` **`scp export wildfire-api-keys to bart@10.10.10.5:c:/scp/`** You can also import the key from an SCP-enabled server by running the following command: `admin@WF-500>` **`scp import wildfire-api-keys from bart@10.10.10.5:c:/scp/my-api-keys`** • After importing API keys, you must load the keys by running the following command: `admin@WF-500#` **`load wildfire api-key mode [merge | replace] from my-api-keys`** If you leave out the `mode` option, the default behavior will merge the new keys. You can use the `replace` option to replace all API keys on the appliance. For example, to replace all API keys, enter the command: `admin@WF-500#` **`load wildfire api-key mode replace from my-api-keys`** • To confirm the keys were loaded, run the following command: `admin@WF-500>` **`show wildfire api-keys all`**

Use the WildFire API on a WildFire Appliance

The following workflow describes how to use the WildFire API to submit sample files to a WildFire appliance for threat analysis. After understanding the basic concepts illustrated in this example, refer to the Palo Alto Networks WildFire Administrator's Guide for details on other useful API functions.

This workflow requires a host computer that has the cURL command line tool installed. You will then send files from the computer to the WildFire appliance using URL syntax.

Use the WildFire API to Submit a File Sample

Step 1 Generate a WildFire API key for the host computer that will perform API functions on the WildFire appliance. For details, see Generate API Keys on the WildFire Appliance.

1. Access the CLI on the WildFire appliance and generate an API key using the following command:

   ```
   admin@WF-500> create wildfire api-key name <key-name>
   ```

2. Run the following command to show the API keys:

   ```
   admin@WF-500> show wildfire api-key all
   ```

3. Make sure the key status is Enabled and then highlight and copy the key. The following screen capture shows an example API key named my-api-key.

```
admin@WF-500-TME> show wildfire api-keys all

+----------------------------------------------------------------------+------------+---------+---------------------+--
| Apikey                                                               | Name       | Status  | Create Time         | L
+----------------------------------------------------------------------+------------+---------+---------------------+--
| D414CC910E93E9E05942A5E6F94DA36777B444543E71761CF5E9ACFA547F7D6F |            | Enabled | 2014-06-19 11:26:08 | 2
| 423868D964ADAB99172B32596C8D9FCE1F67731D8BA273C8AE3E09765F4FE0DF |            | Enabled | 2014-06-23 09:34:17 |
| 87C142CB01CA5BEBE06E226A25C0A473B34050B617073E21E8F1A6BCB8C5C387 | my-api-key | Enabled | 2014-06-23 13:23:37 |
+----------------------------------------------------------------------+------------+---------+---------------------+--
```

Use the WildFire API to Submit a File Sample (Continued)

Step 2 Using the new API key that you generated, submit a sample file to the WildFire appliance.

1. Place a sample file in a folder that can be accessed from the host computer that has the cURL command line tool installed and note the path of the sample file.
2. To submit the file using cURL, run the following command from the host computer:

   ```
   curl -k -F apikey=your-API-key -F file=@local-file-path --remote-name
   https://WF-appliance-IP/publicapi/submit/file
   ```

   The syntax will vary based on the host that you are using. The following examples shows the syntax using a Linux host and a Windows host.

   From a Linux host, run the following command:

   ```
   curl -k -F apikey=87C142CB01CA5BEBE06E226A25C0A473B34050B617073E21E8F1A6BCB8C5C387 -F
   file=@test-wf-api.docx --remote-name https://10.3.4.99/publicapi/submit/file
   ```

   From a Windows host, run the following command: (the only difference is the file path following the @ symbol):

   ```
   curl -k -F apikey=87C142CB01CA5BEBE06E226A25C0A473B34050B617073E21E8F1A6BCB8C5C387 -F
   file=@c://scp/test-wf-api.docx --remote-name https://10.3.4.99/publicapi/submit/file
   ```

3. Verify that the file was successfully submitted to the WildFire appliance. To view a list of recent samples submitted to the appliance, run the following command:

   ```
   admin@WF-500> show wildfire latest samples
   ```

   The following screen capture shows that the firewall successfully submitted the file named test-wf-api.docx to the appliance:

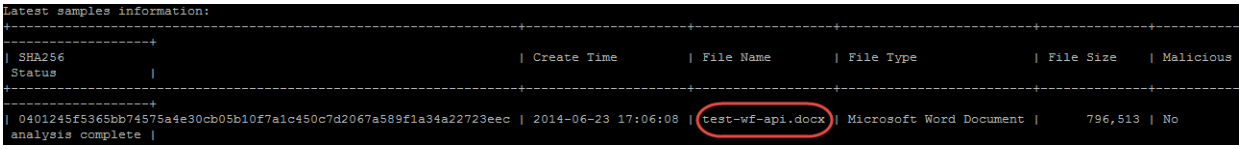

If the sample file does not appear on the appliance, verify connectivity between the host computer and the appliance and confirm that the folder/file path is correct. You can also run `show wildfire status` (status should show `Idle`) and `show wildfire statistics` to verify that the appliance is ready to analyze files. For more information on troubleshooting, refer to the WildFire Administrator's Guide.

URL Filtering Features

- Log HTTP Headers in Web Requests
- Manual Upload of BrightCloud Database

Log HTTP Headers in Web Requests

URL filtering provides visibility and control over web traffic on your network. For improved visibility into web content, you can now configure the URL Filtering Profile to log HTTP header attributes included in a web request. When a client requests a web page, the HTTP header includes the user agent, referer, and x-forwarded-for fields as attribute-value pairs and forwards them to the web server. When enabled for logging HTTP Headers, the firewall logs the following attribute-value pairs in the URL Filtering logs:

Attribute	Description
User-Agent	The web browser that the user used to access the URL, for example, Internet Explorer. This information is sent in the HTTP request to the server.
Referer	The URL of the web page that linked the user to another web page; it is the source that redirected (referred) the user to the web page that is being requested.
X-Forwarded-For	The header field option that preserves the IP address of the user who requested the web page. It allows you to identify the IP address of the user particularly if you have a proxy server on your network, where all requests might seem to originate from the proxy server's IP address.

To view the HTTP header, check for the HTTP Headers widget in the detailed log view in **Monitor > Logs > URL filtering** tab. If there are multiple URLs in a single session, each URL has a separate log with its own set of HTTP headers. The associated headers are grouped together and can be viewed as a set of related logs. Further, to aid in correlating data and analyzing web activity across the network, the HTTP Header options are also displayed with the corresponding threat logs and Wildfire logs.

The HTTP Header fields are available for generating custom reports on the firewall and on Panorama; they are also available for custom log-forwarding to an external syslog server.

Obtain and install a URL filtering license on the firewall.
This feature also requires that you install content update version 454 or later on the firewall/Panorama.

Enable HTTP Header Logging	
Step 1 Create a URL Filtering profile or select an existing one.	1. Select **Objects > Security Profiles > URL Filtering**. Select the default profile and then click **Clone**. The new profile will be named default-1. 2. Select the new profile and rename it.
Step 2 Define how to control access to web content.	In the **Categories** tab, for each category that you want visibility into or control over, select a value from the **Action** column as follows: • If you do not care about traffic to a particular category (that is you neither want to block it nor log it), select **Allow**. • For visibility into traffic to sites in a category, select **Alert**. • To deny access to traffic that matches the category and to enable logging of the blocked traffic, select **Block**.

Enable HTTP Header Logging (Continued)		
Step 3	Specify what to log. The **Log container page only** option is enabled by default so that only the main page that matches the category is logged, not subsequent pages/categories that may be loaded within the container page.	In the **Settings** tab, enable the options for **HTTP Header Logging**. 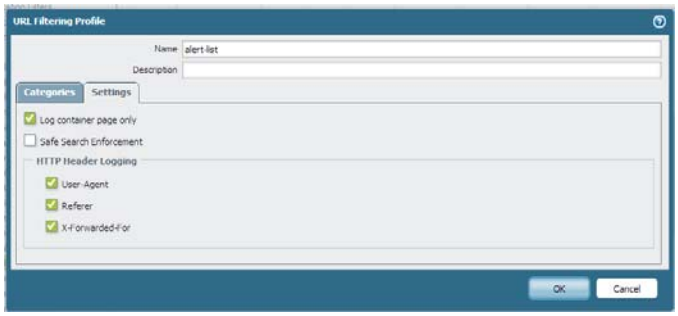
Step 4	Attach the URL Filtering profile to a policy rule.	1. Select **Policies > Security** and select the appropriate security policy rule to modify. 2. Select the **Actions** tab and in the **Profile Setting** section, select the profile you just created from the **URL Filtering** drop-down. (If you don't see drop-downs for selecting profiles, select **Profiles** from the **Profile Type** drop-down.) 3. Click **OK** to save the profile. 4. **Commit** the configuration.
Step 5	View the URL filtering logs.	1. Select **Monitor > Logs > URL Filtering**. 2. Click the details icon next to a specific log to view the HTTP Headers widget in the detailed log view. 3. (Optional) Adjust the log display to include the **Referer**, **User-Agent**, and **X-Forwarded-For** columns.
Step 6	View related logs for easy data correlation between log entries. You will no longer have to close the URL log and find the corresponding entry in the threat logs.	Click into a URL filtering log and scroll/click through the related log entries. When you click a related log, for example a threat log entry, the widget views switch to display the relevant details.

Manual Upload of BrightCloud Database

In deployments where Panorama or a firewall has no direct Internet access, you can now manually upload a BrightCloud database and install it.

Upload and Install a BrightCloud Database		
Step 1	Download the BrightCloud database to a host that has Internet access. Depending on where you will upload the database, Panorama or the firewall must have access to the host.	1. On a host with Internet access, go to the Palo Alto Support website (https://support.paloaltonetworks.com) and log in. 2. In the Resources section, click **Dynamic Updates**. 3. In the BrightCloud Database section, click **Download** and save the database binary file to the host.
Step 2	Upload the database.	1. If you will upload to Panorama, log in to Panorama, select **Panorama > Device Deployment > Dynamic Updates**, then click **Upload**. If you will upload to a firewall, log in to the firewall, select **Device > Dynamic Updates** and click **Upload**. 2. For the **Type**, select **URL Filtering**. 3. Enter the path to the database binary **File** on the host or click **Browse** to find it, then click **OK**. 4. When the Status is **Completed**, click **Close**.
Step 3	Install the database.	1. Click **Install From File**. 2. Select **Type**: **URL Filtering**. If you logged in to a firewall, it automatically selects the file you just uploaded. 3. (Panorama only) In the **File Name** drop-down, select the file you just uploaded. 4. (Panorama only) Select the firewalls on which to install the database. Optionally, if the list of firewalls is long, use the Filters check boxes and search field to filter the list. 5. Click **OK** and, when the Result is **Succeeded**, click **Close**.

GlobalProtect Features

The following sections describe how to configure the GlobalProtect features introduced in PAN-OS 6.1.0:

- Disconnect on Idle
- Disable Browser Access to the Portal Login Page
- Extended SSO Support for GlobalProtect Agents

For details on GlobalProtect features introduced in the GlobalProtect Mobile Security Manager 6.1.0 release, including support for an enterprise app store and VPN for business apps on mobile devices, refer to the GlobalProtect Mobile Security Manager New Features Guide Version 6.1.

As a part of PAN-OS 6.1.0, you can now configure GlobalProtect gateways on VM-Series firewalls deployed in the AWS cloud. By deploying the VM-Series firewall in the AWS cloud, you can quickly and easily deploy GlobalProtect gateways in any region without the expense or IT logistics that are typically required to set up this infrastructure using your own resources. For details, see the Use Case: VM-Series Firewalls as GlobalProtect Gateways in AWS.

Disconnect on Idle

The options to time out GlobalProtect clients have been extended to include settings you can use to log out idle users. You can set the number of minutes after which users will be disconnected from GlobalProtect if there is no traffic going through the VPN.

Configure Timeout Settings for GlobalProtect Clients	
• Define the login lifetime for a single gateway login session.	1. Select **Network > GlobalProtect > Gateways** and select the GlobalProtect gateway configuration to modify. 2. Select **Client Configuration > Tunnel Settings**. 3. For **Login Lifetime**, specify the number of **Minutes**, **Hours**, or **Days** at which the login session would automatically log out.
• Specify the amount of time after which an inactive session is automatically logged out.	1. Select **Network > GlobalProtect > Gateways** and select the GlobalProtect gateway configuration to modify. 2. Select **Client Configuration > Tunnel Settings**. 3. For **Inactivity Logout**, specify the number of **Minutes**, **Hours**, or **Days** at which a client would be logged out of GlobalProtect if the gateway does not receive a HIP check from the client the given amount of time.
• (New) Select the number of minutes after which to log out idle users.	1. Select **Network > GlobalProtect > Gateways** and select the GlobalProtect gateway configuration to modify. 2. Select **Client Configuration > Tunnel Settings**. 3. For **Disconnect on Idle**, specify the number of **Minutes** at which a client is logged out of GlobalProtect if the GlobalProtect app has not routed traffic through the VPN tunnel in the given amount of time.

Disable Browser Access to the Portal Login Page

You can now choose to disable access to the GlobalProtect portal login page from a web browser. You can use this feature to prevent public access to the portal login page and unauthorized attempts to authenticate to the GlobalProtect portal from a browser.

With the portal login page disabled, you can instead use a software distribution tool, such as Microsoft's System Center Configuration Manager (SCCM), to allow your users to download and install the GlobalProtect agent. Enabling this option only prevents the login page from displaying on a web browser and does not affect GlobalProtect agents' or GlobalProtect apps' access to the portal; GlobalProtect agents and apps continue to successfully authenticate and connect to the portal to receive configuration updates.

You can select the new option to **Disable login page** when configuring a portal or modifying a portal configuration.

Disable the Portal Login Page		
Step 1	View the current login page setup.	On the firewall web interface, select **Network > GlobalProtect > Portals > GlobalProtect Portal > Portal Configuration**.
Step 2	Disable the login page.	On the **Portal Configuration** tab, select **Disable login page**.
Step 3	Save your settings.	Click **OK** and **Commit**.

Extended SSO Support for GlobalProtect Agents

With Single Sign-On (SSO), the GlobalProtect agent wraps the user's Windows login credentials to automatically authenticate and connect to the GlobalProtect portal and gateway. SSO has been enhanced in this release so that when a third-party credential provider is being used to wrap the user's Windows login credentials, the GlobalProtect agent can also wrap the third-party credentials. With this feature enabled, users can successfully authenticate to Windows, GlobalProtect, and the third-party credential provider in one step, by using their Windows logon credentials to log on to their Windows system. This extended SSO functionality is supported on Windows 7 and Windows Vista clients.

Use the Windows Registry or the Windows Installer to allow GlobalProtect to wrap third-party credentials:

- Enable SSO Wrapping for Third Party Credentials with the Windows Registry
- Enable SSO Wrapping for Third Party Credentials with the Windows Installer

GlobalProtect SSO wrapping for third-party credential providers (CPs) is dependent on the third-party CP settings and in some cases, GlobalProtect SSO wrapping might not work correctly if the third-party CP implementation does not allow GlobalProtect to successfully wrap their CP.

Enable SSO Wrapping for Third Party Credentials with the Windows Registry

Use the following steps in Windows Registry to enable SSO to wrap third party credentials on Windows 7 and Windows Vista clients.

Use the Windows Registry to Enable SSO Wrapping for Third Party Credentials	
Step 1 Open the Windows Registry and locate the globally unique identifier (GUID) for the third party credential provider that you want to wrap.	1. In the Command Prompt, enter the command `regedit` to open the Windows Registry. 2. Locate currently installed credential providers at the following location: `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Authentication\Credential Providers.` 3. Copy the GUID key for the credential provider that you want to wrap (including the curly brackets— { and } —on either end of the GUID): 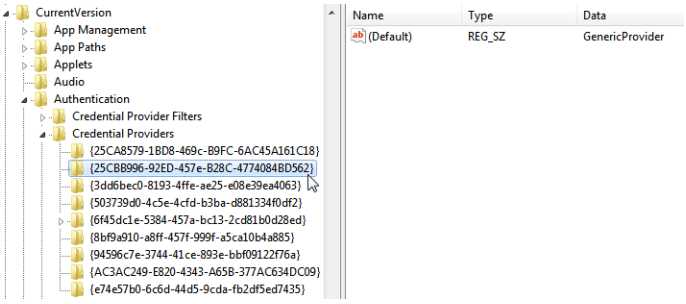

Use the Windows Registry to Enable SSO Wrapping for Third Party Credentials	
Step 2 Enable SSO wrapping for third party credential providers by adding the `wrap-cp-guid` key to the GlobalProtect registry settings.	1. Go to the following Windows Registry location: `HKEY_LOCAL_MACHINE\SOFTWARE\Palo Alto Networks\GlobalProtect`: 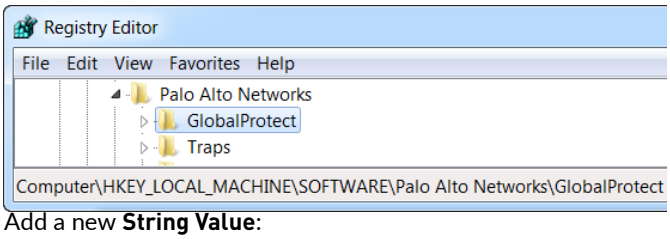 2. Add a new **String Value**: 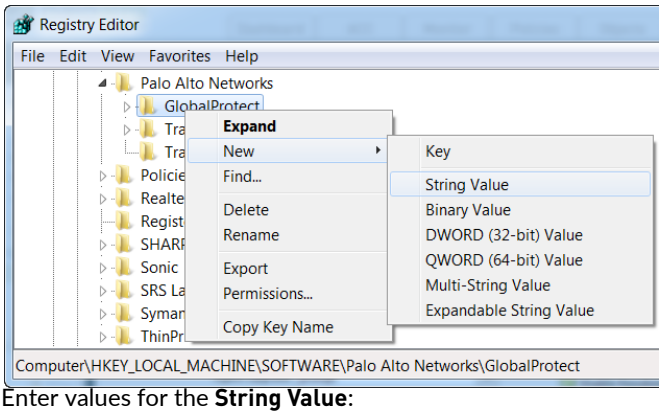 3. Enter values for the **String Value**: • **Name**: `wrap-cp-guid` • **Value data**: {*third party credential provider GUID*} For the **Value data** field, the GUID value that you enter must be enclosed with curly brackets: `{` and `}`. The following is an example of what a third party credential provider GUID in the **Value data** field might look like: `{A1DA9BCC-9720-4921-8373-A8EC5D48450F}` For the new String Value, `wrap-cp-guid` is displayed as the String Value's Name and the GUID is displayed as the Data. 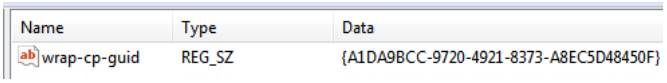
Next Steps...	• SSO wrapping for third party credential providers is successfully set up by completing steps 1 and 2. With this setup, the native Windows logon tile is displayed to users. Users click the tile and log on to the system with their Windows credentials. The single logon authenticates the users to Windows, GlobalProtect, and the third party credential provider. • (Optional) If you want to display two tiles to users at logon, the native Windows tile and the tile for the third party credential provider, continue to Step 3.

Use the Windows Registry to Enable SSO Wrapping for Third Party Credentials	
Step 3	(Optional) Allow the third party credential provider tile to be displayed to users at logon.	Add a second **String Value** with the **Name** `filter-non-gpcp` and enter `no` for the string's **Value data**: 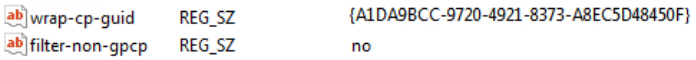 With this string value added to the GlobalProtect registry settings, users will be presented with two logon options when logging on to their Windows system: the native Windows tile and the third party credential provider's tile.

Enable SSO Wrapping for Third Party Credentials with the Windows Installer

Use the following options in Windows Installer (MSIEXEC) to enable SSO to wrap third party credential providers on Windows 7 and Windows Vista clients.

Use the Windows Installer to Enable SSO Wrapping for Third Party Credentials
• Wrap third party credentials and display the native tile to users at logon. Users click the tile and log on to the system with their native Windows credentials. The single logon authenticates users to Windows, GlobalProtect, and the third-party CP. Use the following syntax from the Windows Installer (MSIEXEC): `msiexec.exe /i GlobalProtect.msi WRAPCPGUID="{guid_value}" FILTERNONGPCP="yes"` In the syntax above, the `FILTERNONGPCP` parameter simplifies authentication for the user by filtering the option to log on to the system using the third party credentials.
• If you would like users to have the option to log in with the third party credentials, use the following syntax from the MSIEXEC: `msiexec.exe /i GlobalProtect.msi WRAPCPGUID="{guid_value}" FILTERNONGPCP="no"` In the syntax above, the `FILTERNONGPCP` parameter, which filters out the third party credential provider's logon tile so that only the native tile displays, is set to `"no"`. In this case, both the native Windows tile and the third party credential provider tile is displayed to users when logging on to the Windows system.

Networking Features

LACP

The firewall can now use Link Aggregation Control Protocol (LACP) to detect the physical interfaces between itself and a connected device (peer) and manage those interfaces as a single virtual interface (aggregate group). An aggregate group increases the bandwidth between peers. Enabling LACP provides redundancy within the group: the protocol automatically detects interface failures and performs failover to standby interfaces. Without LACP, you must manually identify interface failures occurring at layers above the physical or occurring between peers that are not directly connected.

The following Palo Alto Networks firewalls support LACP: PA-500, PA-3000 Series, PA-4000 Series, PA-5000 Series, and PA-7050. The firewalls support LACP for HA3 (only on the PA-500, PA-3000 Series, PA-4000 Series, and PA-5000 Series), Layer 2, and Layer 3 interfaces.

Before configuring LACP:

- Determine which physical interfaces connect the LACP peers. This procedure assumes the cabling is complete.
- Determine the optimal LACP Settings for the peers.

This procedure only covers the configuration steps for an LACP peer that is a Palo Alto Networks firewall. Other devices have proprietary methods for configuring LACP.

Perform the following steps to configure LACP on a firewall. In a high availability deployment, configure the primary (for active/active) or active (for active/passive) firewall; the secondary or passive firewall automatically synchronizes with it.

Configure LACP	
Step 1 Add an aggregate group with LACP enabled.	1. Select **Network > Interfaces > Ethernet** and click **Add Aggregate Group**. 2. In the field adjacent to the read-only **Interface Name**, enter a number (1-8) to identify the group. 3. For the **Interface Type**, select **HA**, **Layer 2**, or **Layer 3**. Only select **HA** if the interface is an HA3 link between two firewalls in an active/active deployment. 4. On the **LACP** tab, select **Enable LACP**. 5. For the LACP **Mode**, select whether the firewall is **Passive** (default) or **Active**. 6. For the **Transmission Rate**, select **Fast** or **Slow**. 7. If desired, select **Fast Failover** (under one second) for when interfaces go down. Otherwise, the standard failover (at least three seconds) applies. 8. Enter a **System Priority** number (1-65535, default 32768) that determines whether the firewall or peer overrides the other with respect to port priorities. Note that the lower the number, the higher the priority. 9. For the **Max Ports**, enter the number of interfaces (1-8) that are active. The value cannot exceed the number of interfaces you assign to the group. If the number of assigned interfaces exceeds the number of active interfaces, the remaining interfaces will be in standby mode. 10. By default, each firewall in an HA pair has a unique MAC address. A **Same System MAC Address for Active-Passive HA** configuration is recommended only when the LACP peers are virtualized (appearing to the network as a single device). In such cases, select the check box and select the system-generated **MAC Address**, or enter your own. You must verify the address is globally unique. If the firewalls are not in active/passive HA mode, PAN-OS ignores your selections for these fields. Firewalls in an active/active deployment require unique MAC addresses so PAN-OS automatically assigns them. If the **Interface Type** is **HA**, these fields do not appear. 11. Click **OK**.

Configure LACP (Continued)		
Step 2	Assign interfaces to the aggregate group.	Perform the following steps for each physical interface (1-8) that will belong to the aggregate group. During normal operations, assign more than one interface. You would only assign a single interface if a troubleshooting procedure requires it. 1. Select **Network > Interfaces** and select the interface name. It must be the same type (HA, Layer 2, or Layer 3) as the aggregate group. 2. Change the **Interface Type** to **Aggregate Ethernet**. 3. Select the **Aggregate Group** you just defined. 4. Select the **Link Speed**, **Link Duplex**, and **Link State**. It is a best practice to set the same link speed and duplex values for every interface in the group. For non-matching values, the commit operation displays a warning and PAN-OS defaults to the higher speed and full duplex. 5. Enter an **LACP Port Priority** (1-65535, default 32768). If the number of interfaces you assign exceeds the **Max Ports** value you configured for the group, the port priorities determine which interfaces are active or standby. The lower the numeric value, the higher the priority. 6. Click **OK** and **Commit**.
Step 3	Verify the status of the aggregate group.	1. Select **Network > Interfaces > Ethernet**. 2. Verify that the Features column displays an LACP enabled icon for the aggregate group. 3. Verify that the Link State column displays a green icon for the aggregate group, indicating that all member interfaces are up. If the icon is yellow, at least one member is down but not all. If the icon is red, all members are down. 4. If the Link State icon for the aggregate group is yellow or red, select **Monitor > Logs > System** and review the logs for the LACP subtype to investigate the causes of the interface failures.
Step 4	(Optional) If you selected **HA** as the **Interface Type** for the aggregate group, enable forwarding of packets over the HA3 link.	1. Select **Device > High Availability > Active/Active Config** and edit the Packet Forwarding section. 2. For the **HA3 Interface**, select the aggregate group you configured. 3. Click **OK** and **Commit**.

NAT Capacity Enhancements

NAT is enhanced to increase the number of rules allowed, to display more statistics per resource pool and NAT rule, and to allow configuration of Dynamic IP and Port (DIPP) NAT oversubscription. These enhancements apply to the PA-3000 Series, PA-4000 Series, PA-5000 Series, and PA-7050 platforms.

Increase in Number of NAT Rules Allowed

This release allows more NAT rules; the quantity is based on the platform. Individual rule limits are set for static, Dynamic IP (DIP), and Dynamic IP and Port (DIPP) NAT. The sum of the number of rules used for these NAT types cannot exceed the total NAT rule capacity. For DIPP, the rule limit is based on the device's oversubscription setting (8, 4, 2, or 1) and the assumption of one translated IP address per rule. To see platform- specific NAT rule limits and translated IP address limits, use the Compare Firewalls tool.

Consider the following when working with NAT rules:

- If you run out of pool resources, you cannot create more NAT rules, even if the platform's maximum rule count has not been reached.
- If you consolidate NAT rules, the logging and reporting will also be consolidated. The statistics are provided per the rule, not per all of the addresses within the rule. If you need granular logging and reporting, do not combine the rules.

Additional Dataplane NAT Memory Statistics

The `show running global-ippool` command displays statistics related to NAT memory consumption for a pool. The new columns are Size, which displays the number of bytes of memory that the resource pool is using, and Ratio, which displays the oversubscription ratio (for DIPP pools only). The new lines of pool and memory statistics are explained in the following sample output:

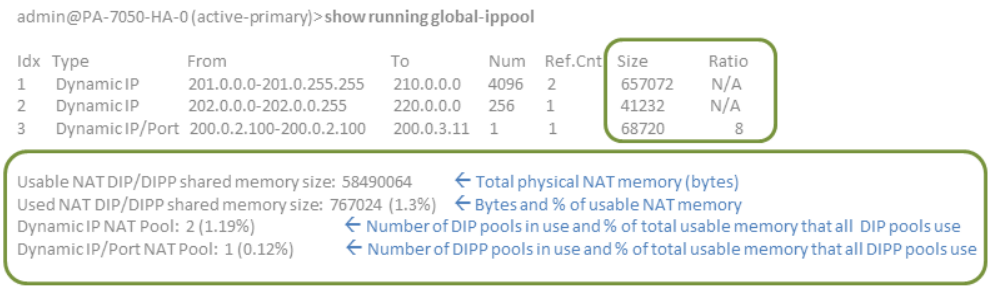

For NAT pool statistics for a virtual system, the `show running ippool` command has two new columns: the memory size used per NAT rule and the oversubscription ratio used (for DIPP rules). The following is sample output for the command.

```
admin@PA-7050-HA-0 vsys1 (active-primary)> show running ippool

VSYS 1 has 4 NAT rules, DIP and DIPP rules:
Rule            Type              Used    Available    Mem Size    Ratio
------------    ---------------   ------  -----------  ----------  -------
nat1            Dynamic IP        0       4096         788144      0
nat2            Dynamic IP        0       256          49424       0
nat3            Dynamic IP/Port   0       638976       100976      4
nat11           Dynamic IP        0       4096         788144      0
```

A new field is added to the output of the `show running nat-rule-ippool rule` command to show the memory (bytes) used per NAT rule. The following is sample output for the command, with the memory usage for the rule encircled.

```
admin@PA-7050-HA-0 (active-primary)> show running nat-rule-ippool rule nat1

VSYS 1 Rule nat1:
Rule: nat1, Pool index: 1,  memory usage: 788144
------------------------------------------------------------
Reserve IP: no
201.0.0.0-201.0.255.255  =>
          210.0.0.0-210.0.15.255

Source          Xlat-Source      Ref.Cnt (F)      TTL(s)

Total IPs in use: 0
Total entries in time-reserve cache: 0
Total freelist left: 4096
```

Dynamic IP and Port NAT Oversubscription

Dynamic IP and Port (DIPP) NAT allows you to use each translated IP address and port pair multiple times (8, 4, or 2 times) in concurrent sessions. This reusability of an IP address and port (known as oversubscription) provides scalability for customers who have too few public IP addresses. The design is based on the assumption that hosts are connecting to different destinations, therefore sessions can be uniquely identified and collisions are unlikely. The oversubscription rate in effect multiplies the original size of the address/port pool to 8, 4, or 2 times the size. For example, the default limit of 64K concurrent sessions allowed, when multiplied by an oversubscription rate of 8, results in 512K concurrent sessions allowed.

The oversubscription rates that are allowed vary based on the platform. The oversubscription rate is global; it applies to the device. This oversubscription rate is set by default and consumes memory, even if you have enough public IP addresses available to make oversubscription unnecessary. You can reduce the rate from the default setting to a lower setting or even 1 (which means no oversubscription). By configuring a reduced rate, you decrease the number of source device translations possible, but increase the DIP and DIPP NAT rule capacities.

If you select **Platform Default**, your explicit configuration of oversubscription is turned off and the default oversubscription rate for the platform applies, as shown in the table below. The **Platform Default** setting allows for an upgrade or downgrade of a software release.

The table below lists the default (highest) oversubscription rate for each platform.

Platform	Default Oversubscription Rate
PA-200	2
PA-500	2
PA-2020	2
PA-2050	2
PA-3020	2
PA-3050	2
PA-3060	2
PA-4020	4
PA-4050	8
PA-4060	8
PA-5020	4
PA-5050	8
PA-5060	8
PA-7050	8
VM-100	1
VM-200	1
VM-300	2
VM-1000-HV	2

Modify the Oversubscription Rate for DIPP NAT

If you have enough public IP addresses that you do not need to use DIPP NAT oversubscription, you can reduce the oversubscription rate and thereby gain more DIP and DIPP NAT rules allowed.

Set NAT Oversubscription		
Step 1	View the DIPP NAT oversubscription rate.	1. Select **Device > Setup > Session > Session Settings.** View the **NAT Oversubscription Rate** setting.
Step 2	Set the DIPP NAT oversubscription rate.	1. Click the Edit icon in the Session Settings section. 2. In the **NAT Oversubscription Rate** drop-down, select **1x**, **2x**, **4x**, or **8x**, depending on which ratio you want. The **Platform Default** setting applies the default oversubscription setting for the platform. If you want no oversubscription, select **1x**. 3. Click **OK** and **Commit** the change.

TCP Session Closing Timers

TCP session closing is enhanced with the renaming of the former TCP Wait timer to the TCP Half Closed timer, and the addition of the TCP Time Wait and Unverified RST timers. No configuration is required; the timers have default values.

TCP Half Closed and TCP Time Wait Timers

In PAN-OS Release 6.1, the TCP connection termination procedure has a TCP Half Closed timer, which is triggered by the first FIN the firewall sees for a session. The timer is named TCP Half Closed because only one side of the connection has sent a FIN. A second timer, TCP Time Wait, is introduced, which is triggered by the second FIN or a RST.

In prior releases, only one TCP wait timer existed, triggered by the first FIN. If the firewall were to have only that one timer, a setting that was too short could prematurely close the half-closed sessions. Conversely, a setting that was too long would make the session table grow too much and possibly use up all of the sessions. Two timers allow you to have a relatively long TCP Half Closed timer and a short TCP Time Wait timer, thereby quickly aging fully closed sessions and controlling the size of the session table.

Figure: TCP Connection Termination Procedure illustrates when the firewall's two timers are triggered during the TCP connection termination procedure.

Figure: TCP Connection Termination Procedure

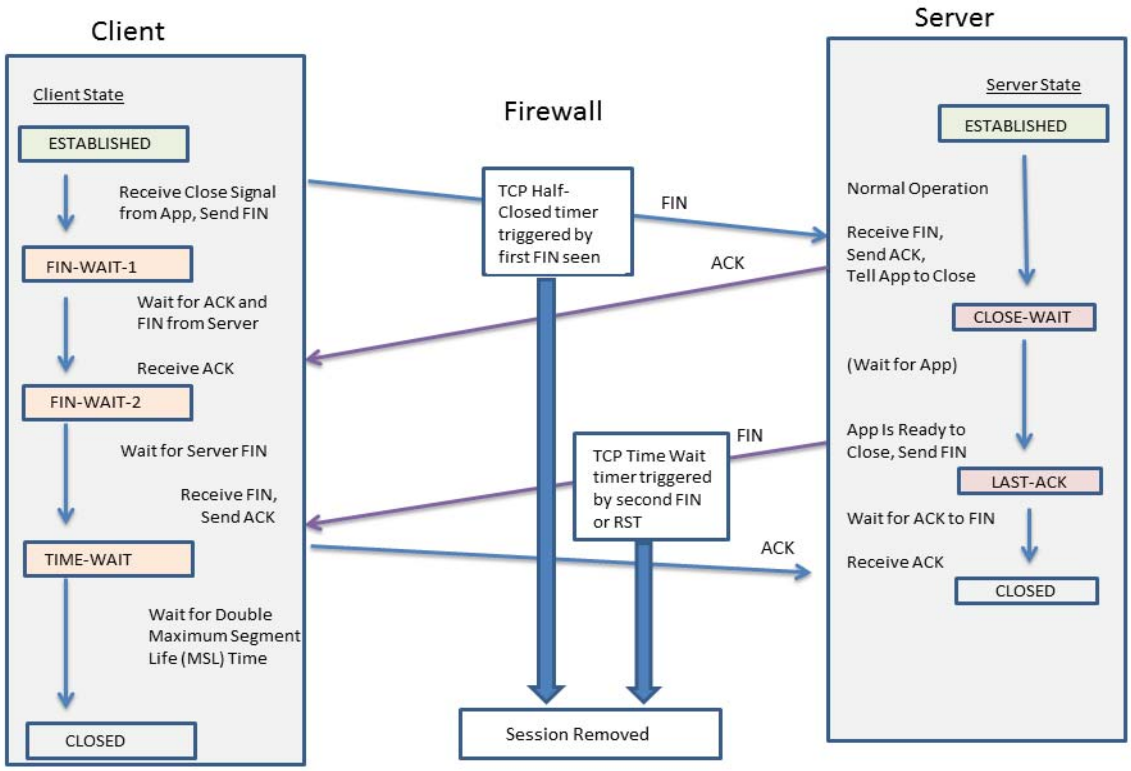

The TCP Time Wait timer should be set to a value less than the TCP Half Closed timer for the following reasons:

- The longer time allowed after the first FIN is seen gives the opposite side of the connection time to fully close the session.
- The shorter Time Wait time is because there is no need for the session to remain open for a long time after the second FIN or a RST is seen. A shorter Time Wait time frees up resources sooner, yet still allows time for the firewall to see the final ACK and possible retransmission of other datagrams.

If you configure a TCP Time Wait timer to a value greater than the TCP Half Closed timer, the commit will be accepted, but in practice the TCP Time Wait timer will not exceed the TCP Half Closed value.

The timers can be set globally or per application. The global settings are used for all applications by default. If you configure TCP wait timers at the application level, they override the global settings.

Unverified RST Timer

If the firewall receives a Reset (RST) packet that cannot be verified (because it has an unexpected sequence number within the TCP window or it is from an asymmetric path), the Unverified RST timer controls the aging out of the session. It defaults to 30 seconds; the range is 1-600 seconds. The Unverified RST timer provides an additional security measure, explained in the second bullet below.

A RST packet will have one of three possible outcomes:

- A RST packet that falls outside the TCP window is dropped.
- A RST packet that falls inside the TCP window but does not have the exact expected sequence number is unverified and subject to the Unverified RST timer setting. This behavior helps prevent denial of service (DoS) attacks where the attack tries to disrupt existing sessions by sending random RST packets to the firewall.
- A RST packet that falls within the TCP window and has the exact expected sequence number is subject to the TCP Time Wait timer setting.

Modify Global TCP Wait Timers or Unverified RST Timer

This task is optional; default timers are in place. The global settings are overridden by TCP Half Closed and TCP Time Wait timers set at the application level, if any.

Modify Global TCP Wait Timers or Unverified RST Timer		
Step 1	Modify the global TCP wait timers or the Unverified RST timer.	1. Select **Device > Setup > Session > Session Timeouts**. 2. Click the Edit icon in the Session Timeouts section. 3. To modify the **TCP Half Closed** timer, highlight the current value and enter a value in the range 1-604800 seconds. The default is 120 seconds. 4. To modify the **TCP Time Wait** timer, highlight the current value and enter a value in the range 1-600 seconds. The default is 15 seconds. Make sure this value is less than the **TCP Half Closed** timer. 5. To modify the **Unverified RST** timer, highlight the current value and enter a value in the range 1-600 seconds. The default is 30 seconds. 6. Click **OK** and **Commit** the change.

Modify Application-Level TCP Wait Timers

These tasks are optional; the global TCP wait timer settings are used for all applications that do not have the timers set at the application level. Any application-based TCP wait timers that are set override the global TCP wait timers. You can configure different timer settings for an application running on different virtual systems.

The steps for a predefined and a custom application differ slightly; both are shown in the following procedure.

Modify Application-Level TCP Wait Timers	
• Modify the application-level TCP wait timers for a predefined application.	1. Select **Objects > Applications** and click the name of the predefined application for which you want to modify timers. 2. In the Options section, to modify the **TCP Half Closed** timer, click **Customize**. In the field (highlight the value if one is present), enter a value in the range 1-604800 seconds. 3. To modify the **TCP Time Wait** timer, click **Customize**. In the field (highlight the value if one is present), enter a value in the range 1-600 seconds. It should be less than the **TCP Half Closed** timer. 4. Click **OK** and **Commit** the change.
• Modify the application-level TCP wait timers for a custom application.	1. Select **Objects > Applications** and click the name of the custom application for which you want to modify timers. 2. On the Advanced tab, to modify the **TCP Half Closed** timer, click in the field (highlight the value if one is present) and enter a value in the range 1-604800 seconds. 3. To modify the **TCP Time Wait** timer, click in the field (highlight the value if one is present) and enter a value in the range 1-600 seconds. It should be less than the **TCP Half Closed** timer. 4. Click **OK** and **Commit** the change.

Session End Reason Logging

The following topics describe session end reasons (for traffic logs) and how to display them:

- Session End Reasons
- Display and Filter Session End Reasons
- Configure a Custom Report with Session End Reasons

Session End Reasons

To troubleshoot connectivity and application availability issues for traffic passing through Palo Alto Networks firewalls, sometimes it helps to know what caused a session to terminate. PAN-OS now provides a *session end reason* field for traffic logs. This field only applies to logs of subtype *end*. For all other subtypes, the value is *not applicable (N/A)*. You can Display and Filter Session End Reasons when monitoring traffic logs. When you configure Syslog and email server profiles, the default traffic log format includes session end reasons. You can also Configure a Custom Report with Session End Reasons.

Sometimes a session can end for multiple reasons. For example, after the firewall drops a session upon detecting a threat, a host might send a TCP FIN message to terminate the same session. Also, a traffic log might record events for multiple sessions (for example, for ping), each with a separate end reason. In these cases, the session end reason field displays only the highest priority reason.

The possible *session end reason* values are as follows, in order of priority (where the first is highest):

Session End Reason	Description
threat	The firewall detected a threat associated with a reset, drop, or block (IP address) action that is defined in a security rule.
policy-deny	The session matched a security policy with a deny or drop action.
tcp-rst-from-client	The client sent a TCP reset to the server.
tcp-rst-from-server	The server sent a TCP reset to the client.
resources-unavailable	The session dropped because of a system resource limitation. For example, the session could have exceeded the number of out-of-order packets allowed per flow or the global out-of-order packet queue.
tcp-fin	One host or both hosts in the connection sent a TCP FIN message to close the session.
tcp-reuse	A session is reused and the firewall closes the previous session.
decoder	The decoder detects a new connection within the protocol (such as HTTP-Proxy) and ends the previous connection.
aged-out	The session aged out.
unknown	This value applies in the following situations: • For logs generated in a PAN-OS release that does not support the session end reason field (releases older than 6.1), the value will be *unknown* after an upgrade to the current PAN-OS release or after the logs are loaded onto the firewall. • In Panorama, logs received from firewalls for which the PAN-OS version does not support session end reasons will have a value of *unknown*.

Display and Filter Session End Reasons

Perform the following steps to display and filter Session End Reasons:

Display and Filter Session End Reasons Logs		
Step 1	Display the traffic logs.	Select **Monitor > Logs > Traffic**. The Session End Reason column is visible by default.
Step 2	Filter the traffic logs by end reason.	1. Create the filter query in one of the following ways: • In the Session End Reason column, click the value of any log that has an end reason you want to use as a query. The value appears in the filter bar above the list. To interpret the query as a negation, type **not** just before the value. You can add multiple queries. The default search type for multiple queries is **and** but you can replace it with **or**. • In the toolbar, click the Add Filter **+** icon. For the **Connector**, select the search type (**and**/**or**). To interpret the query as a negation, select **Negate**. For the **Attribute**, select **Session End Reason**. The **Operator** can be **equal** or **not equal**. For the **Value**, select an end reason. After you click **Add**, the filter bar in the main page displays the new query. Repeat this step for each end reason you want as a query, then click **Close**. 2. In the filter toolbar drop-down, select a time period by which to filter the logs (the default is **Last 24 Hours**). 3. In the filter toolbar, click the Apply Filter → icon. PAN-OS applies the filters to the page.

Configure a Custom Report with Session End Reasons

Perform the following steps to configure a custom report that includes Session End Reasons:

Configure a Custom Report with Session End Reasons		
Step 1	Configure the report parameters.	1. Select **Monitor > Manage Custom Reports** and click **Add**. 2. Enter a **Name** for the report. 3. In the **Database** drop-down, under Detailed Logs, select **Traffic Log**. 4. Optionally, **Group By** the **Session_end_reason** or another column. 5. In the **Available Columns** list, select **Session_end_reason** and any other desired columns, then click the Add **+** icon.
Step 2	To refine the selection criteria by session end reason, configure the **Query Builder** attributes.	Perform the following steps for each desired query rule. 1. For the **Connector**, select the search type (**and**/**or**). To interpret the query as a negation, select **Negate**. 2. For the **Attribute**, select **Session End Reason**. 3. For the **Operator**, select **equal** or **not equal**. 4. For the **Value**, select an end reason. 5. Click **Add**. The query rule appears in the **Query Builder** field.
Step 3	Test and save the report.	1. To test the report settings, click **Run Now**. The report appears in a new tab within the dialog. 2. Modify the settings if necessary, then click **OK** and **Commit**.

Virtualization Features

KVM Support

Kernel-based Virtual Machine (KVM) is an open-source virtualization module for servers running Linux distributions. The VM-Series firewall can be deployed on a Linux server that is running the KVM hypervisor.

System Requirements for VM-Series on KVM

Requirements	Description

Hardware Resources	• vCPU: 2,4,8
	• Memory: 4GB; 5GB for the VM-1000-HV
	• Disk: 40GB
	  ◦ Disk types supported: Virtio and SCSI for best performance; IDE
	  ◦ Disk-controllers: virtio, virt-scsi, IDE
	• Intel-VT or the AMD-V chipset that supports hardware assisted virtualization

Software Versions	• Ubuntu: 12.04 LTS • CentOS/RedHat Enterprise Linux: 6.5 • Open vSwitch: 1.9.3 with bridge compatibility mode

Network Interfaces—Network Interface Cards and Software Bridges	The VM-Series on KVM supports a total of 25 interfaces—1 management interface and a maximum of 24 network interfaces for data traffic.
	VM-Series deployed on KVM supports software-based virtual switches such as the Linux bridge or the Open vSwitch bridge, and direct connectivity to PCI passthrough or an SR-IOV capable adapter.
	• On the Linux bridge and OVS, the e1000 and virtio drivers are supported; the default driver rtl8139 is not supported.
	• For PCI passthrough/SR-IOV support, the VM-Series firewall has been tested for the following network cards:
	  – Intel 82576 based 1G NIC: SR-IOV support on all supported Linux distributions; PCI-passthrough support on all except Ubuntu 12.04 LTS.
	  – Intel 82599 based 10G NIC: SR-IOV support on all supported Linux distributions; PCI-passthrough support on all except Ubuntu 12.04 LTS.
	  – Broadcom 57112 and 578xx based 10G NIC: SR-IOV support on all supported Linux distributions; No PCI-passthrough support.
	  • Drivers: igb; ixgbe; bnx2x
	  • Drivers: igbvf; ixgbevf; bnx2x
	  SR-IOV capable interfaces assigned to the VM-Series firewall, must be configured as Layer 3 interfaces or as HA interfaces.

Options for Attaching the VM-Series on the Network

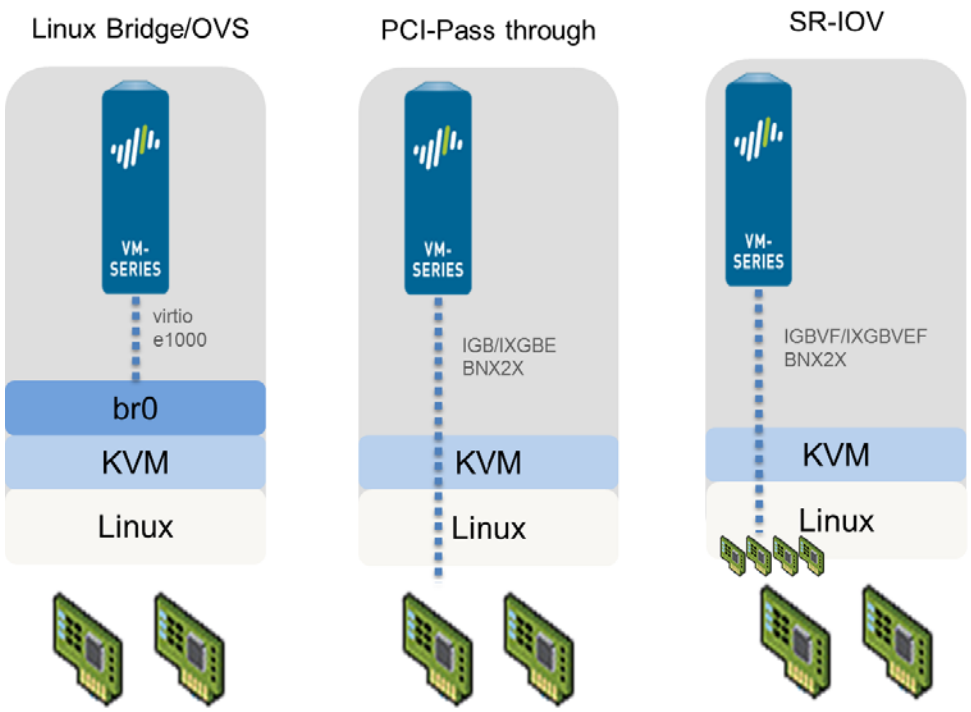

- With a Linux bridge or OVS, data traffic uses the software bridge to connect guests on the same host. For external connectivity, data traffic uses the physical interface to which the bridge is attached.
- With PCI passthrough, data traffic is passed directly between the guest and the physical interface to which it is attached. When the interface is attached to a guest, it is not available to the host or to other guests on the host.
- With SR-IOV, data traffic is passed directly between the guest and the virtual function to which it is attached.

Prerequisites for VM-Series on KVM

Before you install the VM-Series firewall on the Linux server, review the following sections:

▲ Prepare the Linux Server

▲ Prepare to Deploy the VM-Series Firewall on the KVM Hypervisor

Prepare the Linux Server

- Verify the Linux distribution version. For a list of supported versions, see System Requirements for VM-Series on KVM.

- Verify that you have installed and configured KVM tools and packages required for creating and managing virtual machines, for example Libvirt.
- If you want to use a SCSI disk controller to access the disk to which the VM-Series firewall stores data, you must use virsh to attach the virtio-scsi controller to the VM-Series firewall. You can then edit the XML template of the VM-Series firewall to enable the use of the virtio-scsi controller. For instructions, see Enable the Use of a SCSI Controller.

  
  KVM on Ubuntu 12.04 does not support the virtio-scsi controller.
- Verify that the virtualization extensions (VT-d/IOMMU) are enabled in the BIOS. For example, to enable IOMMU, `intel_iommu=on` must be defined in `/etc/grub.conf`. Refer to the documentation provided by your system vendor for instructions.
- Verify that you have set up the networking infrastructure for steering traffic between the guests and the VM-Series firewall and for connectivity to an external server or the Internet. The VM-Series firewall can connect using a Linux bridge, the Open VirtualSwitch, PCI passthrough, or SR-IOV capable network card.
  - Make sure that the link state for all interfaces you plan to use are up, sometimes you have to manually bring them up.
  - Verify the PCI ID of all the interfaces. To view the list, use the command: `Virsh nodedev-list –tree`
  - If using a Linux bridge or OVS, verify that you have set up the bridges required to send/receive traffic to/from the firewall. If not, create bridge(s) and verify that they are up before you begin installing the firewall.
  - If using PCI-passthrough or SR-IOV, verify that the virtualization extensions (VT-d/IOMMU) are enabled in the BIOS. For example, to enable IOMMU, `intel_iommu=on` must be defined in `/etc/grub.conf`. Refer to the documentation provided by your system vendor for instructions.
  - If using PCI-passthrough, ensure that the VM-Series firewall has exclusive access to the interface(s) that you plan to attach to it.

    To allow exclusive access, you must manually detach the interface(s) from the Linux server; Refer to the documentation provided by your network card vendor for instructions.

    To manually detach the interface(s) from the server., use the command:

    `Virsh nodedev-detach <pci id of interface>, for example, pci_0000_07_10_0`

    In some cases, in `/etc/libvirt/qemu.conf`, you may have to uncomment `relaxed_acs_check = 1`.
  - If using SR-IOV, verify that the virtual function capability is enabled for each port that you plan to use on the network card. With SR-IOV, a single Ethernet port (physical function) can be split into multiple virtual functions. A guest can be mapped to one or more virtual functions.

To enable virtual functions, you need to:

1.Create a new file in this location: /etc/modprobe.d/

2. Modify the file using the vi editor to make the functions persistent: vim /etc/modprobe.d/igb.conf

3. Enable the number of number of virtual functions required: options igb max_vfs=4

After you save the changes and reboot the Linux server, each interface (or physical function) in this example will have 4 virtual functions.

Refer to the documentation provided by your network vendor for details on the actual number of virtual functions supported and for instructions to enable it.

Prepare to Deploy the VM-Series Firewall on the KVM Hypervisor

- Purchase the VM-Series model and register the authorization code on the Palo Alto Networks support portal.
- Obtain the qcow2 image and save it on the Linux server. By default, the image is stored in this folder: `/var/lib/libvirt/qemu/images`

  If you plan to deploy more than one instance of the VM-Series firewall, make the required number of copies of the image. Because each instance of the VM-Series firewall maintains a link with the .qcow2 image that was used to deploy the firewall, to prevent any data corruption issues ensure that each image is independent and is not shared with more than one instance of the firewall.

Supported Deployments

The firewall protects traffic across the guests within a Linux server. The VM-Series firewall can be deployed with virtual wire, Layer 2, or Layer 3 interfaces and is placed on the same Linux bridge that connects the guests on the server.

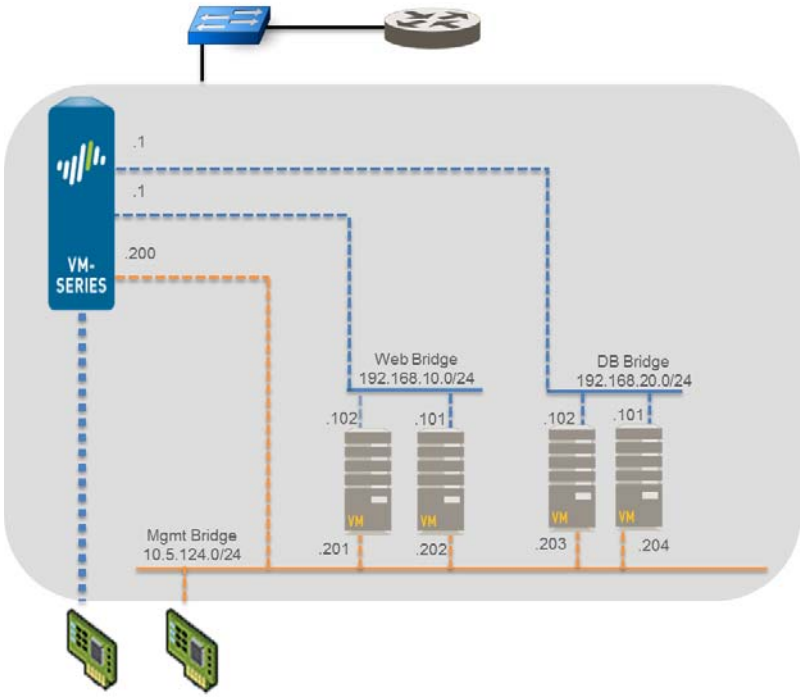

A pair of VM-Series firewalls are deployed in a high availability setup using Layer 3 interfaces; in the illustration an SR-IOV capable adapter is used on the Linux server that hosts the firewalls. The active peer in the HA pair secures traffic that is routed to it from guests deployed on a different Linux server.

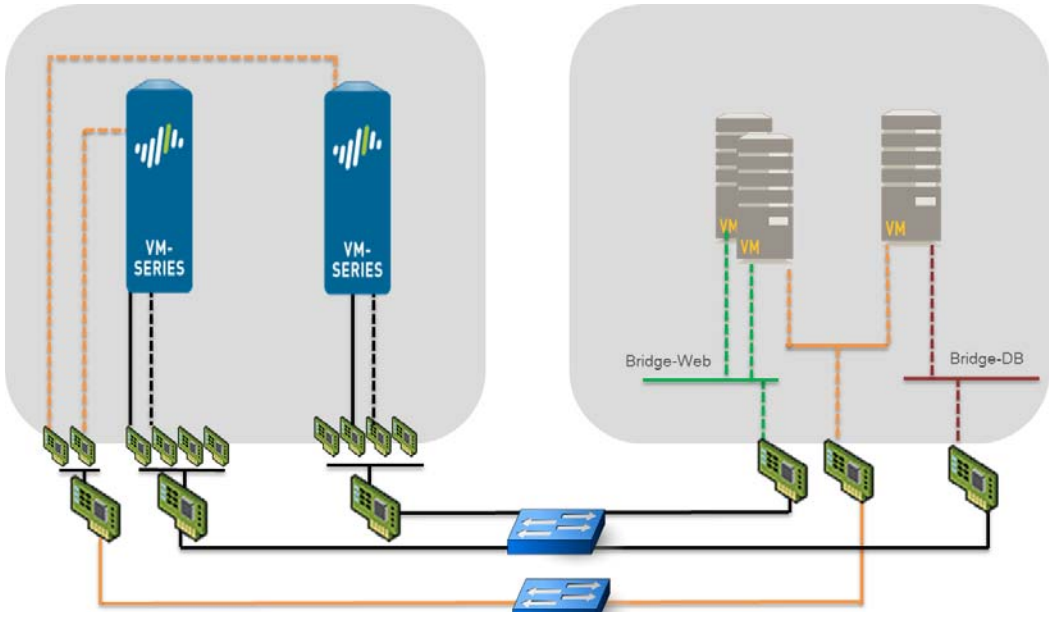

Install the VM-Series Firewall on KVM

The libvirt API that is used to manage KVM includes a host of tools that allow you to create and manage virtual machines. To install the VM-Series firewall on KVM you can use any of the following methods:

- Manually create the XML definition of the VM-Series firewall, then use virsh to import the definition. Virsh is the most powerful tool that allows for full administration of the virtual machine.
- Use virt-install to create the definition for the VM-Series firewall and install it.
- Use the desktop user interface called virt-manager; virt-manager provides a convenient wizard to help you through the installation process.

The following procedure uses virt-manager to install the VM-Series firewall on a server running KVM on RHEL; the instructions for using virsh or virt-install are not included in this document.

Install the VM-Series on KVM	
Step 1 Install the VM-Series firewall.	1. On the Virt-manager, select **Create a new virtual machine**. 2. Add a descriptive **Name** for the VM-Series firewall. 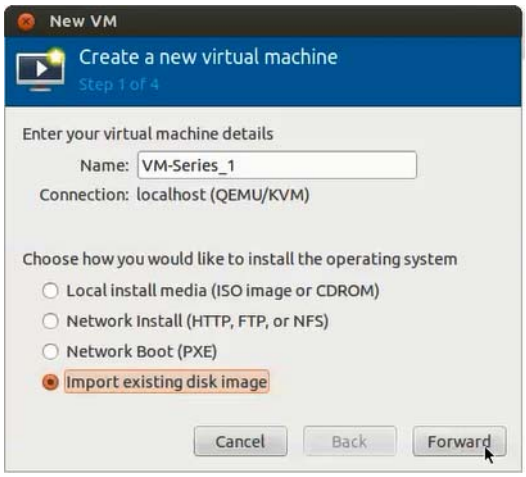 3. Select **Import existing disk image**, browse to the image, and set the **OS Type**: Linux and **Version**: Red Hat Enterprise Linux 6. If you prefer, you can leave the OS Type and Version as Generic. 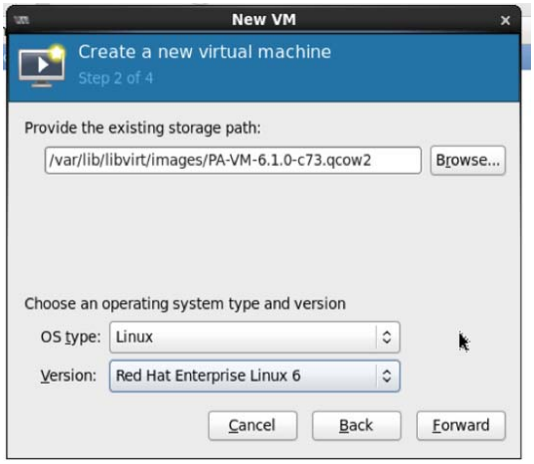 4. Set the **Memory** to 4096 MB; or 5120 MB, if you have purchased the VM-1000-HV license. 5. Set **CPU** to 2,4 or 8.

Install the VM-Series on KVM (Continued)

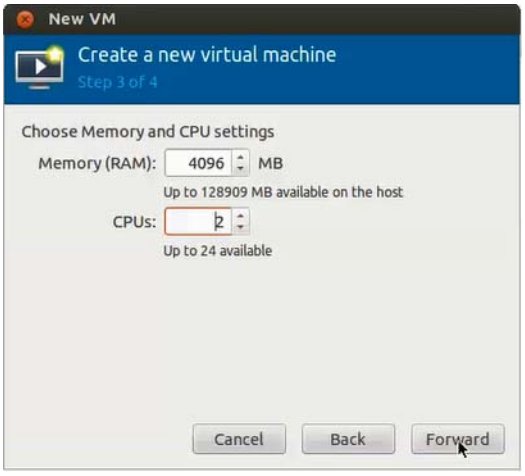

6. Select **Customize configuration before install**.
7. Under Advanced options, select the bridge for the management interface, and accept the default settings.

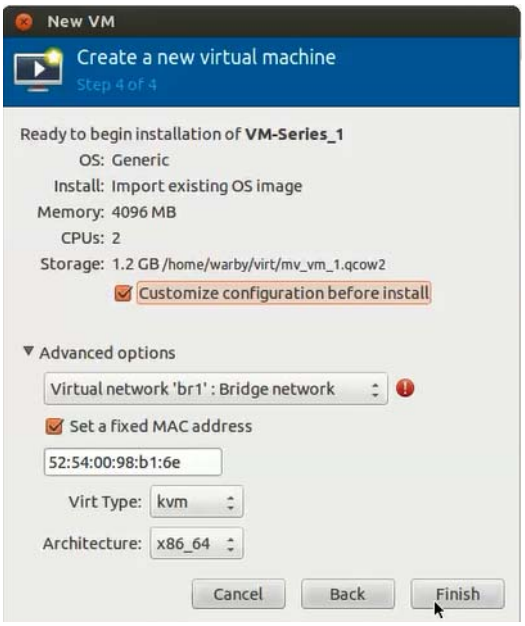

Install the VM-Series on KVM (Continued)

8. To modify disk settings:
   a. Select **Disk**, expand Advanced options, and select **Storage format**—qcow2; **Disk Bus**—IDE or Virtio, based on your setup.

      If If you want to use a SCSI disk bus, see Enable the use of a SCSI Controller.

   b. Expand Performance options, and set **Cache mode** to writethrough. This setting improves installation time and execution speed on the VM-Series firewall.

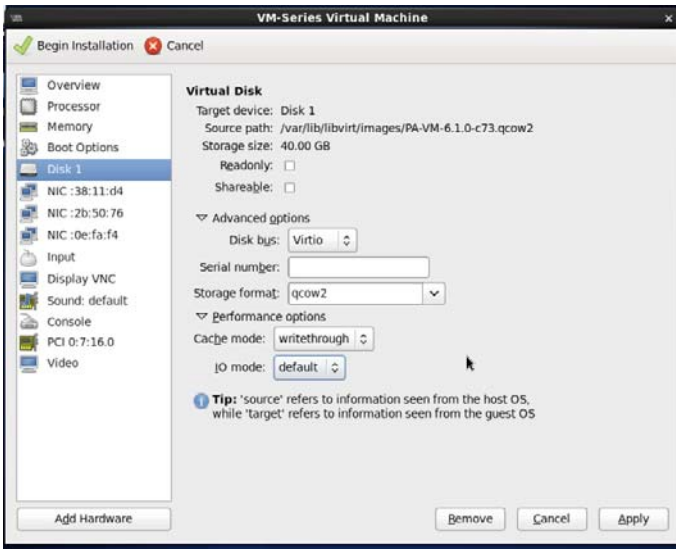

Install the VM-Series on KVM (Continued)

9. To add network adapters for the data interfaces:
   a. Select **Add Hardware > Network** if you are using a software bridge such as the Linux bridge or the Open Virtual Switch.
      - For **Host Device**, enter the name of the bridge or select it from the drop-down.
      - To specify the driver, set **Device Model** to e-1000 or Virtio. These are the only supported virtual interface types.

   b. Select **Add Hardware** > **PCI Host Device** for PCI-passthrough or an SR-IOV capable device.

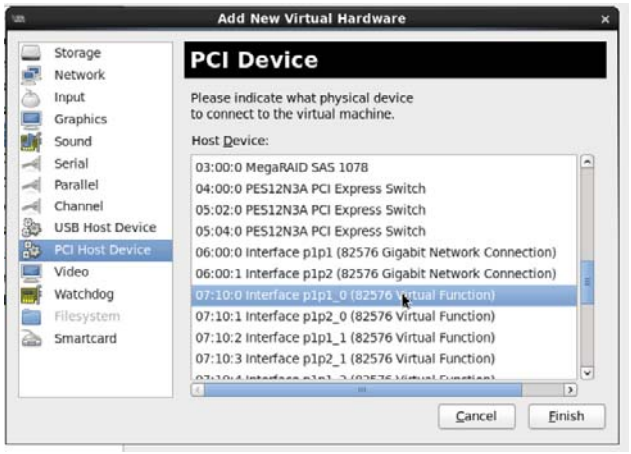

      - In the **Host Device** list, select the interface on the card or the virtual function.
   c. Click **Finish**.
10. Click **Begin Installation** 

.

Install the VM-Series on KVM (Continued)

By default, the XML template for the VM-Series firewall is created and stored at `etc/libvirt/qemu`.	11. Wait 5-7 minutes for the installation to complete. 
Step 2 Configure the network access settings for the management interface.	1. Open a connection to the console. 2. Log in to the firewall with username/password: admin/admin. 3. Enter the following command: `set deviceconfig system ip-address <Firewall-IP> netmask <netmask> default-gateway <gateway-IP> dns-setting servers primary <DNS-IP>` where `<Firewall-IP>` is the IP address you want to assign to the management interface, `<netmask>` is the subnet mask, `<gateway-IP>` is the IP address of the network gateway, and `<DNS-IP>` is the IP address of the DNS server.
Step 3 Verify which ports on the PCI device are mapped to the interfaces on the VM-Series firewall.	To make sure that traffic is handled by the correct interface, use the following command to identify which ports on the PCI device are mapped to the ports on the VM-Series firewall. `admin@PAN-VM>` **`debug show vm-series interfaces all`** Phoenix_interface Base-OS_port Base-OS_MAC PCI-ID mgt eth0 52:54:00:d7:91:52 0000:00:03.0 Ethernet1/1 eth1 52:54:00:fe:8c:80 0000:00:06.0 Ethernet1/2 eth2 0e:c6:6b:b4:72:06 0000:00:07.0 Ethernet1/3 eth3 06:1b:a5:7e:a5:78 0000:00:08.0 Ethernet1/4 eth4 26:a9:26:54:27:a1 0000:00:09.0 Ethernet1/5 eth5 52:54:00:f4:62:13 0000:00:10.0
Step 4 Access the web interface of the VM-Series firewall and configure the interfaces and define security rules and NAT rules to safely enable the applications that you want to secure.	Refer to the PAN-OS Administrator's Guide.

Amazon AWS Support

The VM-Series firewall can be deployed in the Amazon Web Services (AWS) cloud. It can then be configured to secure access to the applications that are deployed on EC2 instances and placed into a Virtual Private Cloud (VPC) in AWS.

About the VM-Series Firewall in AWS

The Amazon Web Service (AWS) is a public cloud service that enables you to run your applications on a shared infrastructure managed by Amazon. These applications can be deployed on scalable computing capacity or EC2 instances in different AWS regions and accessed by users over the Internet. For networking consistency and ease of management of EC2 instances, Amazon offers the Virtual Private Cloud (VPC). A VPC is apportioned from the AWS public cloud, and is assigned a CIDR block from the private network space (RFC 1918). Within a VPC, you can carve public/private subnets for your needs and deploy the applications on EC2 instances within those subnets. You can then safely enable access to the applications within the VPC by deploying the VM-Series firewall on an EC2 instance. The VM-Series firewall can then be configured to secure traffic to and from the EC2 instances within the VPC.

This document assumes that you are familiar with the networking and configuration of the AWS VPC. In order to provide context for the terms used in this section, here is a brief refresher on the AWS terms (some definitions are taken directly from the AWS glossary) that are referred to in this document:

Term	Description
EC2	Elastic Compute Cloud A web service that enables you to launch and manage Linux/UNIX and Windows server instances in Amazon's datacenters.
AMI	Amazon Machine Image An AMI provides the information required to launch an instance, which is a virtual server in the cloud. The VM-Series AMI is an encrypted machine image that includes the operating system required to instantiate the VM-Series firewall on an EC2 instance.
Instance Type	Amazon-defined specifications that stipulate the memory, CPU, storage capacity, and hourly cost for an instance. Some instance types are designed for standard applications, whereas others are designed for CPU-intensive, memory-intensive applications, and so on.

Term	Description
ENI	Elastic Network Interface An additional network interface that can be attached to an EC2 instance. ENIs can include a primary private IP address, one or more secondary private IP addresses, a public IP address, an elastic IP address (optional), a MAC address, membership in specified security groups, a description, and a source/destination check flag.
IP Address Types for EC2 Instances	An EC2 instance can have different types of IP addresses. • Public IP address: An IP address that can be routed across the Internet. • Private IP address: A IP address in the private IP address range as defined in the RFC 1918. You can choose to manually assign an IP address or to auto assign an IP address within the range in the CIDR block for the subnet in which you launch the EC2 instance. If you are manually assigning an IP address, Amazon reserves the first four (4) IP addresses and the last one (1) IP address in every subnet for IP networking purposes. • Elastic IP address (EIP): A static IP address that you have allocated in Amazon EC2 or Amazon VPC and then attached to an instance. Elastic IP addresses are associated with your account, not with a specific instance. They are elastic because you can easily allocate, attach, detach, and free them as your needs change. An instance in a public subnet can have a Private IP address, a Public IP address, and an Elastic IP address (EIP); an instance in a private subnet will have a private IP address and optionally have an EIP.
VPC	Virtual Private Cloud An elastic network populated by infrastructure, platform, and application services that share common security and interconnection.
IGW	Internet gateway provided by Amazon. Connects a network to the Internet. You can route traffic for IP addresses outside your VPC to the Internet gateway.
Subnets	A segment of the IP address range of a VPC to which EC2 instances can be attached. EC2 instances are grouped into subnets based on your security and operational needs. There are two types of subnets: • Private subnet: The EC2 instances in this subnet cannot be reached from the Internet. • Public subnet: The Internet gateway is attached to the public subnet, and the EC2 instances in this subnet can be reached from the Internet.
Security Groups	A security group is attached to an ENI and it specifies the list of protocols, ports, and IP address ranges that are allowed to establish inbound/outbound connections on the interface.  In the AWS VPC, security groups and network ACLs control inbound and outbound traffic; security groups regulate access to the EC2 instance, while network ACLs regulate access to the subnet. Because you are deploying the VM-Series firewall, set more permissive rules in your security groups and network ACLs and allow the firewall to safely enable applications in the VPC.
Route Tables	A set of routing rules that controls the traffic leaving any subnet that is associated with the route table. A subnet can be associated with only one route table.

Term	Description
Key Pair	A set of security credentials you use to prove your identity electronically. The key pair consists of a private key and a public key. When you launch the VM-Series firewall, you must either generate a key pair or select an existing key pair for the VM-Series firewall. The private key is required to access the firewall in maintenance mode.

Deployments Supported in AWS

The VM-Series firewall secures inbound and outbound traffic to and from EC2 instances within the AWS Virtual Private Cloud (VPC). Because the AWS VPC only supports an IP network (Layer 3 networking capabilities), the VM-Series firewall can only be deployed with Layer 3 interfaces.

- Deploy the VM-Series firewall to secure the EC2 instances hosted in the AWS Virtual Private Cloud.
  If you host your applications in the AWS cloud, deploy the VM-Series firewall to protect and safely enable applications for users who access these applications over the Internet. For example, the following diagram shows a VM-Series firewall deployed in the Edge subnet to which the Internet gateway is attached. The application(s) are deployed in the private subnet, which does not have direct access to the Internet.

  When users need to access the applications in the private subnet, the firewall receives the request and directs it to the appropriate application, after verifying security policy and performing Destination NAT. On the return path, the firewall receives the traffic, applies security policy and uses Source NAT to deliver the content to the user. See Use Case: Secure the EC2 Instances in the AWS Cloud.

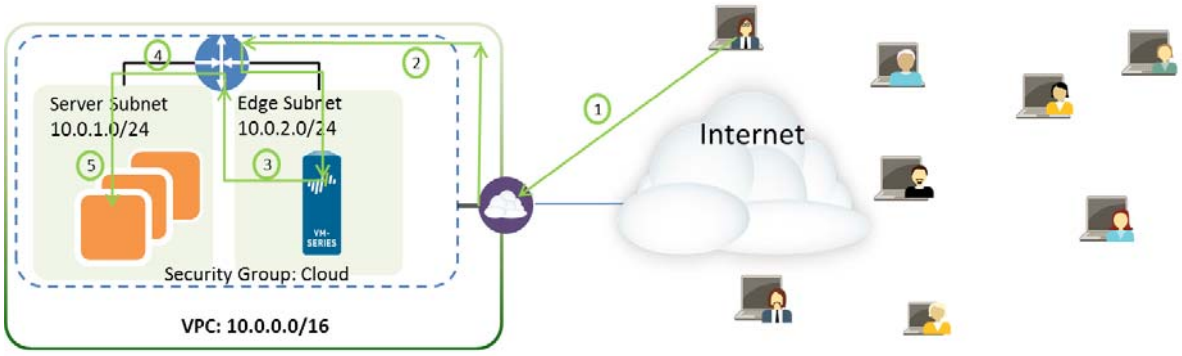

- Deploy the VM-Series firewall for VPN access between the corporate network and the EC2 instances within the AWS Virtual Private Cloud.
  To connect your corporate network with the applications deployed in the AWS Cloud, you can configure the firewall as a termination point for an IPSec VPN tunnel. This VPN tunnel allows users on your network to securely access the applications in the cloud.

  For centralized management, consistent enforcement of policy across your entire network, and for centralized logging and reporting, you can also deploy Panorama in your corporate network. If you need to set up VPN access to multiple VPCs, using Panorama allows you to group the firewalls by region and administer them with ease.

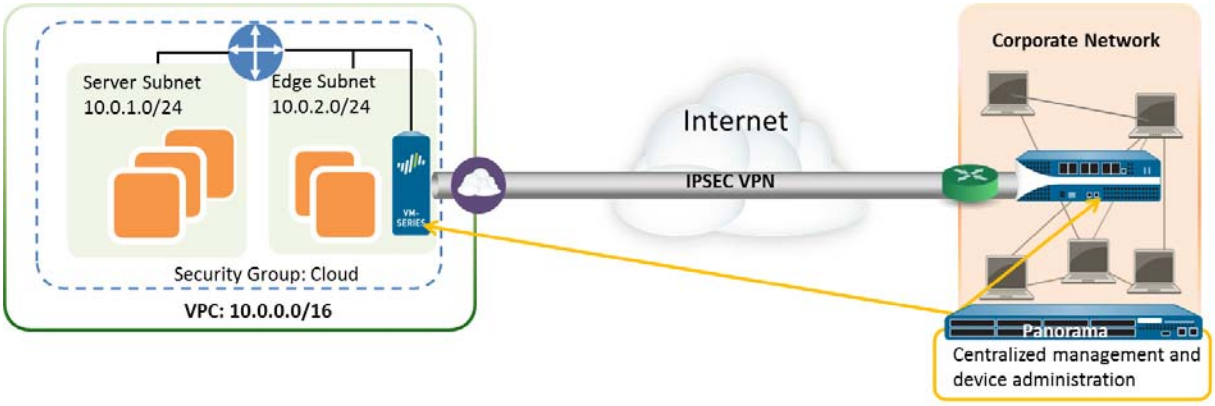

For information on using the VM-Series firewall as a GlobalProtect gateway in AWS, to secure mobile users, refer to the Virtualization Deployment Guide.

Deploy the VM-Series Firewall on AWS

- Obtain the AMI
- Review System Requirements and Limitations for VM-Series on AWS
- Planning Worksheet for the VM-Series in the AWS VPC
- Launch the VM-Series Firewall in AWS

Obtain the AMI

The AMI for the VM-Series firewall is available in the AWS Marketplace with the Bring Your Own License (BYOL) pricing option. For purchasing licenses contact your Palo Alto Networks sales engineer or reseller.

Review System Requirements and Limitations for VM-Series on AWS

Requirement	Details
EC2 instance types	Deploy the VM-Series firewall on any of the following EC2 instance types: • m3.xlarge • m3.2xlarge • c3.xlarge • c3.2xlarge • c3.4xlarge • c3.8xlarge The minimum resource requirements for the VM-Series firewall are: vCPU: 2; Memory: 4GB; 5GB for the VM-1000-HV; Disk: 40GB. If you deploy the VM-Series firewall on an EC2 instance type that does not meet these requirements, the firewall will boot into maintenance mode.
Amazon Elastic Block Storage (EBS)	The VM-Series firewall must use the Amazon Elastic Block Storage (EBS) volume for storage. EBS optimization provides an optimized configuration stack and additional, dedicated capacity for Amazon EBS I/O.
Networking	Because the AWS only supports Layer 3 networking capabilities, the VM-Series firewall can only be deployed with Layer 3 interfaces. Layer 2 interfaces, virtual wire, VLANs, and subinterfaces are not supported on the VM-Series firewall deployed in the AWS VPC.
Interfaces	Support for a total of eight interfaces is available— one management interface and a maximum of seven Elastic Network Interfaces (ENIs) for data traffic. The VM-Series firewall does not support hot attachment of ENIs; to detect the addition or removal of an ENI you must reboot the firewall. Your EC2 instance type selection determines the total number of ENIs you can enable. For example, the c3.8xlarge supports 8 ENIs.
Support entitlement and Licenses	A support account and a valid VM-Series license are required to obtain the Amazon Machine Image (AMI) file, which is required to install the VM-Series firewall in the AWS VPC. The licenses required for the VM-Series firewall— capacity license, support, and subscriptions for Threat Prevention, URL Filtering, WildFire, etc—must be purchased from Palo Alto Networks. To purchase the licenses for your deployment, contact your sales representative.

Planning Worksheet for the VM-Series in the AWS VPC

For ease of deployment, plan the subnets within the VPC and the EC2 instances that you want to deploy within each subnet. Before you begin, use the following table to collate the network information required to deploy and insert the VM-Series firewall into the traffic flow in the VPC:

Configuration Item	Value
VPC CIDR	
Security Groups	
Subnet (public) CIDR	
Subnet (private) CIDR	
Subnet (public) Route Table	
Subnet (private) Route Table	
Security Groups: • Rules for Management Access to the firewall (eth0/0) • Rules for access to the dataplane interfaces of the firewall • Rules for access to the interfaces assigned to the application servers.	
EC2 Instance 1 (VM-Series firewall)  An EIP is only required for the dataplane interface that is attached to the public subnet.	Subnet: Instance Type: Mgmt interface IP: Mgmt interface EIP: Dataplane Interface eth1/1 • Private IP: • EIP (if required): • Security Group: Dataplane interface eth1/2 • Private IP: • EIP (if required): • Security Group:
EC2 Instance 2 (Application to be secured) Repeat these set of values for additional application(s) being deployed.	Subnet: Instance Type: Mgmt Interface IP: Default Gateway: Dataplane Interface 1 • Private IP:

Launch the VM-Series Firewall in AWS

If you have not already registered the capacity auth-code that you received with the order fulfillment email, with your support account, register it now. After completing the registration process, deploy the VM-Series firewall by launching it in the AWS VPC as follows:

Launch the VM-Series Firewall in the AWS VPC		
Step 1	Access the AWS Console.	Log in to the AWS console and select the EC2 Dashboard.
Step 2	Set up the VPC for your network needs. Whether you launch the VM-Series firewall in an existing VPC or you create a new VPC, the VM-Series firewall must be able to receive traffic from the EC2 instances and perform inbound and outbound communication between the VPC and the Internet. Refer to the AWS VPC documentation for instructions on creating a VPC and setting it up for access. For an example with a complete workflow, see Use Case: Secure the EC2 Instances in the AWS Cloud.	1. Create a new VPC or use an existing VPC. 2. Verify that the network and security components are defined suitably. • Enable communication to the Internet. The default VPC includes an Internet gateway, and if you install the VM-Series firewall in the default subnet it has access to the Internet. • Create subnets. Subnets are segments of the IP address range assigned to the VPC in which you can launch the EC2 instances. The VM-Series firewall must belong to the public subnet so that it can be configured to access the Internet. • Create security groups as needed to manage inbound and outbound traffic from the EC2 instances/subnets. • Add routes to the route table for a private subnet to ensure that traffic can be routed across subnets and security groups in the VPC, as applicable.

Launch the VM-Series Firewall in the AWS VPC (Continued)	
Step 3 Launch the VM-Series firewall.  Although you can add additional network interfaces to the VM-Series firewall when you launch, to minimize any issues with deploying the instance attach additional network interfaces after you launch the firewall.  This key pair is required for first time access to the firewall. It is also required to access the firewall in maintenance mode.	1. On the EC2 Dashboard, click **Launch Instance**. 2. Select the VM-Series AMI. To get the AMI, see Obtain the AMI. 3. Launch the VM-Series firewall on an EC2 instance. a. Choose the **EC2 instance type**—m3.xlarge, c3.xlarge, or c3.8xlarge—for allocating the resources required for the firewall, and click **Next**. b. Select the VPC. c. Select the public subnet to which the VM-Series management interface will attach. d. Select **Automatically assign a public IP address**. This allows you to obtain a publicly accessible IP address for the management interface of the VM-Series firewall. You can later attach an Elastic IP address to the management interface; unlike the public IP address that is disassociated from the firewall when the instance is terminated, the Elastic IP address provides persistence and can be reattached to a new (or replacement) instance of the VM-Series firewall without the need to reconfigure the IP address wherever you might have referenced it. e. Select **Launch as an EBS-optimized instance**. f. Accept the default **Storage** settings. g. Skip Tagging. You can add tags later. h. Select an existing **Security Group** or create a new one. This security group is for restricting access to the management interface of the firewall. i. If prompted, select an appropriate **SSD** option for your setup. j. Select **Review and Launch**. Review that your selections are accurate and click **Launch**. k. Select an existing key pair or create a new one, and acknowledge the key disclaimer. l. Download and save the private key to a safe location. You cannot regenerate this key, if lost. It takes 5-7 minutes to launch the VM-Series firewall. You can view the progress on the EC2 Dashboard.When the process completes, the VM-Series firewall displays on the **Instances** page of the EC2 Dashboard.

Launch the VM-Series Firewall in the AWS VPC (Continued)	
Step 4 Configure a new administrative password for the firewall. Using the CLI, you must configure a unique password for the firewall before you can access the web interface of the firewall.	1. Use public IP address to SSH into the Command Line Interface (CLI) of the VM-Series firewall. You will need the private key that you used or created in Step 3-k to access the CLI. 2. Enter the following command to log in to the firewall: **`ssh-i`** *<private_key_name>* **`admin@`***<public-ip_address>* 3. Configure a new password, using the following command and follow the onscreen prompts: **`configure`** **`set mgt-config users admin password`** **`commit`** 4. Terminate the SSH session.
Step 5 Shut down the VM-Series firewall.	1. On the EC2 Dashboard, select **Instances**. 2. From the list, select the VM-Series firewall and click **Actions > Stop**.
Step 6 Create virtual network interface(s) and attach the interface(s) to the VM-Series firewall. The virtual network interfaces are called Elastic Network Interfaces (ENIs) in AWS, and serve as the dataplane network interfaces on the firewall. These interfaces are used for handling data traffic to/from the firewall. You will need at least two ENIs. You can add up to seven ENIs to handle data traffic on the VM-Series firewall; check your EC2 instance type to verify the maximum number supported on it. To detect the newly attached ENIs, the VM-Series firewall requires a reboot. If you have not shutdown the firewall yet, proceed with the license activation process, which triggers a reboot on the firewall. When the firewall reboots, the ENIs will be detected.	1. On the EC2 Dashboard, select **Network Interfaces**, and click **Create Network Interface**. 2. Enter a descriptive name for the interface. 3. Select the subnet. Use the subnet ID to make sure that you have selected the correct subnet. You can only attach an ENI to an instance in the same subnet. 4. Enter the **Private IP** address to assign to the interface or select **Auto-assign** to automatically assign an IP address within the available IP addresses in the selected subnet. 5. Select the **Security group** to control access to the dataplane network interface. 6. Click **Yes, Create**.  7. To attach the ENI to the VM-Series firewall, select the interface you just created, and click **Attach**.  8. Select the **Instance ID** of the VM-Series firewall, and click **Attach**. 9. Repeat the steps above for creating and attaching at least one more ENI to the firewall.

Launch the VM-Series Firewall in the AWS VPC (Continued)		
Step 7	Activate the licenses on the VM-Series firewall. This task is not performed on the AWS management console. Access to the Palo Alto Networks support portal and the web interface of the VM-Series firewall is required for license activation.	See Licensing.
Step 8	Disable Source/Destination check on every firewall dataplane network interface(s). Disabling this attribute allows the interface to handle network traffic that is not destined to the IP address assigned to the network interface.	1. On the EC2 Dashboard, select the network interface, for example eth1/1, in the **Network Interfaces** tab. 2. In the **Action** drop-down, select **Change Source/Dest. Check** 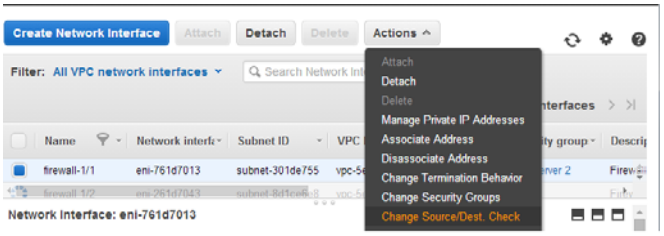 3. Click **Disabled** and **Save** your changes. 4. Repeat these steps for each firewall dataplane interface.

Launch the VM-Series Firewall in the AWS VPC (Continued)	
Step 9 Configure the dataplane network interfaces as Layer 3 interfaces on the firewall. For an example configuration, see Step 14 through Step 17 in Use Case: Secure the EC2 Instances in the AWS Cloud. On the application servers within the VPC, define the dataplane network interface of the firewall as the default gateway.	1. Launch a web browser and enter the Public IP address or EIP assigned to the management interface of the VM-Series firewall. 2. Select **Network > Interfaces > Ethernet**. Before you begin setting up the network interfaces, verify that the Link state for the interface is up. . If the link state is not up, reboot the firewall. 3. Click the link for **ethernet 1/1** and configure as follows: – **Interface Type**: **Layer3** – On the **Config** tab, assign the interface to the default router. – On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. Define a new zone, for example VM_Series_untrust, and then click **OK**. – On the **IPv4** tab, select **Static** or **DHCP Client**. If using **Static** addresses, click **Add** in the IP section, and enter the IP address and network mask for the interface, for example 10.0.0.10/24. Make sure that the IP address matches the ENI IP address that you assigned earlier. If using DHCP, select the **DHCP Client** option; the private IP address that you assigned to the ENI in the AWS management console will be automatically acquired. 4. Click the link for **ethernet 1/2** and configure as follows: – **Interface Type**: **Layer3** – **Security Zone: VM_Series_trust** – **IP address**: Select **Static** or **DHCP Client**. If you selected **Static**, click **Add** in the IP section, and enter the IP address and network mask for the interface. Make sure that the IP address matches the attached ENI IP address that you assigned earlier.
	For DHCP, clear the **Automatically create default route to default gateway provided by server** check box. For an interface that is attached to the private subnet in the VPC, disabling this option ensures that traffic handled by this interface does not flow directly to the Internet gateway on the VPC. 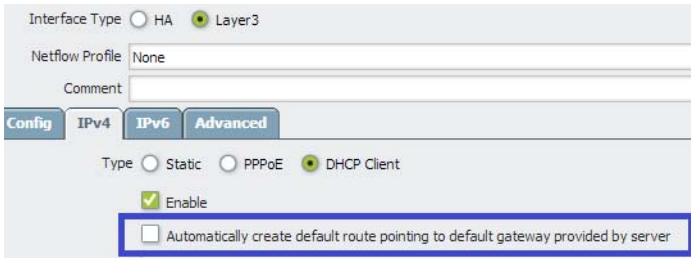

Launch the VM-Series Firewall in the AWS VPC (Continued)	
Step 10 Create NAT rules to allow inbound and outbound traffic from the servers deployed within the VPC	1. From the firewall web interface, select **Policies > NAT**. 2. Create a NAT rule to allow traffic from the dataplane network interface on the firewall to the web server interface in the VPC. 3. Create a NAT rule to allow outbound access for traffic from the web server to the Internet.
Step 11 Create security policies to allow/deny traffic to/from the servers deployed within the VPC.	1. From the firewall web interface, select **Policies > Security**. 2. Click **Add**, and specify the zones, applications and logging options that you would like to execute to restrict and audit traffic traversing through the network.
Step 12 Save the changes to the firewall configuration.	Click **Commit**.
Step 13 Verify that the VM-Series firewall is securing traffic and that the NAT rules are in effect.	1. From the firewall web interface, select **Monitor > Logs > Traffic**. 2. View the logs to make sure that the applications traversing the network match the security policies you implemented.

List of Attributes Monitored on the AWS VPC

The following attributes (or tag names) are available as match criteria for dynamic address groups.

Attribute	Format
Architecture	Architecture.<Architecture string>
Guest OS	GuestOS.<guest OS name>
Image ID	ImageId.<ImageId string>
Instance ID	InstanceId.<InstanceId string>
Instance State	InstanceState.<instance state>
Instance Type	InstanceType.<instance type>
Key Name	KeyName.<KeyName string>
Placement—Tenancy, Group Name, Availability	Placement.Tenancy.<string> Placement.GroupName.<string> Placement.AvailabilityZone.<string>
Private DNS Name	PrivateDnsName.<Private DNS Name>
Public DNS Name	PublicDnsName.<Public DNS Name>
Subnet ID	SubnetID.<subnetID string>
Tag (key, value)	aws-tag.<key>.<value> Maximum of 5 of these tags are supported per instance
VPC ID	VpcId.<VpcId string>

VM Information Sources

With Amazon AWS Support, in addition to the VMware ESXi server and the VMware vCenter server, the firewall can monitor a VPC and retrieve changes in the EC2 instance inventory within the VPC.

Enable VM Monitoring in the AWS VPC

Step 1 Configure the firewall to monitor the VPC.

To see how VM Monitoring can be used with dynamic address groups to secure virtual machines, see Use Case: Use Dynamic Address Groups to Secure New EC2 Instances within the VPC.

1. Select **Device > VM Information Sources**.
2. Click **Add** and enter the following information:
   a. A **Name** to identify the VPC that you want to monitor. For example, VPC-CloudDC.
   b. Set the **Type** to AWS VPC.
   c. In **Source**, enter the URI for the VPC. The syntax is ec2.<*your_region*>.amazonaws.com
   d. Add the credentials required for the firewall to digitally sign API calls made to the AWS services. You need the following:
      - **Access Key ID**: Enter the alphanumeric text string that uniquely identifies the user who owns or is authorized to access the AWS account.
      - **Secret Access Key**: Enter the password and confirm your entry.
   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.

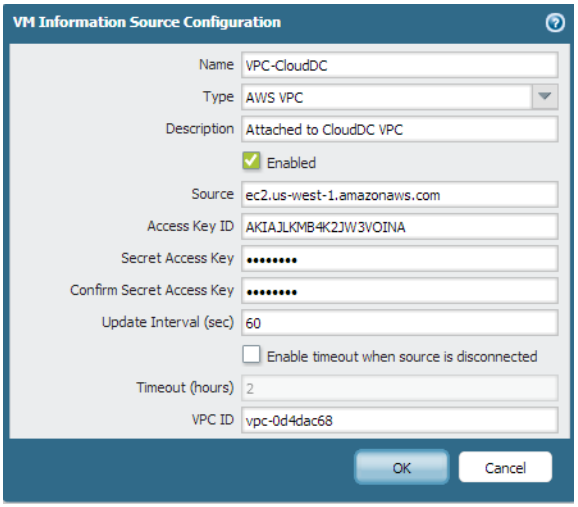

   f. Enter the **VPC ID** that is displayed on the VPC Dashboard in the AWS management console.
   g. Click **OK**, and **Commit** the changes.
   h. Verify that the connection **Status** is connected .

