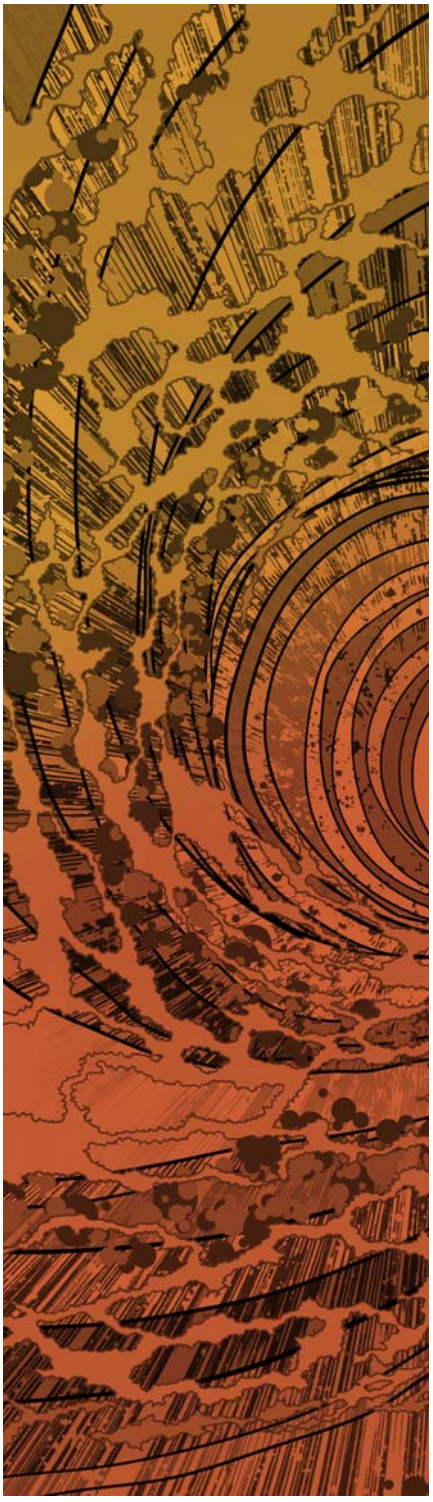

paloalto NETWORKS®

# PAN-OS® New Features Guide

Version 7.0

## Contact Information

Corporate Headquarters:

Palo Alto Networks

4401 Great America Parkway

Santa Clara, CA 95054

www.paloaltonetworks.com/company/contact-us

## About this Guide

This guide describes how to use the new features introduced in PAN-OS 7.0. 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 7.0 release notes, go to https://www.paloaltonetworks.com/documentation/70/pan-os/pan-os-release-notes.html.

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

Palo Alto Networks, Inc.
www.paloaltonetworks.com
© 2016 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at http://www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

**Revision Date: June 8, 2016**

# Table of Contents


Upgrade to PAN-OS 7.0 . . . . . 7
Upgrade/Downgrade Considerations . . . . . 8
Upgrade the Firewall to PAN-OS 7.0 . . . . . 12
Upgrade Firewalls Using Panorama . . . . . 12
Upgrade a Firewall to PAN-OS 7.0 . . . . . 15
Upgrade an HA Firewall Pair to PAN-OS 7.0 . . . . . 17
Downgrade from PAN-OS 7.0 . . . . . 20
Downgrade to a Previous Maintenance Release . . . . . 20
Downgrade to a Previous Feature Release . . . . . 21
Management Features . . . . . 23
All New Application Command Center . . . . . 24
Automated Correlation Engine . . . . . 26
Correlation Objects . . . . . 26
Correlated Events . . . . . 26
Global Find . . . . . 28
Tag Browser . . . . . 29
Configuration Validation Improvements . . . . . 31
Validate a Firewall Configuration . . . . . 31
Validate a Panorama Configuration . . . . . 32
Move and Clone Policies, Objects, and Templates . . . . . 33
Move or Clone a Policy or Object to a Virtual System . . . . . 33
Move or Clone a Policy or Object to a Device Group . . . . . 34
Extended SNMP Support . . . . . 35
SNMP Counter Monitoring . . . . . 35
SNMP Interface MIB for Logical Interfaces . . . . . 35
LLDP MIB . . . . . 35
SaaS Application Usage Report . . . . . 36
Policy Impact Review for New Content Releases . . . . . 38
Review New App-IDs . . . . . 38
Disable or Enable App-IDs . . . . . 40
Prepare Policy Updates For Pending App-IDs . . . . . 41
Virtual System/Device Name in Reports and Logs . . . . . 43
Time-Based Log and Report Deletion . . . . . 44
Configure Time-Based Log and Report Deletion on a Firewall or Panorama . . . . . 44
Configure Time-Based Log Deletion on a Collector Group . . . . . 45
Software Upload Improvements . . . . . 46
Upload and Install Software to a Single Device . . . . . 46
Upload and Install Software to Multiple Firewalls Using Panorama . . . . . 47


© Palo Alto Networks, Inc.

Panorama Features . . . . . 49
Device Group Hierarchy . . . . . 50
Device Group Hierarchy Inheritance and Overrides . . . . . 50
Create a Device Group Hierarchy . . . . . 51
Template Stacks . . . . . 53
Firewall Modes and Overlapping Settings in Stacks . . . . . 53
Configure a Template Stack . . . . . 54
Role-Based Access Control Enhancements . . . . . 56
Firewall Configuration Import into Panorama . . . . . 59
Log Redundancy Within a Collector Group . . . . . 62
Firewall HA State in Panorama . . . . . 63

WildFire Features . . . . . 65
WildFire Grayware Verdict . . . . . 66
WildFire Hybrid Cloud . . . . . 68
WildFire Analysis Profile . . . . . 71

Content Inspection Features . . . . . 73
Configurable Drop Actions in Security Profiles . . . . . 74
Actions in Security Profiles . . . . . 74
Set the Action in a Security Profile . . . . . 75
Blocking of Encoded Content . . . . . 76
Negate Operator for Custom Threat Signatures . . . . . 77

Authentication Features . . . . . 79
Authentication and Authorization Enhancements . . . . . 80
SSL/TLS Service Profiles . . . . . 81
TACACS+ Authentication . . . . . 82
Kerberos V5 Single Sign-On . . . . . 83
Configure Kerberos SSO for Administrator Authentication . . . . . 83
Configure Kerberos SSO for Captive Portal Authentication . . . . . 85
Suite B Cryptography Support . . . . . 88
Suite B Ciphers . . . . . 88
Generate and Assign ECDSA Certificates . . . . . 89
Configure a GlobalProtect IPSec Crypto Profile . . . . . 91
Authentication Server Connectivity Testing . . . . . 92

Decryption Features . . . . . 95
SSL Decryption Enhancements . . . . . 96

User-ID Features . . . . . 99
User Attribution Based on X-Forwarded-For Headers . . . . . 100


© Palo Alto Networks, Inc.

Custom Groups Based on LDAP Filters . . . 101

Virtualization Features . . . 103

Support for High Availability on the VM-Series Firewall . . . 104
HA Timers on the VM-Series Firewalls . . . 104

High Availability for VM-Series in AWS . . . 106

Support for Jumbo Frames . . . 107

Support for Hypervisor Assigned MAC Addresses . . . 108

Networking Features . . . 109

ECMP . . . 110
ECMP Platform, Interface, and IP Routing Support . . . 110
Configure ECMP on a Virtual Router . . . 110

DHCP Options . . . 113

Granular Actions for Blocking Traffic in Security Policy . . . 114

Session-Based DSCP Classification . . . 116

Per-Virtual System Service Routes . . . 119
Customize Service Routes for a Virtual System . . . 119

LLDP . . . 121
Configure LLDP . . . 121

Network Prefix Translation (NPTv6) . . . 123
Create an NPTv6 Policy . . . 123

TCP Split Handshake Drop . . . 126

VPN Features . . . 127

IKEv2 Support for VPN Tunnels . . . 128

IPSec VPN Enhancements . . . 129
Refresh and Restart Behavior for IKE Gateway and IPSec Tunnel . . . 129
Enable or Disable an IKE Gateway or IPSec Tunnel . . . 129
Refresh or Restart an IKE Gateway or IPSec Tunnel . . . 130

GlobalProtect Features . . . 131

Disable Direct Access to Local Networks . . . 132

Static IP Address Allocation . . . 133

Apply a Gateway Configuration to Users, Groups, and/or Operating Systems . . . 135

Welcome Page Management . . . 136

RDP Connection to a Remote Client . . . 137

Simplified GlobalProtect License Structure . . . 138

SSL/TLS Service Profiles for GlobalProtect Portals and Gateways . . . 139

GlobalProtect IPSec Crypto Profiles for GlobalProtect Client Configurations . . . 140

Licensing Features . . . 141

Support for Usage-Based Licensing in AWS . . . . . .142
Launch the VM-Series Firewall in the AWS-VPC . . . . . .142
Register the Usage-Based Model of the VM-Series Firewall in AWS . . . . . .143
Self-Service License & Subscription Management . . . . . .145

 © Palo Alto Networks, Inc.

# Upgrade to PAN-OS 7.0


- Upgrade/Downgrade Considerations
- Upgrade the Firewall to PAN-OS 7.0
- Downgrade from PAN-OS 7.0

# Upgrade/Downgrade Considerations

Table: PAN-OS 7.0 Upgrade/Downgrade Considerations lists the new features that have upgrade and/or downgrade impacts. Make sure you understand the changes that will occur in the configuration prior to upgrading to or downgrading from PAN-OS 7.0. For additional information about this release, refer to the Release Notes.

**Table: PAN-OS 7.0 Upgrade/Downgrade Considerations**

| Feature | Upgrade Considerations | Downgrade Considerations |
|---|---|---|
| Template Stacks | Panorama template configurations will no longer have multiple virtual systems mode, operational mode (normal, FIPS, or CC), or VPN mode settings. | All templates will have an operational mode set to normal, VPN mode set to enabled, and multiple virtual systems mode set to enabled. |
| Role-Based Access Control Enhancements | Panorama creates an access domain configuration named <administrator_name>_AD for any access control settings that are associated with an administrator account. Panorama associates the access domain with the role assigned to the account. | Panorama populates the access control settings in an administrator account with the values from the first listed access domain in that account. Panorama also assigns the first listed role to the account. |
| Log Redundancy Within a Collector Group | Log redundancy is disabled by default. | Before downgrading Panorama, disable log redundancy in Collector Groups to avoid log data loss. After disabling, only one copy of the logs will be available for queries. |
| Authentication and Authorization Enhancements | • PAN-OS moves the **User Domain**, **Kerberos Realm**, and **Retrieve User Group** values from server profiles to the authentication profiles that reference them.<br>• In authentication profiles, the **Username Modifier** field is set to **None**. If you leave the field at this value:<br>  • For RADIUS authentication, the device normalizes the username to the NetBIOS format (domain\user).<br>  • For LDAP and Kerberos authentication, the device removes any domain that the user enters during login.<br>• After upgrading, make sure that any authentication profile selected for global administrative access to the web interface references a RADIUS server profile. Any other type of server profile will cause a commit failure. | • PAN-OS converts any periods to underscores in the names of authentication profiles and sequences.<br>• In LDAP server profiles, a blank **Login Attribute** defaults to sAMAccountName in PAN-OS 7.0. However, after a downgrade, the blank field will have no value. To avoid login failures after a downgrade, manually enter sAMAccountName in any LDAP server profile that has a blank **Login Attribute**. |


© Palo Alto Networks, Inc.

| Feature | Upgrade Considerations | Downgrade Considerations |
|---|---|---|
| SSL/TLS Service Profiles | PAN-OS creates an SSL/TLS service profile for each certificate that was assigned to a device service, and assigns the profile to that service. The profile name is <certificate_name> - ssl-tls-service-profile. If no certificate was assigned for a service, PAN-OS sets the **SSL/TLS Service Profile** value to **None** for that service. | PAN-OS replaces each SSL/TLS service profile that was assigned to a device service with the certificate associated with that profile. |
| Suite B Cryptography Support | LSVPN tunnel interfaces between a GlobalProtect firewall and a satellite firewall cannot pass traffic if you upgrade the satellite firewall to a PAN-OS 7.0.0–7.0.6 release while the GlobalProtect firewall runs a PAN-OS 6.1 or earlier release. The workaround is to upgrade both firewalls to a PAN-OS 7.0 or later release. If you cannot, yet, upgrade the GlobalProtect firewall to a PAN-OS 7.0 or later release, then you must upgrade the satellite firewall to PAN-OS 7.0.7 or a later release. | • LSVPN tunnel interfaces between a GlobalProtect firewall and a satellite firewall cannot pass traffic if you downgrade the GlobalProtect firewall to a PAN-OS 6.1 or earlier release while the satellite firewall runs a PAN-OS 7.0.0–7.0.6 release. The workaround is to upgrade the satellite firewall to PAN-OS 7.0.7 or a later release before downgrading the GlobalProtect firewall to a PAN-OS 6.1 or earlier release. If you don't want to upgrade the satellite firewall, the other option is to downgrade both firewalls to a PAN-OS 6.1 or earlier release.<br>• When you initiate a downgrade on a device that uses ECDSA certificates, PAN-OS displays a warning, prompting you to remove those certificates and any references to them (for example, in SSL/TLS service profiles) before performing the downgrade.<br>• In profiles that use Diffie-Hellman groups (for example, IPSec Crypto profiles), DH group 14 replaces DH group 19 or 20.<br>• In profiles that reference Suite B algorithms (for example, IPSec Crypto profiles), algorithm aes-256-cbc replaces aes-256-gcm and algorithm aes-128-cbc replaces aes-128-gcm.<br>• PAN-OS removes GlobalProtect IPSec Crypto profiles from gateway configurations. |
| Policy Impact Review for New Content Releases | | You cannot successfully downgrade to a previous PAN-OS release version when the most recent content release has been Downloaded to the firewall, but is not yet installed. To downgrade to any previous PAN-OS software release, first select **Device > Dynamic Updates** and **Install** the latest content release version. |

| Feature | Upgrade Considerations | Downgrade Considerations |
|---|---|---|
| WildFire Hybrid Cloud and WildFire Analysis Profile | Palo Alto Networks highly recommends that you save the current Panorama configuration before upgrading to Panorama 7.0: see the Downgrade Considerations for details on the impact to WildFire.<br>Following the upgrade to PAN-OS 7.0, update the firewall WildFire configuration based on the WildFire deployment you plan to enable or continue using (public, private, or hybrid cloud):<br>• Forward files to the WildFire public cloud only:<br>If the firewall was configured to forward files to the WildFire public cloud before the upgrade to PAN-OS 7.0, no action is required to continue to forward files to the WildFire public cloud only.<br>• Forward files to a WildFire appliance only:<br>Upgrade the WildFire appliance and make sure to complete Step 9, a required step to enable the firewall to forward files to a WildFire appliance.<br>• Forward files to a WildFire Hybrid Cloud:<br>To enable the firewall to forward files to the WildFire public cloud or a WildFire appliance (based on the new WildFire Analysis profile settings), set up a WildFire Hybrid Cloud.<br>Review the new WildFire Analysis Profile for details on new WildFire settings and changes to behavior introduced with PAN-OS 7.0. | When prompted during the downgrade process, load a Panorama configuration that was saved before the upgrade to Release 7.0 to ensure that any File Blocking profiles with a rule **Action** set to **forward** or **continue-and-forward** (used for WildFire forwarding) will be available. |
| Tag Browser | On upgrade, the maximum number of tags that the firewall and Panorama can support is now increased from 2,500 to 10,000. This limit is enforced across the firewall/Panorama and is not allocated by virtual system or device group. | To prevent a commit failure on downgrade to 6.1, delete the tags in excess of 2,500; on downgrade to 6.0, delete the tags in excess of 1,024. |

 © Palo Alto Networks, Inc.

| Feature | Upgrade Considerations | Downgrade Considerations |
|---|---|---|
| Federal Information Processing Standard (FIPS) Mode | If your firewall is running a PAN-OS 6.1 or earlier release and is in FIPS mode, you must Enable FIPS and Common Criteria Support using Set CCEAL4 Mode before you upgrade to PAN-OS 7.0.1 or a later release. If you do not change to CCEAL4 mode before you upgrade, the firewall will enter maintenance mode because FIPS mode is not supported as of PAN-OS 7.0.1.<br>  After you change from FIPS mode to CCEAL4 mode, you will need to import the saved configuration backup that you created prior to the mode change. If the configuration contains IKE and IPSec crypto profiles that use 3DES, you will need to delete the profiles and create new profiles using AES because 3DES is not supported in CCEAL4 mode. | |

# Upgrade the Firewall to PAN-OS 7.0

How you upgrade to PAN-OS 7.0 depends on whether you have standalone firewalls or firewalls in a high availability (HA) configuration and whether, for either scenario, your firewalls are managed by Panorama. Review the Release Notes and then follow the procedure specific to your configuration:

- Upgrade Firewalls Using Panorama
- Upgrade a Firewall to PAN-OS 7.0
- Upgrade an HA Firewall Pair to PAN-OS 7.0

When upgrading firewalls that you manage with Panorama or firewalls that are configured to forward content to a WF-500 appliance, you must first upgrade Panorama and its Log Collectors and upgrade the WF-500 appliance, 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 deployed in a high availability (HA) configuration.

| Upgrade Firewalls Using Panorama | | |
|---|---|---|
| Step 1 | Save a backup of the current configuration file on each managed firewall you plan to upgrade.<br>  Although the firewall will automatically create a backup of the configuration, it is a best practice to create and externally store a backup prior to upgrading. | 1. Log in to Panorama, select **Panorama > Setup > Operations**, and click **Export Panorama and devices config bundle** to generate and export the latest configuration backup of Panorama and of each managed device.<br>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.<br>  Make sure the firewalls you plan to upgrade are running content release version 497 or later. | 1. Select **Panorama > Device Deployment > Dynamic Updates**.<br>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.<br>3. **Download** the desired version. After a successful download, the link in the Action column changes from **Download** to **Install**.<br>4. Click **Install**, select the devices on which you want to install the update, and click **OK**. |

 © Palo Alto Networks, Inc.

## Upgrade Firewalls Using Panorama (Continued)

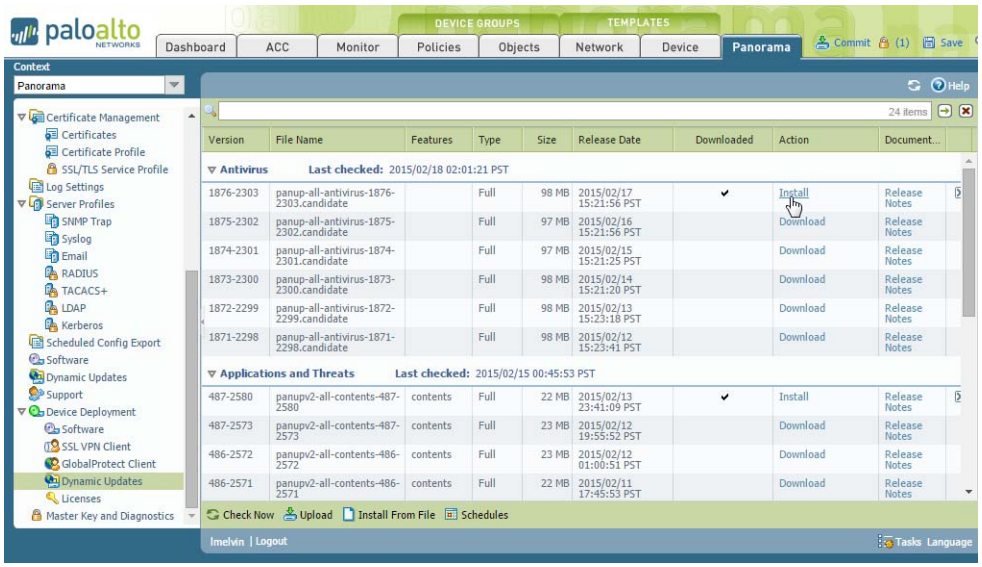


| | | |
|---|---|---|
| Step 3 | Determine the software upgrade path.<br>You cannot skip any major release versions on the path to your desired PAN-OS version. For example, if you want to upgrade from PAN-OS 5.0.13 to PAN-OS 7.0.2, you must:<br>• Download and install PAN-OS 6.0.0 and reboot.<br>• Download and install PAN-OS 6.1.0 and reboot.<br>• Download PAN-OS 7.0.1 (7.0.1 is the base image for the 7.0 release; not 7.0.0).<br>• Download and install PAN-OS 7.0.2 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.<br>2. Select **Device > Software**.<br>3. Check which version has a check mark in the Currently Installed column and proceed as follows:<br>• If PAN-OS 6.1.0 or later is currently installed, continue to Step 4.<br>• If a version earlier than PAN-OS 6.1.0 is currently installed, follow the upgrade path to 6.1.0 before you upgrade to 7.0. 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.<br>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 7.0.1, download the images that have File Name PanOS_3000-7.0.1 and PanOS_5000-7.0.1. After a successful download, the link in the Action column changes to **Install**. |

| Upgrade Firewalls Using Panorama (Continued) | |
|---|---|
| Step 5 Install the software updates on the firewalls.<br>  <br>To avoid downtime when updating the software on HA firewalls, update one peer at a time.<br>For active/active firewalls, it doesn't matter which HA peer you update first.<br>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:<br>• **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**.<br>• **Active/active HA firewalls:**<br>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.<br>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**.<br>• **Active/passive HA firewalls**—In this example, the active firewall is named fw1 and the passive firewall is named fw2:<br>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.<br>b. Access fw1, select **Device > High Availability > Operational Commands**, and click **Suspend local device**.<br>c. Access fw2 and, on the **Dashboard**, **High Availability** widget, verify that the **Local** firewall state is **active** and the **Peer** firewall is **suspended**.<br>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.<br>e. Access fw1, select **Device > High Availability > Operational Commands**, and click **Make local device functional**. Wait two minutes before proceeding.<br>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**. |

 © Palo Alto Networks, Inc.

**Upgrade Firewalls Using Panorama (Continued)**

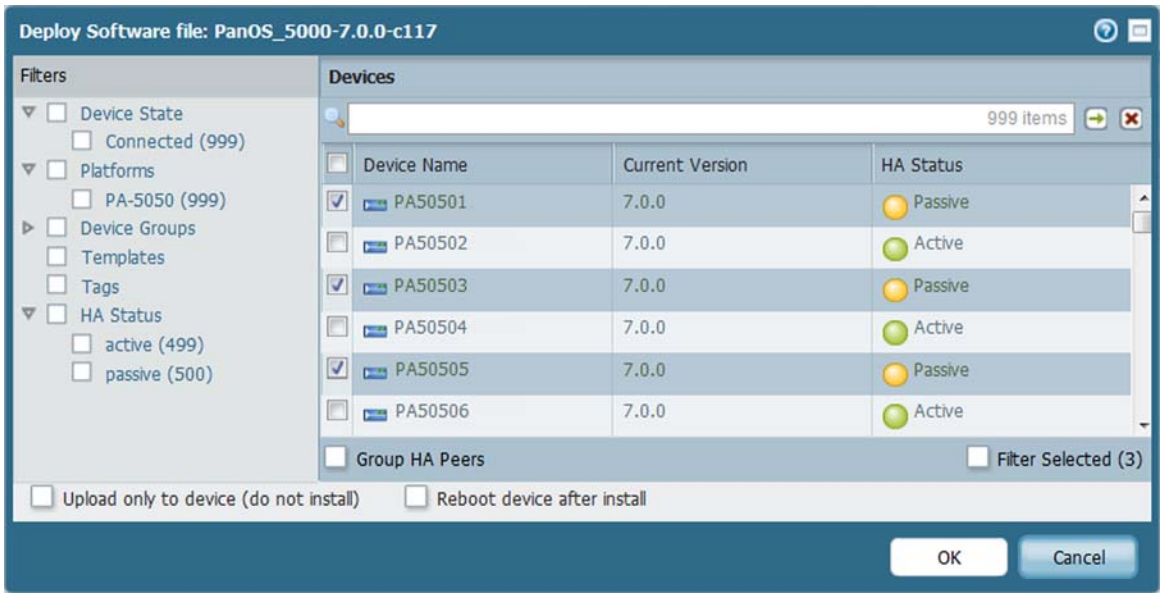


| | | |
|---|---|---|
| Step 6 | Verify the software and content release version running on each managed firewall. | 1. On Panorama, select **Panorama > Managed Devices**.<br>2. Locate the firewalls and review the content and software versions in the table. |

| | Device Name | Status: Device State | Status: HA Status | Software Version | Apps and Threat | Antivirus | URL Filtering | WildFire |
|---|---|---|---|---|---|---|---|---|
| | AcmeCorp-Europe (1/1 Devices Connected): Shared > Acme-Corp > AcmeCorp-Perimeter > AcmeCorp-Europe | | | | | | | |
| | PA50502 | Connected | Active | 7.0.0 | 227-852 | 321-450 | 3435 | 66437-75388 |
| | PA50501 | | Passive | 7.0.0 | 227-852 | 321-450 | 3435 | 66437-75388 |
| | PA50503 | Connected | Passive | 7.0.0 | 227-852 | 321-450 | 3435 | 66437-75388 |
| | PA50504 | | Active | 7.0.0 | 227-852 | 321-450 | 3435 | 66437-75388 |
| | PA50506 | Connected | Active | 7.0.0 | 227-852 | 321-450 | 3435 | 66437-75388 |
| | PA50505 | | Passive | 7.0.0 | 227-852 | 321-450 | 3435 | 66437-75388 |

## Upgrade a Firewall to PAN-OS 7.0

Review the Release Notes and then use the following procedure to upgrade a firewall that is not in an HA configuration to PAN-OS 7.0.

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.<br>Although the firewall will automatically create a backup of the configuration, it is a best practice to create and externally store a backup prior to upgrading. | 1. Select **Device > Setup > Operations** and click **Export named configuration snapshot**.<br>2. Select the XML file that contains your running configuration (for example, **running-config.xml**) and click **OK** to export the configuration file.<br>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 the firewall is running content release version 497 or later. | 1. Select **Device > Dynamic Updates**.<br>2. Check the **Applications and Threats** or **Applications** section to determine what update is currently running.<br>3. If the firewall is not running the required update or later, click **Check Now** to retrieve a list of available updates.<br>4. Locate the desired update and click **Download**.<br>5. After the download completes, click **Install**. |
| Step 3 Determine the upgrade path.<br>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.<br>For example, if you want to upgrade from PAN-OS 5.0.13 to PAN-OS 7.0.2, you must:<br>• Download and install PAN-OS 6.0.0 and reboot.<br>• Download and install PAN-OS 6.1.0 and reboot.<br>• Download PAN-OS 7.0.1 (7.0.1 is the base image for the 7.0 release; not 7.0.0).<br>• Download and install PAN-OS 7.0.2 and reboot. | 1. Select **Device > Software**.<br>2. Check which version has a check mark in the Currently Installed column and proceed as follows:<br>• If PAN-OS 6.1.0 or later is currently installed, continue to Step 4.<br>• If a version of PAN-OS prior to 6.1.0 is currently installed (as shown here), follow the upgrade path to 6.1.0 before you can upgrade to 7.0. Refer to the release notes for your currently installed PAN-OS version for upgrade instructions.<br> 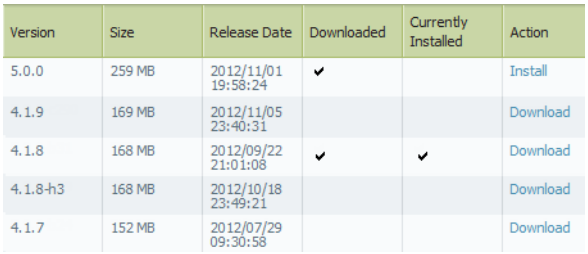 <br>Version / Size / Release Date / Downloaded / Currently Installed / Action<br>5.0.0 / 259 MB / 2012/11/01 19:58:24 / ✔ / / Install<br>4.1.9 / 169 MB / 2012/11/05 23:40:31 / / / Download<br>4.1.8 / 168 MB / 2012/09/22 21:01:08 / ✔ / ✔ / Download<br>4.1.8-h3 / 168 MB / 2012/10/18 23:49:21 / / / Download<br>4.1.7 / 152 MB / 2012/07/29 09:30:58 / / / Download |
| Step 4 Install PAN-OS 7.0.<br>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.<br>2. Locate the version you want to upgrade to and then click **Download**.<br>3. After the download completes, click **Install**.<br>4. After the install completes, reboot using one of the following methods:<br>• If you are prompted to reboot, click **Yes**.<br>• If you are not prompted to reboot, select **Device > Setup > Operations** and click **Reboot Device** in the Device Operations section. |

 © Palo Alto Networks, Inc.

| Upgrade PAN-OS (Continued) | | |
|---|---|---|
| Step 5 | Verify that the firewall is passing traffic. | Select **Monitor > Session Browser**. |

## Upgrade an HA Firewall Pair to PAN-OS 7.0

Review the Release Notes and then use the following procedure to upgrade a pair of firewalls in a high availability (HA) configuration. This procedure applies to both active/passive and active/active configurations.

When upgrading peers in an HA configuration, you must upgrade each firewall 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 one feature release to the next consecutive feature release, PAN-OS 6.1.x to PAN-OS 7.0 in this case. 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 as describe in the following procedure.

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.<br>  <br>Although the firewall will automatically create a backup of the configuration, it is a best practice to create and externally store a backup prior to upgrading. | Perform these steps on each firewall in the pair:<br>1. Select **Device > Setup > Operations** and click **Export named configuration snapshot**.<br>2. Select the XML file that contains your running configuration (for example, **running-config.xml**) and click **OK** to export the configuration file.<br>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 497 or later. | 1. Select **Device > Dynamic Updates**.<br>2. Check the **Applications and Threats** or **Applications** section to determine what update is currently running.<br>3. If the firewall is not running the required update or later, click **Check Now** to retrieve a list of available updates.<br>4. Locate the content release **Version** you want to install and click **Download**.<br>5. After the download completes, click **Install**. |

| Upgrade PAN-OS (Continued) | |
|---|---|
| Step 3 Determine the upgrade path.<br>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.<br>For example, if you want to upgrade from PAN-OS 5.0.13 to PAN-OS 7.0.2, you must:<br>• Download and install PAN-OS 6.0.0 and reboot.<br>• Download and install PAN-OS 6.1.0 and reboot.<br>• Download PAN-OS 7.0.1 (7.0.1 is the base image for the 7.0 release; not 7.0.0).<br>• Download and install PAN-OS 7.0.2 and reboot. | 1. Select **Device > Software**.<br>2. Check which version has a check mark in the Currently Installed column and proceed as follows:<br>• If PAN-OS 6.1.0 or later is currently installed, continue to Step 4.<br>• If a version of PAN-OS prior to 6.1.0 is currently installed (as shown here), follow the upgrade path to 6.1.0 before you can upgrade to 7.0. Refer to the Release Notes for your currently installed PAN-OS version for upgrade instructions.<br> 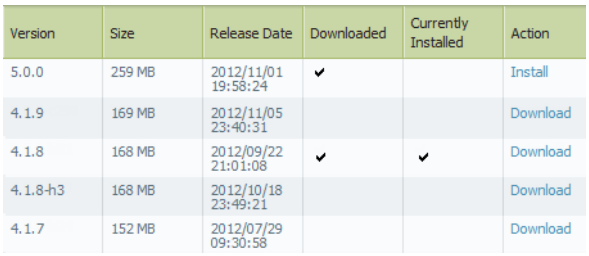 <br>Version / Size / Release Date / Downloaded / Currently Installed / Action<br>5.0.0 / 259 MB / 2012/11/01 19:58:24 / ✔ / / Install<br>4.1.9 / 169 MB / 2012/11/05 23:40:31 / / / Download<br>4.1.8 / 168 MB / 2012/09/22 21:01:08 / ✔ / ✔ / Download<br>4.1.8-h3 / 168 MB / 2012/10/18 23:49:21 / / / Download<br>4.1.7 / 152 MB / 2012/07/29 09:30:58 / / / Download |
| Step 4 Install PAN-OS 7.0 on the passive device (active/passive) or on the active-secondary device (active/active).<br>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.<br>2. Locate the version you want to upgrade to and then click **Download**.<br>3. After the download completes, click **Install**.<br>4. After the install completes, reboot using one of the following methods:<br>• If you are prompted to reboot, click **Yes**.<br>• 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**.<br>2. Click **Suspend local device**.<br>3. Select **Dashboard** and verify that the state of the passive device changes to **active** in the **High Availability** widget.<br>4. Verify that the firewall that took over as active or active-primary is passing traffic by selecting **Monitor > Session Browser**.<br>5. (Optional) If you have session synchronization enabled and you are not upgrading directly from PAN-OS 6.1.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. |

 © Palo Alto Networks, Inc.

| Upgrade PAN-OS (Continued) | | |
|---|---|---|
| Step 6 | Install PAN-OS 7.0 on the other device in the pair.<br>  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.<br>2. Locate the version you want to upgrade to and then click **Download**.<br>3. After the download completes, click **Install**.<br>4. After the install completes, reboot using one of the following methods:<br>• If you are prompted to reboot, click **Yes**.<br>• 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.<br>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.<br>  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.<br>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:<br>• (Active device(s) only) To verify that the active devices are passing traffic, run **`show session all`**.<br>• 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:<br>• In an active/passive configuration, only the active device will show packets transmitted and the passive device will only show packets received.<br>• In an active/passive configuration, only the active device will show packets transmitted and the passive device will only show packets received.<br>  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.<br>• In the active/active configuration, you will see packets received and packets transmitted on both devices. |

# Downgrade from PAN-OS 7.0

The way you downgrade from PAN-OS 7.0 depends on whether you are downgrading to a previous feature release (where the first or second digit in the PAN-OS version changes, for example 7.0 to 6.1 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 7.0.4 to 7.0.2). 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 6.1.4 to 6.1.2 is considered a maintenance release downgrade because only the third digit in the release version is different.

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

| Downgrade to a Previous Maintenance Release | |
|---|---|
| Step 1 Save a backup of the current configuration file.<br>  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**.<br>2. Select the XML file that contains your running configuration (for example, **running-config.xml**) and click **OK** to export the configuration file.<br>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. |

 © Palo Alto Networks, Inc.

| Downgrade to a Previous Maintenance Release | |
|---|---|
| Step 2 Install the previous maintenance release image.<br>  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**.<br>2. Locate the version to which you want to downgrade. If the image has not yet been downloaded, click **Download**.<br>3. After the download completes, click **Install**.<br>4. After the install completes, reboot using one of the following methods:<br>• If you are prompted to reboot, click **Yes**.<br>• 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 7.0 to any version earlier than PAN-OS 5.0.5 is not supported because the log management subsystem has been significantly enhanced between PAN-OS 6.0 and PAN-OS 5.0. 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.<br>  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**.<br>2. Select the XML file that contains your running configuration (for example, **running-config.xml**) and click **OK** to export the configuration file.<br>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.<br>  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**.<br>2. Locate the version to which you want to downgrade. If the image has not yet been downloaded, click **Download**.<br>3. After the download completes, click **Install**.<br>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 7.0.1 and want to downgrade to PAN-OS 6.1.3, select autosave-6.1.3.<br>5. After the install completes, reboot using one of the following methods:<br>• If you are prompted to reboot, click **Yes**.<br>• If you are not prompted to reboot, select **Device > Setup > Operations** and click **Reboot Device** in the Device Operations section. |

 © Palo Alto Networks, Inc.

# Management Features


- All New Application Command Center
- Automated Correlation Engine
- Global Find
- Tag Browser
- Configuration Validation Improvements
- Move and Clone Policies, Objects, and Templates
- Extended SNMP Support
- SaaS Application Usage Report
- Policy Impact Review for New Content Releases
- Virtual System/Device Name in Reports and Logs
- Time-Based Log and Report Deletion
- Software Upload Improvements


© Palo Alto Networks, Inc.

# All New Application Command Center

The Application Command Center (ACC) is an interactive, graphical summary of the applications, users, URLs, threats, and content traversing your network.The ACC uses the firewall logs to provide visibility into traffic patterns and actionable information on threats. The new ACC layout includes a tabbed view of network activity, threat activity, and blocked activity and each tab includes pertinent widgets for better visualization of traffic patterns on your network. The graphical representation allows you to interact with the data and visualize the relationships between events on the network so that you can uncover anomalies or find ways to enhance your network security rules. For a personalized view of your network, you can also add a custom tab and include widgets that allow you to drill down into the information that is most important to you.

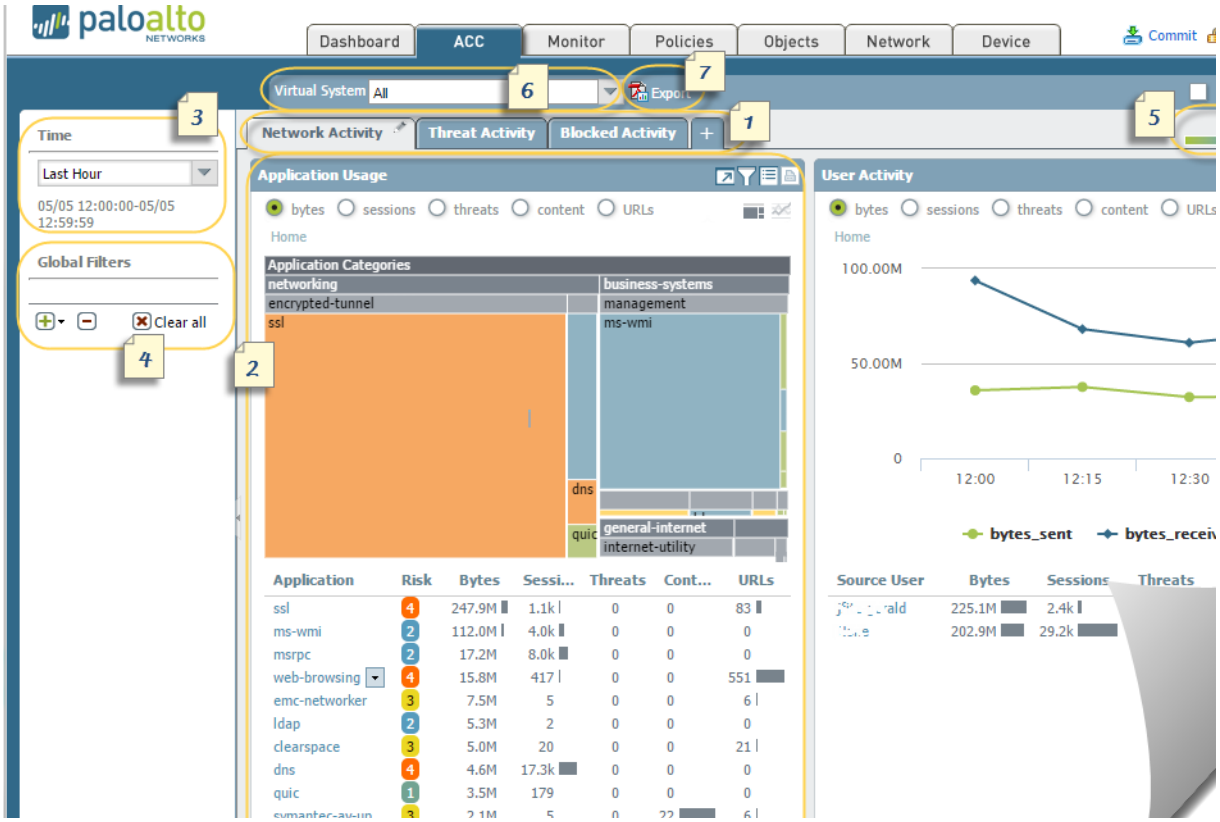


| ACC - First Look | | |
|---|---|---|
| 1 | **Tabs** | The ACC includes three predefined tabs that provide visibility into network traffic, threat activity, and blocked activity. |
| 2 | **Widgets** | Each tab includes a default set of widgets that best represent the events/trends associated with the tab. The widgets allow you to survey the data using the following filters:<br>• bytes—in and out,<br>• sessions<br>• content—files and data<br>• URL categories<br>• threats— malicious, benign, and count. |

 © Palo Alto Networks, Inc.

| ACC - First Look | | |
|---|---|---|
|  | **Time** | The charts or graphs in each widget provide a real-time and historic view. You can choose a custom range or use the predefined time periods that range from the last 15 minutes up to the last 30 days or last 30 calendar days.<br>The time period used to render data, by default, is the last hour updated in 15 minute intervals. The date and time interval are displayed onscreen, for example at 11:40 is: 01/12 10:30:00-01/12 11:29:59 |
|  | **Global Filters** | The global filters allow you to apply a filter across all tabs. The charts/graphs apply the selected filters before rendering the data. |
|  | **Risk Factor** | The risk factor (1=lowest to 5=highest) indicates the relative security risk on your network. The risk factor uses a variety of factors such as the type of applications seen on the network and their associated risk levels, the threat activity and malware as seen through the number of blocked threats, compromised hosts or traffic to malware hosts/domains. |
|  | **Source** | The data source used for the display.<br>On the firewall, if enabled for multiple virtual systems, you can use the **Virtual System** drop-down to change the ACC display to include all virtual systems or just a selected virtual system.<br>On Panorama, the **Data Source** can be **Panorama data** or **Remote Device Data**. Remote Device Data is only available when all the managed firewalls are on PAN-OS 7.0.1 or later.<br>When the data source is Panorama, you can filter the display for a specific device group. |
|  | **Export** | You can export the widgets displayed in the current tab as a PDF. |

# Automated Correlation Engine

The automated correlation engine is an analytics tool that uses the logs on the firewall to detect actionable events on your network. The engine correlates a series of related threat events that, when combined, indicate a likely attack on your network. It pinpoints areas of risk, such as compromised hosts on the network, allows you to assess the risk and take action to prevent exploitation of network resources. The automated correlation engine uses *correlation objects* to analyze the logs for patterns and when a match occurs, it generates a *correlated event*.

The automated correlation engine is supported on the following platforms only:

- Panorama—M-Series appliance and the virtual appliance
- PA-7000 Series
- PA-5000 Series
- PA-3000 Series

▲ Correlation Objects

▲ Correlated Events

## Correlation Objects

A correlation object is a definition file that specifies patterns to match against, the data sources to use for the lookups, and the time period within which to look for these patterns. A pattern is a boolean structure of conditions that queries the following data sources (or logs) on the firewall: application statistics, traffic, traffic summary, threat summary, threat, data filtering, and URL filtering. Each pattern has a severity rating, and a threshold for the number of times the pattern match may occur within a defined time limit to indicate malicious activity. When the match conditions are met, a correlation event is logged.

To view the correlation objects that are currently available, select **Monitor > Automated Correlation Engine > Correlation Objects**. All the objects in the list are enabled by default.

## Correlated Events

A correlation event is logged when the patterns and thresholds defined in a correlation object match the traffic patterns on your network. You can view and analyze the logs generated for each correlated event in the **Monitor > Automated Correlation Engine > Correlated Events** tab.

 © Palo Alto Networks, Inc.


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

| Tab | Description |
|---|---|
| **Match Information** | Object Details: Presents information on the Correlation Objects that triggered the match. |
| | Match Details: A summary of the match details that includes the match time, last update time on the match evidence, severity of the event, and an event summary. |
| **Match Evidence** | Presents all the evidence that corroborates the correlated event. It lists detailed information on the evidence collected for each session. |

For a graphical display of the correlated events, see the compromised hosts widget on **ACC >Threat Activity**. The compromised hosts widget aggregates the correlated events and sorts them by severity. It displays the source IP address/user who triggered the event, the correlation object that was matched and the number of times the object was matched. The match count link allows you to jump to the match evidence details.

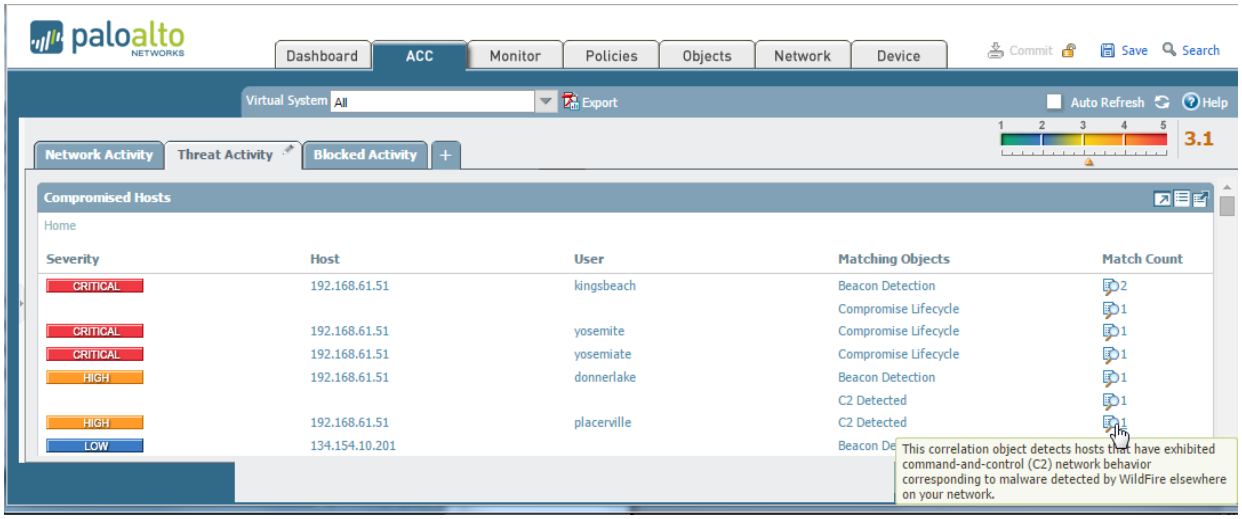

# Global Find

To make the management of your Palo Alto Networks devices more efficient, a new Global Find feature is introduced to enable you to search the candidate configuration on a firewall or Panorama for a particular string, such as an IP address, object name, policy rule name, threat ID, or application name. The search results are grouped by category and provide links to the configuration location in the web interface, so that you can quickly and easily find all of the places where the string is referenced. For example, if you temporarily deny an application that is defined in multiple security policy rules and you now want to allow that application, you can search on the application name and quickly locate all referenced polices to change the action back to allow.

Watch the video.

Global Find will not search dynamic content (such as logs, address ranges, or allocated DHCP addresses). In the case of DHCP, you can search on a DHCP server attribute, such as the DNS entry, but you cannot search for individual addresses allocated to users. Global Find also does not search for individual user or group names identified by User-ID unless the user/group is defined in a policy. In general, you can only search content that the firewall writes to the configuration.

## Use Global Find

- Launch Global Find by clicking the **Search** icon located on the upper right of the web interface.

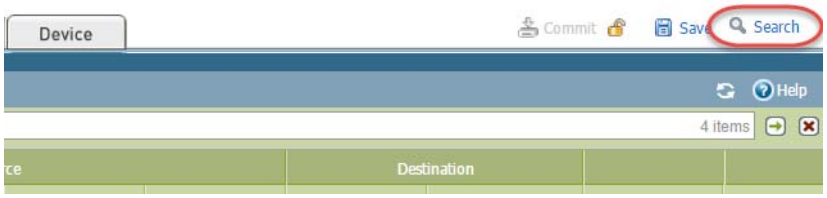


- To access the Global Find from within a configuration area, click the drop-down next to an item and click **Global Find** as follows:

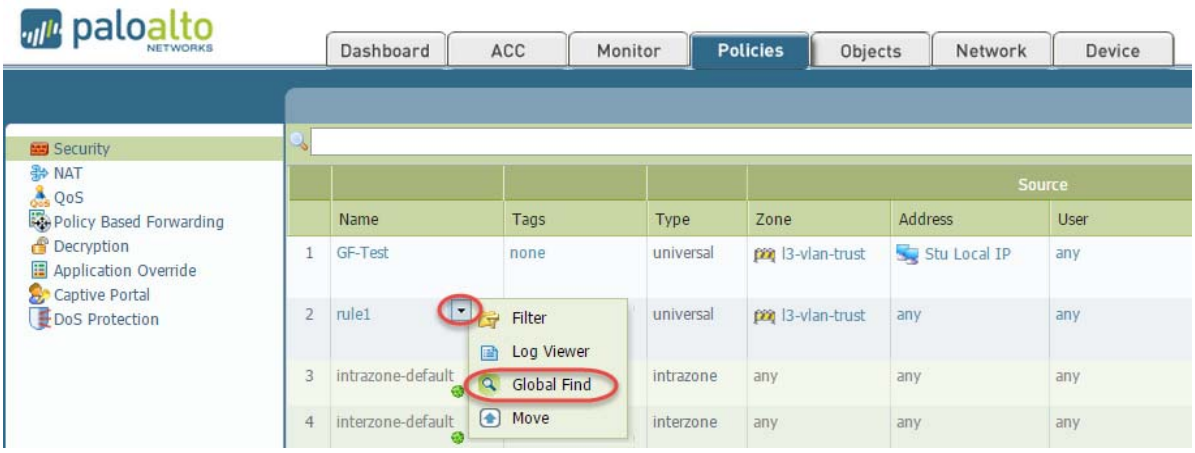


 © Palo Alto Networks, Inc.

# Tag Browser

The tag browser provides a way to view all the tags used within a rulebase. In rulebases with a large number of rules, the tag browser simplifies the display by presenting the tags, the color code, and the rule numbers in which the tags are used.

The tag browser also allows you to group rules using the first tag applied to the rule. As a best practice, use the first tag to identify the primary purpose for a rule. For example, the first tag can identify a rule by a high-level function such as high-risk applications, personal applications, or IT sanctioned applications. In the tag browser, when you **Filter by first tag in rule**, you can easily identify gaps in coverage and move rules or add new rules within the rulebase. All the changes are saved to the candidate configuration until you commit the changes on the firewall and make them a part of the running configuration.

For devices that are managed by Panorama, the tags applied to pre-rules and post-rules that have been pushed from Panorama display in a green background and are demarcated with green lines so that you can identify these tags from the local tags on the device.

The maximum number of tags that the firewall and Panorama support is now increased from 2,500 to 10,000. This limit is enforced across the firewall/Panorama and is not allocated by virtual system or device group.

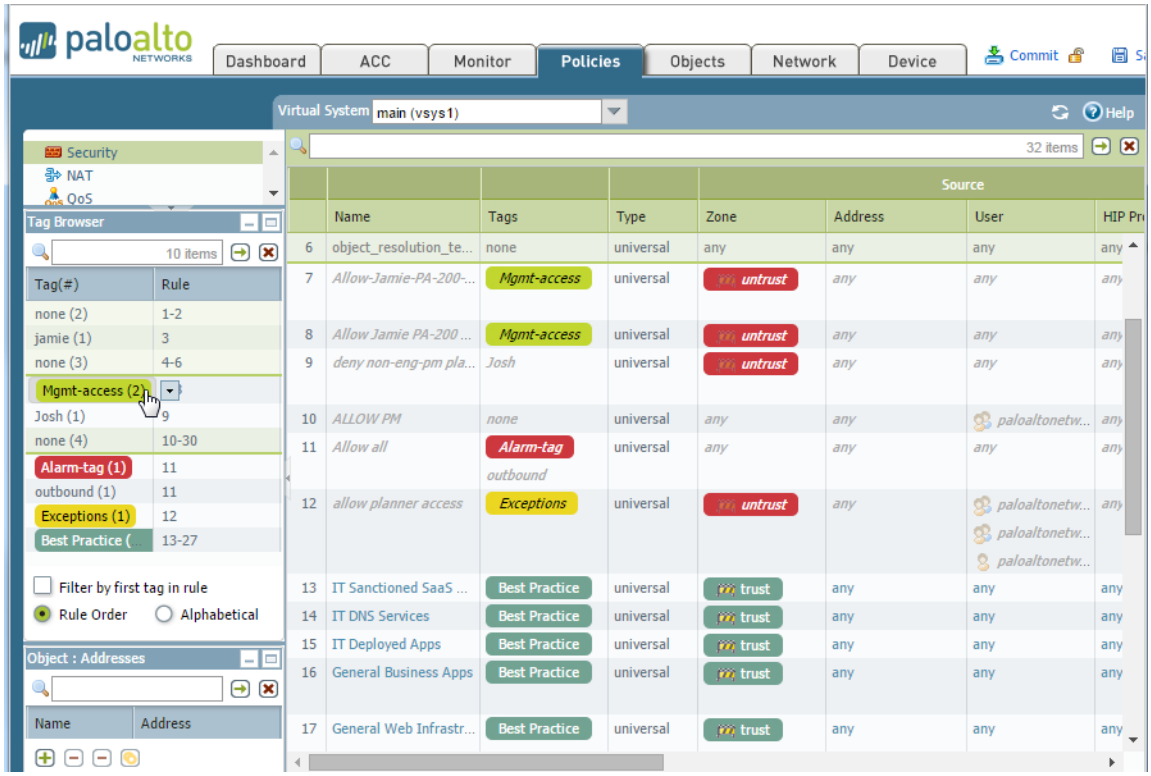

| Use the Tag Browser | |
|---|---|
| • Explore the tag browser. | 1. Access the **Tag Browser** on the left pane of the **Policies >** tab. The tag browser displays the tags that have been used in the rules for the selected rulebase, for example **Policies > Security**.<br>2. **Tag (#)**—Displays the label and the rule number or range of numbers in which the tag is used contiguously. Hover over the label to see the location where the rule was defined. It can be inherited from a shared location, a device group, or a virtual system.<br>3. **Rule**—Lists the rule number or range of numbers associated with the tags.<br>4. Sort the tags.<br>• **Filter by first tag in rule**—Sorts rules using the first tag applied to each rule in the rulebase. This view is particularly useful if you want to narrow the list and view related rules that might be spread around the rulebase. For example if the first tag in each rule denotes its function—best practices, administration, web-access, data center access, proxy—you can narrow the result and scan the rules based on function.<br>• **Rule Order**—Sorts the tags in the order of appearance within the selected rulebase. When displayed in order of appearance, tags used in contiguous rules are grouped. The rule number with which the tag is associated is displayed along with the tag name.<br>• **Alphabetical**—Sorts the tags in alphabetical order within the selected rulebase. The display lists the tag name and color (if a color is assigned) and the number of times it is used within the rulebase.<br>The label **None** represents rules without any tags; it does not display rule numbers for untagged rules. When you select **None**, the right pane is filtered to display rules that have no tags assigned to them.<br>5. **Clear**—Clears the filter on the currently selected tags in the search bar.<br>6. **Search bar**—To search for a tag, enter the term and click the green arrow icon to apply the filter. The tag browser also displays the total number of tags in the rulebase and the number of selected tags.<br>7. Expand or collapse the tag browser.  |
| Refer the PAN-OS Administrator's Guide for details on creating and applying tags and using the tag browser. | |

 © Palo Alto Networks, Inc.

# Configuration Validation Improvements

You can now Use the Web Interface to perform a syntactic validation (of configuration syntax) and semantic validation (whether the configuration is complete and makes sense) of a firewall or Panorama candidate configuration before committing it. The results display all of the errors and warnings of a full commit or virtual system commit, including rule shadowing and application dependency warnings. Possible errors could be an invalid route destination or a missing account and password that are required to query a server. Such validation significantly reduces failures at commit time.

The new **Validate Changes** method of validating a configuration (using the **Commit** button, as shown in the task below) replaces the former method of validating (using **Device > Setup > Operations > Validate**). The former method was limited to syntactic validation.

Only one commit or validate function can be run at one time on either the firewall or Panorama.

The predefined Admin Roles of superuser, device, and virtual system include the **Validate** option as an allowed task. Therefore, you do not need to specifically allow validation in predefined roles. You can control validation in custom admin roles as well. Validation is enabled by default. Alternatively, you can Restrict Admin Access to Validation Functions. You can also create a custom admin role that allows validation on Panorama.

- ▲ Validate a Firewall Configuration
- ▲ Validate a Panorama Configuration

## Validate a Firewall Configuration

| Validate a Firewall Configuration | |
|---|---|
| Step 1 Validate a firewall configuration. | 1. After you have made one or more configuration changes, click **Commit**.<br>2. Click **Advanced** to select specific types of changes:<br>• Click **Include Device and Network configuration** to include device and network changes in the validation.<br>• Click **Include Policy and Object configuration** (not available on multiple virtual system firewalls) to include policy and object changes in the validation.<br>• Click **Include Shared Object configuration** (on multiple virtual system firewalls only) to include shared object changes in the validation.<br>3. If your platform supports multiple virtual systems, and if you click **Include Virtual System configuration,** click **All virtual systems** or **Select one or more virtual systems**, in which case, select the virtual systems you want validated.<br>4. Click **Validate Changes**. Alternatively, from any screen that has the **Validate Changes** button, click **Validate Changes**. |

| Validate a Firewall Configuration (Continued) | | |
|---|---|---|
| Step 2 | View the validation results. | The Validate window displays the percentage of validation completed. The **Result** indicates **OK** if the validation succeeded. The **Details** indicate any configuration errors or warnings.<br>On the Task Manager, the **Status** indicates **Completed** and the Result is displayed. |

## Validate a Panorama Configuration

First you validate and/or commit the candidate configuration on Panorama, and then you validate the configuration that Panorama will push to the device group or template for firewalls. Thus, you can independently validate for Panorama, device groups, and templates.

| Validate a Panorama Configuration | | |
|---|---|---|
| Step 1 | Validate a Panorama candidate configuration. | 1. After making one or more configuration changes, click **Commit**.<br>2. For **Commit Type**, select **Panorama**.<br>3. Click **Validate Changes**.<br>The **Result** is **OK** if the validation succeeds. The **Details** indicate any errors or warnings. |
| Step 2 | Validate a candidate configuration for a device group or template to be pushed to the firewall. | 1. Click **Commit**.<br>2. For **Commit Type**, select **Template** or **Device Group** and select a template or device group from the list.<br>3. (Optional) Click **Merge with Device Candidate Config** if desired.<br>4. (Optional) Click **Include Device and Network Templates** if desired.<br>5. Click **Validate Changes**. The Job Status might indicate something similar to "validation succeeded with warnings." Click on the status phrase to open the Details window, which indicates any errors or warnings. |

 © Palo Alto Networks, Inc.

## Validate a Panorama Configuration (Continued)

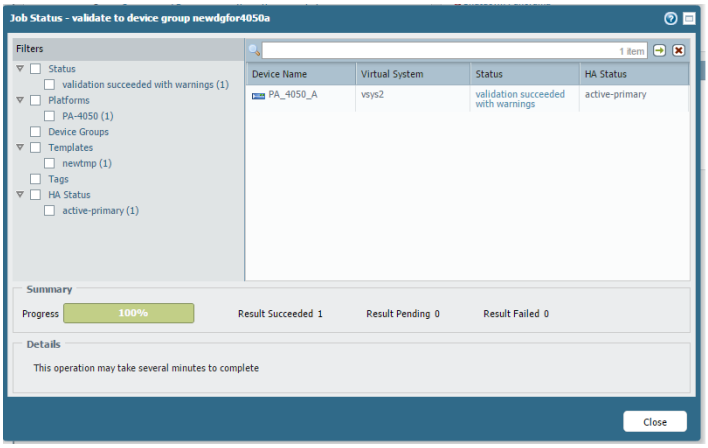

# Move and Clone Policies, Objects, and Templates

You can now move or clone policy rules and objects to a different virtual system, device group, or the Shared location. This saves you the effort of deleting, recreating, or renaming rules and objects when only a move or copy is needed. Moving and cloning is particularly useful for cleaning up device groups after a Firewall Configuration Import into Panorama. You can also clone templates and template stacks now in the same way as other configurations in the **Panorama** tab (select the item and click **Clone**).

- Move or Clone a Policy or Object to a Virtual System
- Move or Clone a Policy or Object to a Device Group

## Move or Clone a Policy or Object to a Virtual System

On a firewall, if a policy rule or object that you will move or clone from a virtual system (vsys) has references to objects in that vsys, move or clone the referenced objects also. If the references are to shared objects, you don't have to include those when moving or cloning. You can perform a Global Find to check for references.

| Move or Clone a Policy or Object to a Virtual System | |
|---|---|
| Step 1 | Log in to the firewall and select the policy type (for example, **Policy > Security**) or object type (for example, **Objects > Addresses**). |
| Step 2 | Select the **Virtual System** and select one or more policy rules or objects. |
| Step 3 | Perform one of the following steps:<br>• Select **Move > Move to other vsys** (for policy rules).<br>• Click **Move** (for objects).<br>• Click **Clone** (for policy rules or objects). |
| Step 4 | In the **Destination** drop-down, select the new virtual system or **Shared**. The default is the **Virtual System** selected in Step 2. |
| Step 5 | (Policies only) Select the **Rule order**:<br>• **Move top** (default)—The rule will come before all other rules.<br>• **Move bottom**—The rule will come after all other rules.<br>• **Before rule**—In the adjacent drop-down, select the rule that comes after the Selected Rules.<br>• **After rule**—In the adjacent drop-down, select the rule that comes before the Selected Rules. |
| Step 6 | **Error out on first detected error in validation** is enabled by default, which means the firewall will display the first error it finds and stop checking for more errors. For example, an error occurs if the **Destination** vsys doesn't have an object that the policy rule you are moving references. When you move or clone many items at once, selecting this check box can simplify troubleshooting errors one at a time. If you clear the check box, the firewall will find all the errors before displaying them. Regardless of this setting, the firewall won't move or clone anything until you fix all the errors for all the selected items. |
| Step 7 | Click **OK** to start the error validation. If the firewall finds errors, fix them and retry the move or clone operation. If the firewall doesn't find errors, it performs the operation. After the operation finishes, click **Commit**. |

 © Palo Alto Networks, Inc.

## Move or Clone a Policy or Object to a Device Group

On Panorama, if a policy rule or object that you will move or clone from a device group has references to objects that are not available in the target device group (**Destination**), move or clone the referenced objects also. In a Device Group Hierarchy, remember that referenced objects might be available through inheritance. For example, shared objects are available in all device groups. You can perform a Global Find to check for references. If you move or clone an overridden object, be sure that overrides are enabled for that object in the parent device group of the **Destination** (see Step 4 under Create a Device Group Hierarchy).

| Move or Clone a Policy or Object to a Device Group | |
|---|---|
| Step 1 | Log in to Panorama and select the policy type (for example, **Policy > Security**) or object type (for example, **Objects > Addresses**). |
| Step 2 | Select the **Device Group** and select one or more policy rules or objects. |
| Step 3 | Perform one of the following steps:<br>• Select **Move > Move to other device group** (for policy rules).<br>• Click **Move** (for objects).<br>• Click **Clone** (for policy rules or objects). |
| Step 4 | In the **Destination** drop-down, select the new device group or **Shared**. The default is the **Device Group** selected in Step 2. |
| Step 5 | (Policies only) Select the **Rule order**:<br>• **Move top** (default)—The rule will come before all other rules.<br>• **Move bottom**—The rule will come after all other rules.<br>• **Before rule**—In the adjacent drop-down, select the rule that comes after the Selected Rules.<br>• **After rule**—In the adjacent drop-down, select the rule that comes before the Selected Rules. |
| Step 6 | **Error out on first detected error in validation** is enabled by default, which means Panorama will display the first error it finds and stop checking for more errors. For example, an error occurs if the **Destination** device group doesn't have an object that the policy rule you are moving references. When you move or clone many items at once, selecting this check box can simplify troubleshooting errors one at a time. If you clear the check box, Panorama will find all the errors before displaying them. Regardless of this setting, Panorama won't move or clone anything until you fix all the errors for all the selected items. |
| Step 7 | Click **OK** to start the error validation. If Panorama finds errors, fix them and retry the move or clone operation. If Panorama doesn't find errors, it performs the operation. |
| Step 8 | Click **Commit**, for the **Commit Type** select **Panorama**, then click **Commit** again. |
| Step 9 | Click **Commit**, for the **Commit Type** select **Device Group**, select the original and destination device groups, then click **Commit** again. |

# Extended SNMP Support

PAN-OS support for Simple Network Management Protocol (SNMP) now includes the following features. To access the latest MIBs, refer to SNMP MIB Files.

- SNMP Counter Monitoring
- SNMP Interface MIB for Logical Interfaces
- LLDP MIB

## SNMP Counter Monitoring

You can now track global counters related to Denial of Service (DoS), IP fragmentation, TCP state, and dropped packets. Tracking these counters enables you to monitor traffic irregularities that result from DoS attacks, device or connection faults, or resource limitations. Monitoring such irregularities is useful for maintaining the health and security of your network. Previously, you had to use the device CLI or XML API to monitor global counters. The counters belong to a new panGlobalCounters MIB. In a MIB browser, the path is **panCommonMib > panCommonObjs > panSys > panGlobalCounters**.

## SNMP Interface MIB for Logical Interfaces

The PAN-OS implementation of the interfaces and IfMIB have been extended to support all logical interfaces on the firewall, including tunnels, aggregate groups, L2 subinterfaces, L3 subinterfaces, loopback interfaces, and VLAN interfaces. This is in addition to the SNMP Interface MIB support on physical interfaces. The VPN tunnel status can be now monitored.

## LLDP MIB

Palo Alto Networks firewalls now support the LLDP v2 MIB (OID 1.3.111.2.802.1.1.13) for monitoring Link Layer Discovery Protocol (LLDP) events. For example, you can check the lldpV2StatsRxPortFramesDiscardedTotal object to see the number of LLDP frames that were discarded for any reason.

 © Palo Alto Networks, Inc.

# SaaS Application Usage Report

A new predefined report is introduced to provide visibility into Software as a Service (SaaS) application usage. The SaaS Application Usage report enables you to assess and subsequently mitigate the risks to your enterprise's data when taking advantage of SaaS applications. The report will also help you assess risks to the security of your enterprise network, such as the delivery of malware through SaaS applications adopted by your users.

This report, which uses the SaaS application characteristic (see **Objects > Applications**), lists the top SaaS applications (up to 100) running on your network on a given day.

| View the SaaS Report | |
|---|---|
| Step 1 | Select **Monitor > Reports**. |
| Step 2 | Expand the **Application Reports** section in the right-hand frame and select **SaaS Application Usage**. |
| Step 3 | Select a date for which to view SaaS application traffic from the calendar. The report displays in the middle pane. |

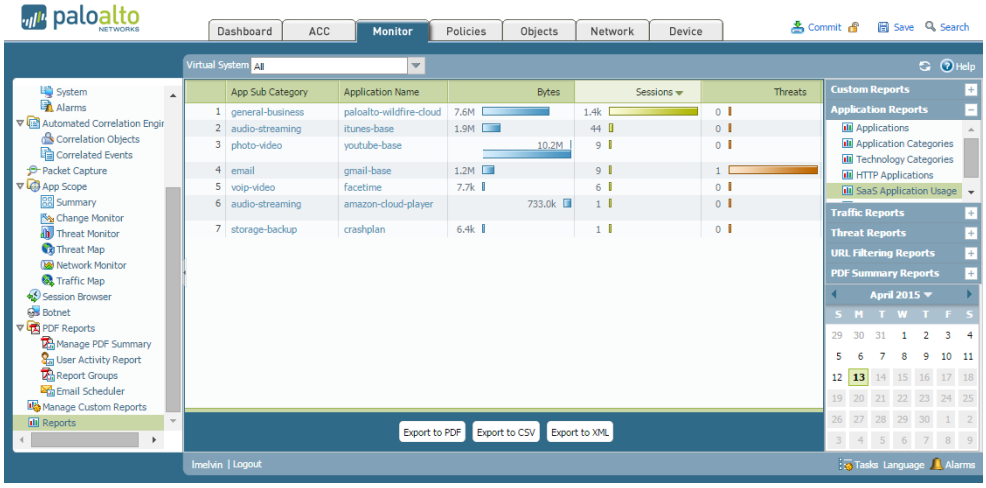


| | |
|---|---|
| Step 4 | Use the report to gain visibility into the SaaS application traffic that is running on your network. The report identifies the application name and subcategory of each SaaS application and details the number of sessions and bytes for each application on the selected date. In addition, the report identifies the number of threats detected in each of the applications. |
| Step 5 | (Optional) Export the report to PDF, CSV, or XML for archive or analysis by clicking the corresponding button. |

## View the SaaS Report (Continued)

Step 6 To investigate any suspicious traffic, click the application name or category to view more details in the All New Application Command Center.

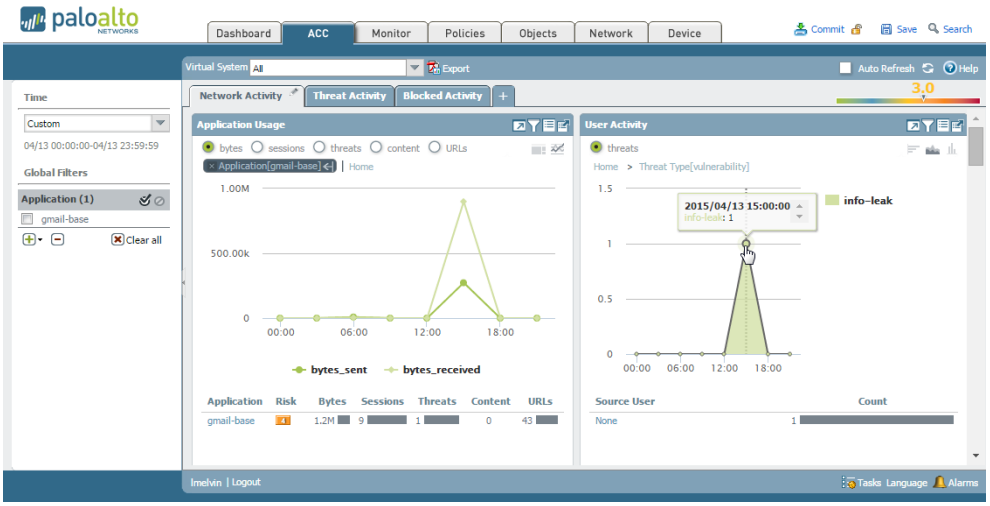


Step 7 (Optional) You can also use the new SaaS characteristic to generate custom reports that help you monitor SaaS activity on your network by entering `characteristic-of-name eq 'is-saas'` in the Query Builder section of the custom report builder.

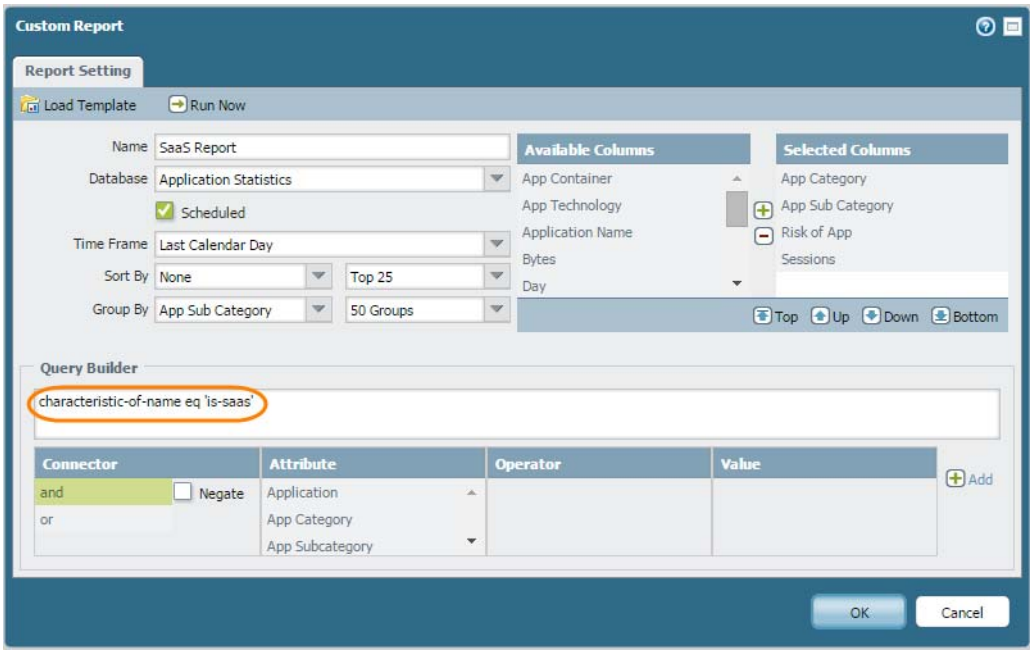


 © Palo Alto Networks, Inc.

# Policy Impact Review for New Content Releases

Before installing a content release, you can now review the policy impact for new App-IDs and stage any necessary policy updates. This enables you to assess the treatment an application receives both before and after the new content is installed and then prepare any related policy updates to take effect at the same time that the content update is installed. This feature specifically includes the capability to modify existing security policies using *pending* App-IDs. Pending App-IDs are application signatures contained in a downloaded content release (prior to installing the new content) or signatures that you have manually disabled. You can simultaneously update your security policy rules and install and/or enable pending App-IDs, to allow for a seamless shift in policy enforcement. The option to install threat signatures immediately, but to delay installing App-ID signatures allows you to be protected against the latest threats, while providing the flexibility to enable the pending App-IDs after you've had the chance to prepare any policy changes.

The following options enable you to assess the impact of new App-IDs on your existing policy rules, disable (and enable) App-IDs, and seamlessly update policies to secure and enforce newly-identified applications:

- ▲ Review New App-IDs
- ▲ Disable or Enable App-IDs
- ▲ Prepare Policy Updates For Pending App-IDs

## Review New App-IDs

After downloading a new content release version, you can review the new App-IDs included in the content version and assess the impact of the new App-IDs on existing policy rules:

- Review the list of new App-IDs that are available since the last installed content release version.
- Review the impact of new App-IDs on existing policy rules.

## Review New App-IDs

Review the list of new App-IDs that are available since the last installed content release version

Select **Device > Dynamic Updates** and click the **Apps** link in the **Features** column to view details on newly-identified applications:

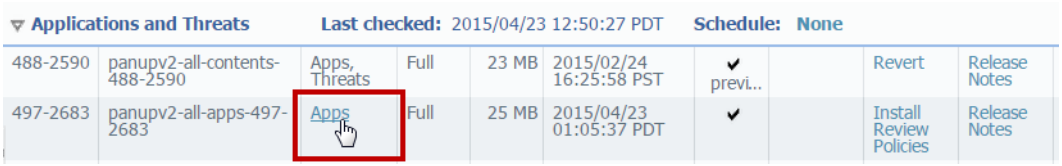


A list of App-IDs shows all new App-IDs introduced from the content version installed on the firewall, to the selected **Content Version**.

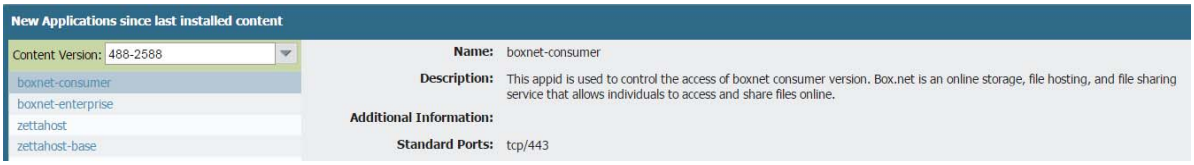


App-ID details that you can use to assess possible impact to policy enforcement include:

- **Depends on**—Lists the application signatures that this App-ID relies on to uniquely identify the application. If one of the application signatures listed in the **Depends On** field is disabled, the dependent App-ID is also disabled.
- **Previously Identified As**—Lists the App-IDs that matched to the application before the new App-ID was installed to uniquely identify the application.
- **App-ID Enabled**—All App-IDs display as enabled when a content release is downloaded, unless you choose to manually disable the App-ID signature before installing the content update (see Disable or Enable App-IDs).

Multi-vsys firewalls display App-ID status as **vsys-specific**. This is because the status is not applied across virtual systems and must be individually enabled or disabled for each virtual system. To view the App-ID status for a specific virtual system, select **Objects > Applications**, select a **Virtual System**, and select the App-ID.

 © Palo Alto Networks, Inc.

**Review New App-IDs (Continued)**

Review the impact of new App-IDs on existing policy rules

1. Select **Device > Dynamic Updates**.
2. You can review the policy impact of new content release versions that are downloaded to the firewall. **Download** a new content release version, and click the **Review Policies** in the Action column. The **Policy review based on candidate configuration** dialog allows you to filter by **Content Version** and view App-IDs introduced in a specific release (you can also filter the policy impact of new App-IDs according to **Rulebase** and **Virtual System**).


3. Select a new App-ID from the **Application** drop-down to view policy rules that currently enforce the unidentified application traffic. Use the detail provided in the policy review to plan policy rule updates to take effect when the application is uniquely identified.

You can continue to Prepare Policy Updates For Pending App-IDs, or you can directly add the new App-ID to policy rules that the application was previously matched to by continuing to use the policy review dialog.

In the following example, the new App-ID adobe-cloud is introduced in a content release. Adobe-cloud traffic is currently identified as SSL and web-browsing traffic. Policy rules configured to enforce SSL or web-browsing traffic are listed to show what policy rules will be affected when the new App-ID is installed. In this example, the rule Allow SSL App currently enforces SSL traffic. To continue to allow adobe-cloud traffic when it is uniquely identified, and no longer identified as SSL traffic.


Add the new App-ID to existing policy rules, to allow the application traffic to continue to be enforced according to your existing security requirements when the App-ID is installed.

In this example, to continue to allow adobe-cloud traffic when it is uniquely identified by the new App-ID, and no longer identified as SSL traffic, add the new App-ID to the security policy rule Allow SSL App.

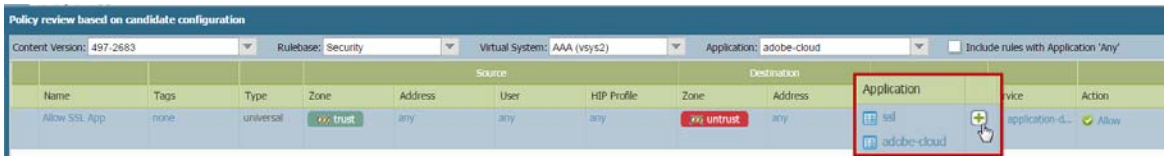


The policy rule updates take effect only when the application updates are installed.

## Disable or Enable App-IDs

You can now choose to disable new App-IDs introduced in a content release, in order to immediately benefit from protection against the latest threats while continuing to have the flexibility to later enable App-IDs after preparing necessary policy updates. You can disable all App-IDs introduced in a content release, set scheduled content updates to automatically disable new App-IDs, or disable App-IDs for specific applications.

Policy rules referencing App-IDs only match to and enforce traffic based on enabled App-IDs.

| Disable and Enable App-IDs | |
|---|---|
| Disable all App-IDs in a content release or for scheduled content updates. | • To disable all new App-IDs introduced in a content release, select **Device > Dynamic Updates** and **Install** an Application and Threats content release. Before the release is installed, a prompt is displayed enabling you to **Disable new apps in content update**. Select the check box to disable apps and continue installing the content release; this allows you to be protected against threats, and gives you the option to enable the applications at a later time.<br>• On the **Device > Dynamic Updates** page, select Schedule. Choose to **Disable new apps in content update** for downloads and installations of content releases. |
| Disable App-IDs for one application or multiple applications at a single time. | • To quickly disable a single application or multiple applications at the same time, click **Objects > Applications**. Select one or more application check box and click **Disable**.<br>• To review details for a single application, and then disable the App-ID for that application, select **Objects > Applications** and **Disable App-ID**. You can use this step to disable both pending App-IDs (where the content release including the App-ID is downloaded to the firewall but not installed) or installed App-IDs. |
| Enable App-IDs. | Enable App-IDs that you previously disabled by selecting **Objects > Applications**. Select one or more application check box and click **Enable** or open the details for a specific application and click **Enable App-ID**. |

## Prepare Policy Updates For Pending App-IDs

You can now stage seamless policy updates for new App-IDs. Release versions prior to PAN-OS 7.0 required you to install new App-IDs (as part of a content release) and then make necessary policy updates. This allowed for a period during which the newly-identified application traffic was not enforced, either by existing rules (that the traffic had matched to before being uniquely identified) or by rules that had yet to be created or modified to use the new App-ID.

*Pending* App-IDs can now be added to policy rules to prevent gaps in policy enforcement that could occur during the period between installing a content release and updating security policy. Pending App-IDs includes App-IDs that have been manually disabled, or App-IDs that are downloaded to the firewall but not installed. Pending App-IDs can be used to update policies both before and after installing a new content release. Though they can be added to policy rules, pending App-IDs are not enforced until the App-IDs are both installed and enabled on the firewall.

The names of App-IDs that have been manually disabled display as gray and italicized, to indicate the disabled status:

- Disabled App-ID listed on the **Objects > Applications** page:


- Disabled App-ID included in a security policy rule:

 © Palo Alto Networks, Inc.


App-IDs that are included in a downloaded content release version might have an App-ID status of enabled, but App-IDs are not enforced until the corresponding content release version is installed.

## Perform Seamless Policy Updates for New App-IDs

To install the content release version now and then update policies:

Do this to benefit from new threat signatures immediately, while you review new application signatures and update your policies.

1. Select **Device > Dynamic Updates** and **Download** the latest content release version.
2. Review New App-IDs to assess the policy impact of new App-IDs.
3. **Install** the latest content release version. Before the content release is installed, you are prompted to **Disable new apps in content update**. Select the checkbox and continue to install the content release. Threat signatures included in the content release will be installed and effective, while new or updated App-IDs are disabled.
4. Select **Policies** and update **Security**, **QoS**, and **Policy Based Forwarding** rules to match to and enforce the now uniquely identified application traffic, using the pending App-IDs.
5. Select **Objects > Applications** and select one or multiple disabled App-IDs and click **Enable**.
6. **Commit** your changes to seamlessly update policy enforcement for new App-IDs.

To update policies now and then install the content release version:

1. Select **Device > Dynamic Updates** and **Download** the latest content release version.
2. Review New App-IDs to assess the policy impact of new App-IDs.
3. While reviewing the policy impact for new App-IDs, you can use the **Policy Review based on candidate configuration** to add a new App-ID to existing policy rules: .
4. The new App-ID is added to the existing rules as a disabled App-ID.
5. Continue to review the policy impact for all App-IDs included in the latest content release version by selecting App-IDs in the **Applications** drop-down. Add the new App-IDs to existing policies as needed. Click **OK** to save your changes.
6. **Install** the latest content release version.
7. **Commit** your changes to seamlessly update policy enforcement for new App-IDs.

# Virtual System/Device Name in Reports and Logs

You can now create a PAN-OS or Panorama report, or view or search logs, based on a virtual system name or device name, which are more user-friendly attributes than the virtual system ID or device ID. Now you need not personally map a virtual system name to its ID, or map a device name to its serial number, in order to view or search logs, or create reports.

In **Monitor > Manage Custom Reports**, when adding or revising a report, **Virtual System Name** and **Device Name** are available in the **Available Columns** field, the **Group By** field, and the **Attribute** field of Query Builder if they are relevant to the **Database** selected for the report.

For information on generating custom reports on the firewall, see Generate Custom Reports.

For PAN-OS or Panorama, when viewing logs, you can now add a log filter based on a virtual system name or device name. In **Monitor > Logs**, after choosing a log type (such as Traffic, Threat, etc.), click on the **+** to add a log filter. **Virtual System Name** and **Device Name** are available in the **Attribute** field if they are relevant to the type of log selected.

When a security policy rule is created to export or forward logs, the **Virtual System Name** and **Device Name** are included in the log.

Because Device Name and Virtual System Name log fields are introduced in PAN-OS 7.0, when forwarding logs to Panorama from firewalls running PAN-OS releases prior to PAN-OS 7.0, Panorama will display the device name in the Serial Number column for readability purposes. PAN-OS and Panorama cannot generate reports based on Device Name or Virtual System Name for logs received from devices running a release prior to PAN-OS 7.0.

For information on forwarding logs to Panorama, in the Panorama Administrator's Guide, see Configure Log Forwarding to Panorama.

 © Palo Alto Networks, Inc.

# Time-Based Log and Report Deletion

For logs and reports that the firewall, Panorama, and Log Collectors generate, you can now configure automatic deletion based on time, not just on size quotas. This is useful in deployments where periodically deleting monitored data is desired or necessary. For example, deleting user data after a certain period might be mandatory in your organization for legal reasons.

If a firewall is in Common Criteria (CC) mode, it won't delete logs unless they were first exported.

- Configure Time-Based Log and Report Deletion on a Firewall or Panorama
- Configure Time-Based Log Deletion on a Collector Group

## Configure Time-Based Log and Report Deletion on a Firewall or Panorama

On the firewall, the expiration periods apply to all its virtual systems. On Panorama, the expiration periods apply to the logs and reports that a Panorama management server and its managed collectors generate.

The firewall and Panorama synchronize expiration periods across high availability (HA) pairs. Because only the active HA peer generates logs, the passive peer has no logs or reports to delete unless failover occurs and it starts generating logs.

To configure expiration periods for the logs that managed collectors receive from firewalls, see Configure Time-Based Log Deletion on a Collector Group.

| Configure Time-Based Log and Report Deletion on a Firewall or Panorama | |
|---|---|
| Step 1 | On the firewall, select **Device > Setup > Management** and edit the Logging and Reporting Settings.<br>On Panorama, select **Panorama > Setup > Management** and edit the Logging and Reporting Settings. |
| Step 2 | In the **Log Storage** tab, enter the expiration period (**Max Days**) for each log type and each summary log type (1-2,000 days). By default, the fields are blank for all log types, which means the logs never expire. |
| Step 3 | In the **Log Export and Reporting** tab, enter the **Report Expiration Period** (1-2,000 days). By default, the field is blank, which means reports never expire. |
| Step 4 | On the firewall, click **OK** and **Commit**.<br>On Panorama, click **OK** and **Commit**, for the **Commit Type** select **Panorama**, then click **Commit** again. |

## Configure Time-Based Log Deletion on a Collector Group

The expiration periods you configure for a Collector Group apply to the logs that its managed collectors receive from firewalls.

To configure expiration periods for the logs that Panorama and managed collectors generate, and for the reports that Panorama generates, see Configure Time-Based Log and Report Deletion on a Firewall or Panorama.

| Configure Time-Based Log Deletion on a Collector Group | |
|---|---|
| Step 1 | Select **Panorama > Collector Groups** and click the Name of the Collector Group. |
| Step 2 | Click the **Log Storage** link. |
| Step 3 | Enter the expiration period (**Max Days**) for each log type and each summary log type (1-2,000 days). By default, the fields are blank for all log types, which means the logs never expire. |
| Step 4 | Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, then click **Commit** again. |
| Step 5 | Click **Commit**, for the **Commit Type** select **Collector Group**, select the Collector Group, then click **Commit** again. |

 © Palo Alto Networks, Inc.

# Software Upload Improvements

Devices now display details (for example, version, size, and release date) about uploaded software updates that enable you to check, before installing an update, that it is the intended one. Installing uploaded software and synchronizing it across high availability (HA) peers now involves fewer steps. This makes software deployment easier when a device doesn't have access to the external network.

You can upload up to two software images on a device. If you upload a third, the device deletes the oldest image.

- Upload and Install Software to a Single Device
- Upload and Install Software to Multiple Firewalls Using Panorama

## Upload and Install Software to a Single Device

This procedure applies to PAN-OS and Panorama software.

| Upload and Install Software to a Single Device | |
|---|---|
| Step 1 | Log in to a computer that has access to the Software Update site and download the update to that computer. |
| Step 2 | For a PAN-OS update, log in to the firewall and select **Device > Software**.<br>For a Panorama update, log in to Panorama and select **Panorama > Software**. |
| Step 3 | Click **Upload** and **Browse** to the update. |
| Step 4 | (High availability devices only) Select the **Sync To Peer** check box to push the imported software image to the secondary HA peer. As of this release, the web interface of the secondary peer displays the update without requiring you to click **Check Now**. |
| Step 5 | Click **OK** to upload the image. The Software page then displays the same details for the uploaded image as it does for downloaded images (except Release Notes). The Available column displays Uploaded, the Action column displays an **Install** link, and a delete icon appears in the far right column for the update. |
| Step 6 | Click **Install** to install the software. The device logs you out while it reboots to complete the installation. |

## Upload and Install Software to Multiple Firewalls Using Panorama

You can upload and install the current PAN-OS release or earlier releases to firewalls that are managed devices on Panorama.

| Upload and Install Software to Multiple Firewalls Using Panorama | |
|---|---|
| Step 1 | Log in to a computer that has access to the Software Update site and download the update to that computer. |
| Step 2 | Log in to Panorama and select **Panorama > Device Deployment > Software**. |
| Step 3 | Click **Upload** and **Browse** to the update. |
| Step 4 | Click **OK** to upload the image. The Software page then displays the same details for the uploaded image as it does for downloaded images (except Release Notes). The Available column displays Uploaded, the Action column displays an **Install** link, and a delete icon appears in the far right column for the update. |
| Step 5 | Click **Install** and select the firewalls. If you select the **Reboot device after install** check box, the firewall will reboot during installation. The installation can't finish without a reboot. Click **OK** to start the installation. |

 © Palo Alto Networks, Inc.

# Panorama Features


- Device Group Hierarchy
- Template Stacks
- Role-Based Access Control Enhancements
- Firewall Configuration Import into Panorama
- Log Redundancy Within a Collector Group
- Firewall HA State in Panorama


© Palo Alto Networks, Inc.

# Device Group Hierarchy

You can now create nested device groups in a tree hierarchy, with lower-level groups inheriting the settings of higher-level groups, and all groups inheriting settings from the Shared location at the top of the hierarchy. This enables you to organize devices based on function and location without redundant configuration. For example, you could configure shared settings that are global to all firewalls, configure device groups with function-specific settings at the first level, and configure device groups with location-specific settings at lower levels (see Figure: Device Group Hierarchy). Without a hierarchy, you would have to configure both function- and location-specific settings for every device group in a single level under Shared.

- Device Group Hierarchy Inheritance and Overrides
- Create a Device Group Hierarchy

## Device Group Hierarchy Inheritance and Overrides

The device group hierarchy can have up to four levels of device groups, with Shared above them all. At the bottom level, a device group can have parent, grandparent, and great-grandparent device groups (*ancestors*), from which it inherits policies and objects. At the top level, a device group can have child, grandchild, and great-grandchild device groups (*descendants*).

**Figure: Device Group Hierarchy**

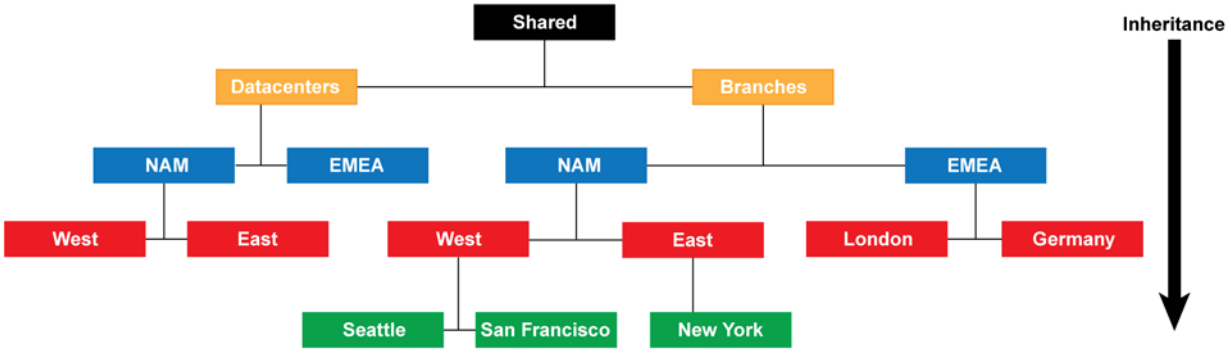


Firewalls in a device group hierarchy evaluate policies in the following order: shared pre-rules, device group pre-rules, local firewall rules, device group post-rules, shared post-rules, intrazone-default, and interzone-default. If a firewall inherits policies from device groups at different levels in the hierarchy, it evaluates pre-rules in the order of highest to lowest device group level and evaluates post-rules from lowest to highest level.

If an inherited object needs different values in a particular device group, you can override the ancestor object. If necessary, you can revert to ancestor values at any time.

 © Palo Alto Networks, Inc.

# Create a Device Group Hierarchy

| Create a Device Group Hierarchy | |
|---|---|
| Step 1 Plan the device group hierarchy. | 1. Decide the device group levels, and which firewalls and virtual systems you will assign to each device group. You can assign any one firewall or virtual system (vsys) to only one device group. If a device group will be just an organizational container for lower level device groups, you don't need to assign devices to it.<br>2. Remove firewall or vsys assignments from existing device groups if those assignments don't fit your planned hierarchy.<br>a. Select **Panorama > Device Groups** and click the device group Name.<br>b. In the Devices section, clear the check boxes of firewalls and virtual systems you want to remove, then click **OK**.<br>3. If necessary, add more managed devices that you will assign to device groups. |
| Step 2 Configure the top-level device groups. | For each top-level device group:<br>1. Select **Panorama > Device Groups** and click **Add**.<br>2. Enter a **Name** to identify the device group.<br>3. In the Devices section, select check boxes to assign devices and virtual systems to the device group.<br>4. Leave the **Parent Device Group** option at **Shared** (the default) and click **OK**. |
| Step 3 Configure lower levels of device groups. | • For new device groups at each lower level, repeat Step 2 but set the **Parent Device Group** to a device group at the previous level.<br>• For each existing device group, in the **Device Groups** page, click the device group Name to edit it, select a **Parent Device Group**, and click **OK**.<br>  If you move a device group to a different parent, all its descendant device groups move with it, along with all devices, policies, and objects associated with the device group and its descendants. If the new parent is in another access domain, the moved device group will no longer have membership in the original access domain. If the new access domain has read-write access for the parent device group, it will also have read-write access for the moved device group. If the new access domain has read-only access for the parent, it will have no access for the moved device group. To reconfigure access for device groups, see Configure an Access Domain. |

| Create a Device Group Hierarchy (Continued) | | |
|---|---|---|
| Step 4 | Configure, move, and clone objects and policy rules as necessary to account for inheritance in the device group hierarchy. | • Configure objects for use in shared or device group policy.<br>When you add an object, the **Disable Override** check box is cleared by default, which means you can override the object in descendant device groups. To disable overrides for the object, select the check box.<br>You can edit objects only at their *location*: the device group to which they are assigned. Descendant device groups inherit read-only instances of the objects from that location. To change the values of inherited objects in a descendant device group, you have to override them (Step 5).<br>• Create or edit policies.<br>• Change object and policy locations: see Move or Clone a Policy or Object to a Device Group. |
| Step 5 | Override any object that needs different values in a device group than the values it inherits from an ancestor.<br>After overriding an object, you can override it again in descendant device groups. However, you can never override shared or predefined (default) objects.<br>In the **Objects** tab, inherited objects have a green icon in the Name column, and the Location column displays the ancestor device group. | 1. In the **Objects** tab, select the object type (for example, **Objects > Addresses**).<br>2. Select the **Device Group** that will have the override instance.<br>3. Select the object and click **Override**.<br>4. Edit the values that you want to differ from the ancestor object. You can't edit the **Name** or **Shared** settings.<br>5. Click **OK**. The Name column displays a yellow-overlapping-green icon for the object to indicate it is overridden. |
| Step 6 | Save and commit your changes.<br>Perform a Panorama and device group commit after any change to the hierarchy. If a template references objects in a device group (for example, interfaces referencing addresses), and a firewall assigned to the template is no longer assigned to that device group because of a hierarchy change, perform a template commit also. | 1. Click **Commit**, for the **Commit Type** select **Panorama**, then click **Commit** again.<br>2. Click **Commit**, for the **Commit Type** select **Device Group**, select all the device groups you added or changed, then click **Commit** again. |

 © Palo Alto Networks, Inc.

# Template Stacks

You can now define a template stack, which is a combination of templates. By assigning firewalls to a stack, you can push all the necessary settings to them without the redundancy of adding every setting to every template. For example, you could assign data center firewalls in the Asia-Pacific (APAC) region to a stack that has one template with global settings, one template with APAC-specific settings, and one template with data center-specific settings. To manage firewalls in an APAC branch office, you could then re-use the global and APAC-specific templates by adding them to another stack that includes a template with branch-specific settings (see Figure: Template Stacks).

- Firewall Modes and Overlapping Settings in Stacks
- Configure a Template Stack

## Firewall Modes and Overlapping Settings in Stacks

Templates have a configurable priority order in template stacks that ensures Panorama pushes only one value for any duplicate setting. Panorama evaluates the templates listed in a stack from top to bottom, with higher templates having priority. In the following example of a data center stack, the data center template has a higher priority than the global template, so Panorama pushes the idle timeout value from the data center template and ignores the value from the global template.

**Figure: Template Stacks**

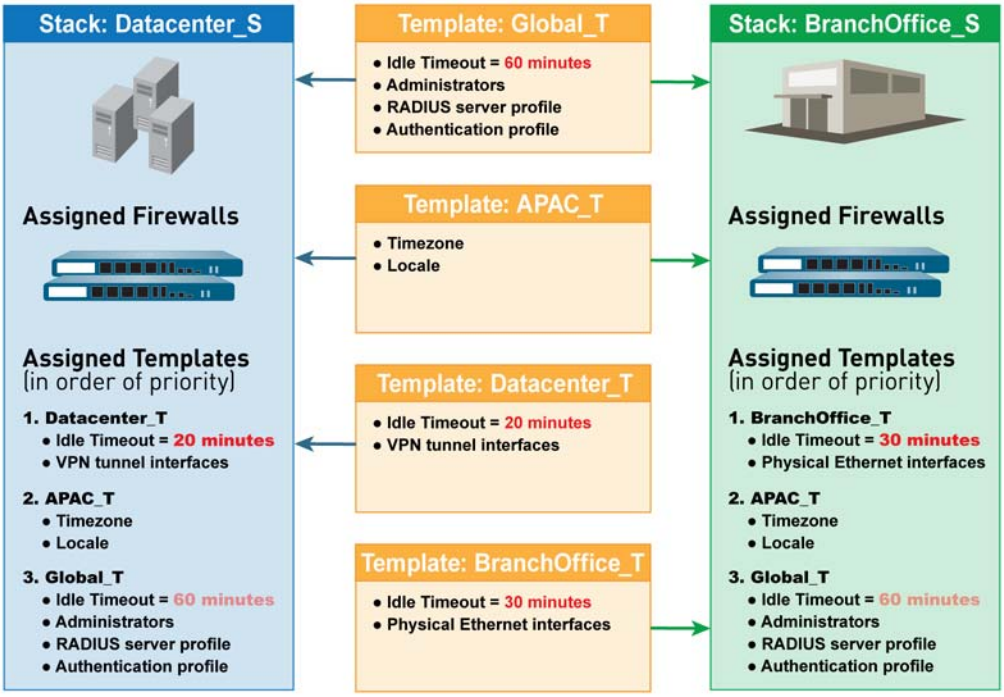


You can assign firewalls that have non-matching modes (VPN mode, multiple virtual systems mode, or operational mode) to the same template or stack. Panorama pushes mode-specific settings only to firewalls that support those modes. As an exception, you can configure Panorama to push the settings of the default virtual system in a template to firewalls that don't support virtual systems or have none configured.

# Configure a Template Stack

| Configure a Template Stack | | |
|---|---|---|
| Step 1 | Plan the template stack and configure the templates.<br>Panorama doesn't validate template combinations, so avoid ordering templates in a way that creates invalid relationships. For example, consider a stack in which the ethernet1/1 interface is Layer 3 in Template_A but is Layer 2 with a VLAN in Template_B. If Template_A has a higher priority, Panorama will push ethernet1/1 as a Layer 3 type but assigned to a VLAN. For details, see Firewall Modes and Overlapping Settings in Stacks.<br>A template configuration can't reference a configuration in another template, even if both templates are in the same stack. | In any template that has virtual system (vsys)-specific configurations that you want Panorama to push to firewalls that don't support virtual systems or have none configured, select a **Default VSYS** when you configure the template. Only one vsys can be the default in a template. The predefined vsys1 is set as the default unless you change it. If you select **None**, Panorama won't push vsys-specific configurations from that template to firewalls that don't support virtual systems or have none configured. |
| Step 2 | Create a template stack. | 1. Select **Panorama > Templates** and click **Add Stack**.<br>2. Enter a unique **Name** to identify the stack.<br>3. For each template the stack will combine (up to 16), click **Add** and select the template. The dialog lists the added templates in order of priority with respect to duplicate settings, where values in the higher templates override those that are lower in the list. To change the order, select a template and click **Move Up** or **Move Down**.<br>4. In the Devices section, select check boxes to assign firewalls. You can't assign individual virtual systems, only an entire firewall. You can assign firewalls running PAN-OS 5.0 and later. You can assign any firewall to only one template or stack. After you finish selecting, click **OK**. |

 © Palo Alto Networks, Inc.

| Configure a Template Stack (Continued) | |
|---|---|
| Step 3 Edit the **Network** and **Device** settings as necessary.<br>  While Panorama pushes mode-specific settings only to firewalls that support those modes, this selective push doesn't adjust mode-specific values. For example, if a template has firewalls in Federal Information Processing Standards (FIPS) mode and an IKE Crypto profile that uses non-FIPS algorithms, the template commit will fail. To avoid such errors, use the **Mode** drop-down in the **Network** and **Device** tabs to filter mode-specific features and value options. | 1. Depending on the settings you want to configure, select the **Network** or **Device** tab and select the stack in the **Template** drop-down. The tab settings are read-only when you select a stack.<br>2. Filter the tabs to display only the mode-specific settings you want to edit:<br>• In the **Mode** drop-down, select or clear the **Multi VSYS**, **Operational Mode**, and **VPN Mode** options.<br>• Set all the **Mode** options to reflect the mode configuration of a particular firewall by selecting it in the **Device** drop-down.<br>3. You can edit settings only at the template level, not at the stack level. To identify and access the template that contains the setting you want to edit:<br>• If the page displays a table, select **Columns > Template** in the drop-down of any column header. The Template column displays the source template for each setting. If multiple templates have the same setting, the Template column displays the higher priority template. Click the template name in this column: the **Template** drop-down changes to that template, at which point you can edit the setting.<br>• If the page doesn't display a table, hover over the template icon for a setting: a tooltip displays the source template. If multiple templates have the same setting, the tooltip displays the higher priority template. In the **Template** drop-down, select the template specified in the tooltip to edit the setting.<br>4. Edit the settings as needed.<br>5. Click **Commit**, for the **Commit Type** select **Panorama**, then click **Commit** again.<br>6. Click **Commit**, for the **Commit Type** select **Template**, select the template stack, then click **Commit** again. |

# Role-Based Access Control Enhancements

You can now map access domains to custom roles in Device Group and Template administrator accounts to enforce the separation of information among the functional or regional areas of your organization. An access domain defines access to specific device groups, templates, and firewalls (through context switching). A role defines access to specific configuration settings, logs, and reports within Panorama and firewall contexts. By combining each access domain with a role, you can control administrator access to specific information within each area. For example, you can restrict an administrator to monitoring activities for data center firewalls but allow that administrator to set policies for test lab firewalls. In this release, you can also enable a custom administrator to access more types of logs and reports, and to filter them by device group in addition to access domain.

For details on how to organize device groups and templates into functional and regional areas, see Device Group Hierarchy and Template Stacks.

The following topics describe how to configure access domains and administrators managed locally. For RADIUS management, refer to the article RADIUS Vendor-Specific Attributes (VSAs).

| Configure Role-Based Access Control | |
|---|---|
| Step 1 Configure an Access Domain.<br>Panorama supports up to 4,000 access domains. | 1. Select **Panorama > Access Domain** and click **Add**.<br>2. Enter a **Name** to identify the access domain.<br>3. (New) Select an access privilege for **Shared Objects**:<br>• **write**—Administrators can perform all operations on shared objects. This is the default value.<br>• **read**—Administrators can display and clone shared objects but cannot perform any other operations on them. When adding non-shared objects or cloning shared objects, the destination has to be a device group within the access domain, not the Shared location.<br>• **shared-only**—Administrators can add objects only to Shared. Administrators can display, edit, and delete shared objects, but cannot move or clone them. A consequence of this option is that administrators can't perform any operations on non-shared objects other than to display them. An example of why you might select option is if your organization requires all objects to be in a single, global repository.<br>4. In the **Device Groups** tab, click the icons to enable read-write or read-only access for device groups in the access domain.<br>If you set the **Shared Objects** access to **shared-only**, Panorama applies read-only access to the objects in any device groups for which you specify read-write access.<br>5. In the **Templates** tab, for each template you want to assign to the access domain, click **Add** and select the template from the drop-down.<br>6. In the **Device Context** tab, select check boxes to assign firewalls to the access domain, then click **OK**. |

 © Palo Alto Networks, Inc.

| Configure Role-Based Access Control (Continued) | | |
|---|---|---|
| Step 2 | Configure a custom role profile.<br>  For Device Group and Template roles, you can now configure access to App-Scope reports, PDF reports (including the email scheduler), custom reports, and the new Automated Correlation Engine features. | 1. Select **Panorama > Admin Roles** and click **Add**.<br>2. Enter a **Name** for the profile and select the **Role** type (**Device Group and Template**, in this example).<br>3. (New) Configure access to Panorama (**Web UI**) and firewalls (**Context Switch UI**) for each functional area by toggling the adjacent icon to the desired setting: Enable, Read Only, or Disable. In this example, be sure to enable read-write access in both tabs to all **Dashboard**, **ACC**, and **Monitor** information, in addition to any other desired features. When you finish, click **OK**. |
| Step 3 | Configure a Device Group and Template administrator and map the role to an access domain.<br>  Only Device Group and Template administrators use access domains. | 1. Select **Panorama > Administrators** and click **Add**.<br>2. Enter a user **Name** for the administrator.<br>3. Enter a **Password** and re-enter it in the **Confirm Password** field.<br>4. Select the **Administrator Type** (in this example, **Device Group and Template Admin**).<br>5. (New) In the Access Domain to Administrator Role section, add access domains and map each one to a role profile:<br>a. Click **Add** and select an Access Domain from the drop-down.<br>b. Click the adjacent Admin Role cell and select a role profile from the drop-down.<br>6. Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, then click **Commit** again. |
| Step 4 | Check that the administrator you created can filter logs and reports using the assigned access domains and device groups.<br>At the bottom of the web interface, a new **Access Domain** drop-down lists (alphabetically) the access domains. Initially it defaults to the top entry.<br>Near the top of the top of the web interface, the **Device Group** drop-down initially defaults to **All**, which specifies all device groups within the selected **Access Domain**. The drop-down lists the device groups in hierarchical order (see Device Group Hierarchy).<br>  For administrators with a predefined or custom Panorama role, **All** specifies every device group system-wide. | 1. Log in to Panorama as the Device Group and Template administrator you just created.<br>2. (New) In the **Dashboard**, select an **Access Domain** and **Device Group** and check that the information in the Applications and Logs widgets reflects your selections.<br>3. Select the **ACC** tab and check that it displays activities only within the selected **Access Domain** and **Device Group**.<br>4. Select the **Monitor** tab and check that it displays logs, correlation objects, and reports only for the selected **Access Domain** and **Device Group**.<br>  Only administrators with custom or dynamic Panorama roles can access predefined reports, and only if the **Device Group** is set to **All**.<br>Any PDF reports or report groups that you add are only available to the selected **Device Group**.<br>When you configure an email schedule, the **Email Profile** drop-down displays only profiles for which the Location is Panorama or the selected **Device Group**. If the **Device Group** is **All**, the **Email Profile** drop-down displays only the profiles for which the Location is Panorama. The email schedules you add are only available to the selected **Device Group**. |

| Configure Role-Based Access Control (Continued) | | |
|---|---|---|
| Step 5 | Configure a custom report with a device group filter.<br>For Device Group and Template administrators, the selected **Access Domain** controls which device groups are available as a query filter in custom reports. For Panorama administrators, all device groups are available. The report configuration and results are only available in the selected **Device Group**. | 1. Select **Monitor > Manage Custom Reports**.<br>2. (New) Select the **Access Domain** that contains the device groups you will query.<br>3. (New) Select the **Device Group** where the report will be available.<br>4. Click **Add** and enter a **Name** for the report.<br>5. Select a **Database** and (optionally) select the **Scheduled** check box (to run the report regularly).<br>6. Define the filtering criteria: **Time Frame**, **Sort By** order, **Group By** preference, and **Selected Columns** to display.<br>7. Add the columns you want in the report to the Selected Columns list.<br>8. Repeat the following steps for each device group query:<br>a. Select a **Connector**.<br>b. For the **Attribute**, select **Device Group**.<br>c. For the **Operator**, select **equal** or **not equal**.<br>d. For the **Value**, enter a string that contains all or part of the device group name. To specify multiple device groups that have a common character string, enter that string as a value.<br>e. Click **Add**.<br>9. Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, then click **Commit** again. |
| Step 6 | Check that the administrator can switch context to firewalls. | 1. The **Context** drop-down lists only the firewalls for which you configured **Context Switch UI** access in the role profile (Step 2). Select a firewall in the drop-down to access its web interface.<br>2. Check that you can access only the firewall features for which you configured read-write or read-only access in the role profile. |

 © Palo Alto Networks, Inc.

# Firewall Configuration Import into Panorama

You can now import firewall configurations into Panorama instead of recreating them. Panorama provides the option to import objects from the Shared location on the firewall into the Shared location in Panorama. To contain the imported network and device settings, Panorama creates a template. To contain the imported policies and non-shared objects, Panorama creates device groups: one for each virtual system (vsys) if the firewall has multiple, or a single device group otherwise. The device groups will be one level below Shared in the device group hierarchy, though you can reassign them to other parent device groups after the import (see Step 3 under Create a Device Group Hierarchy). You can import PAN-OS 5.0 or later firewall configurations.

Be sure that the content versions (for example, Applications and Threats database) on Panorama are the same as, or higher than, those on the firewall from which you will import a configuration. Otherwise, imported policy rules might reference content items (for example, an application) that Panorama doesn't have.

This procedure applies only to the initial import of firewall configuration, not to re-importing it. To re-import a configuration, refer to Migrate a Firewall to Panorama Management.

A firewall doesn't lose logs when you import its configuration into Panorama.

Migrating firewalls to Panorama management involves pre-import planning and post-import testing.

| Import a Firewall Configuration into Panorama | | |
|---|---|---|
| Step 1 | Add the firewall as a managed device, if it wasn't already added. | 1. Log in to Panorama, select **Panorama > Managed Devices** and click **Add**.<br>2. Enter the serial number of the firewall and click **OK**.<br>If you will import multiple firewall configurations, enter the serial number of each one on a separate line. This is the only step you can perform for all the firewalls at once. Perform the remaining steps in sequence for each firewall before importing the next configuration.<br>3. Click **Commit**, for the **Commit Type** select **Panorama**, then click **Commit** again. |
| Step 2 | Establish a connection from the firewall to Panorama. | 1. Log in to the firewall, select **Device > Setup**, and edit the Panorama Settings.<br>2. In the **Panorama Servers** field, enter the IP address of the Panorama management server.<br>3. Click **OK** and **Commit**. |

| Import a Firewall Configuration into Panorama (Continued) | |
|---|---|
| Step 3 (New) Import the firewall configuration into Panorama. | 1. From Panorama, select **Panorama > Setup > Operations** and click **Import device configuration to Panorama**.<br>2. Select the **Device**.<br>Panorama can't import a configuration from a firewall that is assigned to an existing device group or template.<br>3. Enter a **Template Name**. If the firewall is in multi-vsys mode, the field is blank. Otherwise, the default value is the firewall name.<br>4. (Optional) For a multi-vsys firewall, in the **Device Group Name Prefix** field, add a character string to prepend to the device group names. This can be useful for distinguishing the device groups that Panorama creates for this import from other device groups. For example, if you will import multiple firewall configurations, you can enter a prefix for each firewall to identify its associated device groups.<br>5. Edit the **Device Group** names, if desired. If the firewall is in multi-vsys mode, each device group has a vsys name by default. Otherwise, the default value is the firewall name.<br>**Import devices' shared objects into Panorama's shared context** is enabled by default, which means Panorama imports objects that belong to Shared in the firewall to Shared in Panorama. If you clear the check box, Panorama copies shared firewall objects into device groups instead of Shared. This could create duplicate objects, so selecting the check box is a best practice in most cases.<br>6. Select a **Rule Import Location** for the imported policy rules: **Pre Rulebase** or **Post Rulebase**. Regardless of your selection, Panorama imports default security rules (intrazone-default and interzone-default) into the post rulebase.<br>If Panorama has a rule with the same name as a firewall rule that you import, Panorama displays both rules. However, rule names must be unique: you have to delete one of the rules before performing a commit on Panorama.<br>7. Click **OK**. Panorama displays the import status, result, details about your selections, details about what was imported, and any warnings. Click **Close**. |

 © Palo Alto Networks, Inc.

| Import a Firewall Configuration into Panorama (Continued) | | |
|---|---|---|
| Step 4 | Fine-tune the imported configuration. | 1. In Panorama, select **Panorama > Config Audit**, select the **Running config** and **Candidate config** for the comparison, click **Go**, and review the output.<br>2. Based on the config audit, and any warnings that Panorama displayed after the import, update the device group and template configurations as needed. For example, you might have to delete redundant objects and policy rules, Move or Clone a Policy or Object to a Device Group, or move firewalls to different device groups or templates.<br>3. Click **Commit**, for the **Commit Type** select **Panorama**, then click **Commit** again. Panorama creates a device configuration bundle named <firewall_name>_import.tgz, in which all policies and objects are removed. You will use this bundle in the next step. |
| Step 5 | (New) Push the device configuration bundle to the firewall to remove all policies and objects from the local configuration. This step is necessary to prevent device group commit errors. | 1. In Panorama, select **Panorama > Setup > Operations** and click **Export or push device config bundle**.<br>2. Select the **Device** from which you imported the configuration, click **OK**, and click **Push & Commit**. Panorama pushes the bundle and initiates a commit on the firewall. |
| Step 6 | Commit your changes to the device groups and templates. | 1. In Panorama, click **Commit** and for the **Commit Type** select **Device Group**.<br>2. Select the **Merge with Device Candidate Config**, **Include Device and Network Templates** and **Force Template Values** check boxes.<br>3. Select the device groups and templates that contain the imported firewall configurations, then click **Commit**. |

# Log Redundancy Within a Collector Group

You can now enable log duplication for a Collector Group so that each log will have two copies and each copy will reside on a different Log Collector. This redundancy ensures that, if any one Log Collector becomes unavailable, no logs are lost: you can see all the logs forwarded to the Collector Group and run reports for all the log data. Panorama displays (in reports) and forwards (to external services) only one copy of each log. Note that enabling redundancy doubles the log processing traffic in a Collector Group, and therefore reduces its maximum logging rate by half, as each Log Collector must distribute a copy of each log it receives.

Because enabling redundancy creates more logs, this configuration requires more storage capacity. When a Collector Group runs out of space, it deletes older logs. To avoid losing logs, increase storage capacity by adding Log Collectors to the Collector Group. For example, if you need two Log Collectors to store four weeks of logs, after enabling redundancy you will need four Log Collectors.

To optimize log collection performance, configure the M-Series appliance to use separate interfaces for device management, device log collection, and Collector Group communication. By default, the MGT interface performs all three functions.

| Configure Log Redundancy Within a Collector Group | |
|---|---|
| Step 1 | Select **Panorama > Collector Groups** and click the name of the Collector Group. Note that log redundancy is available only if the Collector Group has multiple Log Collectors and each Log Collector has the same number of disks. |
| Step 2 | (New) Select the **Enable log redundancy across collectors** check box. |
| Step 3 | Click **OK** and **Commit**, for the **Commit Type** select **Panorama**, then click **Commit** again. |
| Step 4 | Click **Commit**, for the **Commit Type** select **Collector Group**, select the Collector Group, then click **Commit** again. A popup warns you that distributing log copies across the Log Collectors for redundancy can take hours for each terabyte of logs, and deletes older logs if current space runs out. During the distribution process, the maximum logging rate is reduced. Click **Yes** to start the process. In the **Panorama > Collector Groups** page, the Redistribution State column indicates the process status. |

 © Palo Alto Networks, Inc.

# Firewall HA State in Panorama

The Panorama web interface now displays the high availability (HA) state of firewalls (for example, active or passive) in places where knowing that state is useful. For example, to avoid downtime when installing a PAN-OS software update, you would first upgrade the passive firewall peer and then the active peer. In the web interface, the icons of firewalls that are in HA mode have colors to indicate state, as follows. Panorama updates the icons within 10 seconds of a state change.

- Green—Active (normal traffic-handling operational state).
- Yellow—Passive (normal backup state) or the firewall is initiating. A Palo Alto Networks device remains in the initiating state for up to 60 seconds after bootup.
- Red—Non-functional (an error state), suspended (an administrator disabled the firewall), or tentative (for a link or path monitoring event in an active/active configuration).

Panorama displays the HA state, and enables you to filter by that state, when you:

- Switch device context by selecting a firewall in the **Context** drop-down.
- **Commit** device group or template changes. The icons appear in an HA Status column.
- Deploy software or content updates to firewalls using the **Panorama > Device Deployment** pages or the **Panorama > Managed Devices** page (**Install** option). The icons appear in an HA Status column.
- View managed devices (**Panorama > Managed Devices**). Clear the **Group HA Peers** check box to filter by HA state; otherwise, the page displays both peers even if only one matches the state string (for example, active). The icons appear in an HA Status column.
- Assign firewalls to various configurations, for example: templates, template stacks, device groups, and Collector Groups.
- Push a firewall configuration bundle using the **Panorama > Setup > Operations** page, **Export or push device config bundle** link (see Firewall Configuration Import into Panorama).
- Filter logs or custom reports by **Device SN** (serial number), as in the following example. Note that you can't use HA state as a filter for logs and reports, only as a filter for the firewall selector popup when filtering by **Device SN**.

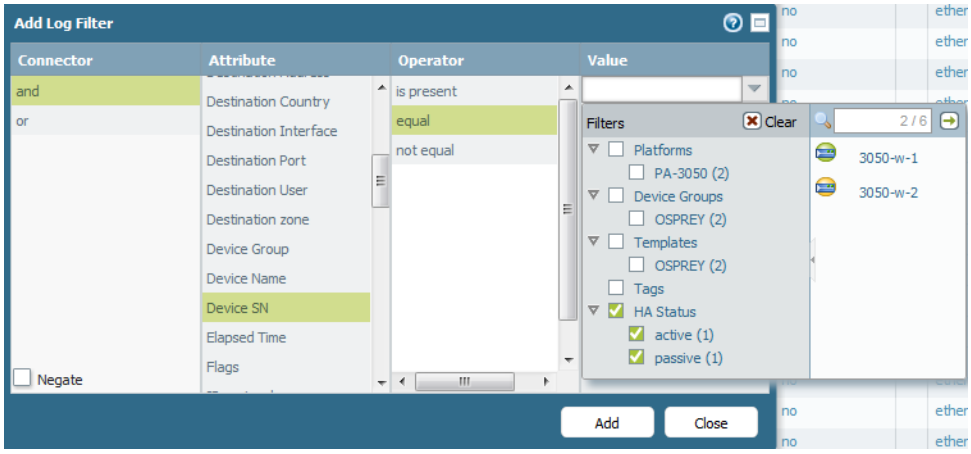

 © Palo Alto Networks, Inc.

# WildFire Features

- WildFire Grayware Verdict
- WildFire Hybrid Cloud
- WildFire Analysis Profile

To use the features introduced in PAN-OS 7.0 in a WildFire private cloud deployment, you must first upgrade the WF-500 appliance you are using to host a private cloud. You must also upgrade the WF-500 appliance before upgrading the firewalls that are configured to forward samples to it.

# WildFire Grayware Verdict

The WildFire grayware verdict classifies files that behave similarly to malware, but are not malicious in nature or intent. A grayware verdict might be assigned to files that do not pose a direct security threat, but display otherwise obtrusive behavior (for example, installing unwanted software, changing various system settings, or reducing system performance). Examples of grayware software can typically include adware, spyware, and Browser Helper Objects (BHOs). The grayware verdict allows you to quickly distinguish malicious files on the network from grayware, and to prioritize accordingly.

Antivirus signatures are not generated for grayware and security policies cannot be enforced based on the grayware verdict. However, logs and reports can continue to alert to endpoints downloading grayware, enabling you to take any necessary action.

| Enable the Grayware Verdict and Monitor Grayware | | |
|---|---|---|
| Step 1 | Change the WildFire Public Cloud setting on the firewall to point to the WildFire cloud server. | 1. On the firewall running PAN-OS 7.0 that is configured to forward samples to WildFire, select **Device > Setup > WildFire**.<br>2. Click the Edit icon and in the **WildFire Public Cloud** field, enter `wildfire.paloaltonetworks.com`. |
| Step 2 | Enable grayware reporting. | 1. Select **Device > Setup > WildFire** and edit the **General Settings**.<br>2. Select the **Report Grayware Files** check box:<br> 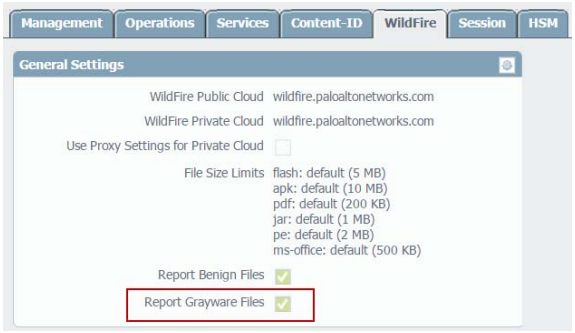 <br>3. Click **OK** and **Commit** the changes. |



 © Palo Alto Networks, Inc.

| Enable the Grayware Verdict and Monitor Grayware (Continued) | |
|---|---|
| Step 3 Monitor grayware. | Monitor Grayware Using the Firewall<br>Select **Monitor > WildFire Submissions**. The **Verdict** column displays the WildFire analysis result for logged entries.<br>Enter `( category eq grayware )` in the filter area to view logs for grayware:<br> 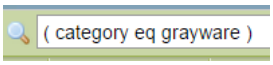 <br>Monitor Grayware Using the WildFire Portal<br>• The WildFire portal **Dashboard**, in addition to displaying statistics for malware and benign files, now includes grayware statistics for each device listed:<br> 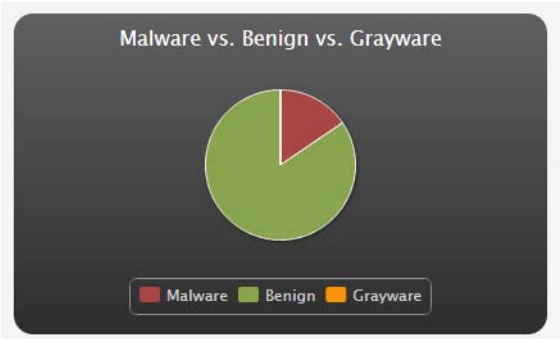 <br>• Select **Reports** and filter for WildFire reports using the Grayware Verdict:<br> 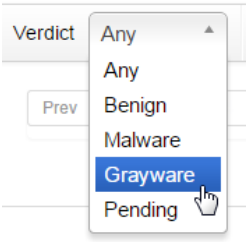 <br>• Open the WildFire Analysis Report for a file to view the Verdict for the file:<br>  |
| Step 4 Setup alerts and notifications for grayware. | Setup Alerts for Grayware From the WildFire Portal<br>Select **Settings** and select the **Grayware** checkbox for devices from which you want to receive grayware alerts.<br>Forward Firewall Grayware Logs<br>Select **Objects > Log Forwarding** and forward grayware logs as SNMP traps, syslog messages, or email notifications. |

# WildFire Hybrid Cloud

With the WildFire hybrid cloud feature, you can now configure the firewall to send files to two distinct locations for analysis: the WildFire cloud or the WF-500 appliance. Samples (files and email links) allowed by your security policy, but unknown to the firewall, can be analyzed locally on your network or in the WildFire public cloud based on the file type, application, the transmission direction of the file (upload or download). Forward certain samples to the WildFire public cloud for analysis, such as Portable Executables (PEs), while continuing to forward samples that you want to be analyzed locally to the WF-500 appliance (such as PDFs). Offloading files to the WildFire public cloud for analysis allows you to benefit from a prompt verdict for files that have been previously processed in the WildFire public cloud and frees up WF-500 appliance capacity to process sensitive content. In a WildFire hybrid cloud deployment, you can also use the WildFire public cloud to process certain file types that are not supported for analysis by the WildFire appliance, such as Android Application Package (APK) files.

The WildFire hybrid cloud feature also introduces a WildFire Analysis Profile to be used to forward files for WildFire analysis in place of the File Blocking profiles that were used in previous release versions.

## Enable WildFire Hybrid Cloud

Before you begin:

Make sure that you upgrade the WF-500 appliance to PAN-OS 7.0 before you upgrade the firewalls connected to the appliance. Both the appliance and connected firewalls must be running PAN-OS 7.0 in order to enable this feature.

| | | |
|---|---|---|
| Step 1 | Review and modify the **WildFire Analysis** profile settings that were migrated from existing **File Blocking** profiles during the upgrade to PAN-OS 7.0. | 1. Select **Objects > Security Profiles > WildFire Analysis** and review all profile rules.<br>2. Check that the **Analysis** location for samples matched to the rule is set correctly for each rule (to either **public-cloud** or **private-cloud** as needed).<br>After the upgrade to PAN-OS 7.0, all WildFire Analysis profiles are set to forward files to the **public-cloud** by default. For files that you want to be forwarded to the WildFire appliance for analysis, change the default Analysis setting for the profile to **private-cloud**.<br>  If the firewall was configured to forward files to a WF-500 appliance before the upgrade to PAN-OS 7.0, a check is in place to prevent any forwarding to the **WildFire Public Cloud** until you manually change the server setting to point to wildfire.paloaltonetworks.com. (See the next step). |

 © Palo Alto Networks, Inc.

| Enable WildFire Hybrid Cloud (Continued) | |
|---|---|
| Step 2 Review and update migrated WildFire settings following the upgrade to PAN-OS 7.0.<br>Do not perform this step until you have completed Step 1. Confirm that WildFire Analysis profiles settings are configured correctly, specifically the **Analysis** location that defines where files are forwarded for WildFire analysis. | Select **Device > Setup > WildFire** to review and edit WildFire server settings:<br>1. Update **WildFire Public Cloud** settings. For PAN-OS 7.0, enter `wildfire.paloaltonetworks.com`.<br>If the firewall was configured to forward files to a WF-500 appliance before the upgrade to PAN-OS 7.0, a check is in place to prevent the firewall from forwarding files to the public cloud, until you modify this field. Files or links matched to a WildFire analysis profile with the public-cloud setting will begin to be forwarded for WildFire public cloud analysis when this field is updated.<br>2. Review and update **WildFire Private Cloud** settings. Enter the IP address or FQDN of a WF-500 appliance. |
| Step 3 Continue to create new **WildFire Analysis** profile rules, specifying files to be forwarded to the WF-500 appliance and files to be forwarded to the WildFire public cloud. | 1. Select **Objects > Security Profiles > WildFire Analysis** and click **Add**.<br>2. Enter a **Name** for the WildFire Analysis profile, for example WildFire_default. Optionally enter a **Description**, such as Default WildFire file forwarding settings.<br>3. Enter a rule **Name**, such as local-PDF-analysis.<br>4. Select criteria to determine what type of traffic or files you want to be analyzed. File analysis can be performed based on **Applications**, **File Types**, or the transmission **Direction** of the file (upload or download). For example, select the **File Type** *PDF* and set the **Direction** to *both* to apply this rule to all PDFs identified by the firewall.<br>5. Set **Analysis** to **private-cloud** to forward files to the WF-500 appliance for analysis or to **public-cloud** to forward files to the WildFire public cloud.<br>In this example, you could select **private-cloud** to enable the firewall to forward all PDFs to the WF-500 appliance for analysis. This ensures that the PDFs are analyzed locally and any sensitive data in the PDFs is contained to the private cloud hosted by a WF-500 appliance. |

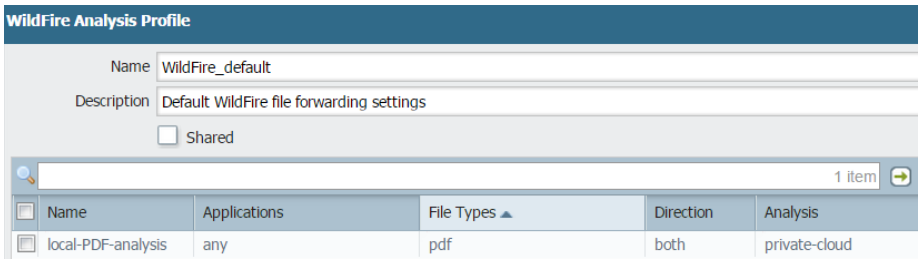

## Enable WildFire Hybrid Cloud (Continued)

Continue by adding more rules to the WildFire analysis profile to specify where to forward other types of traffic for analysis. For example, create a second rule to forward Portable Executables (PEs), flash files, and Android Application Package (APK) files to the WildFire Cloud for analysis. In this case, the option to perform file analysis in two locations enables local file analysis for sensitive data (PDFs), while file types that might be considered less sensitive such as PEs and flash files can be analyzed by the WildFire Cloud. This offers the significant benefits of increasing the WF-500 appliance capacity to process PDFs, enabling quick verdicts to be delivered for PE, flash, and APK files already analyzed by the WildFire public cloud, and enabling WildFire public cloud analysis for file types that are not supported to be analyzed by the WF-500 appliance (such as APK files).

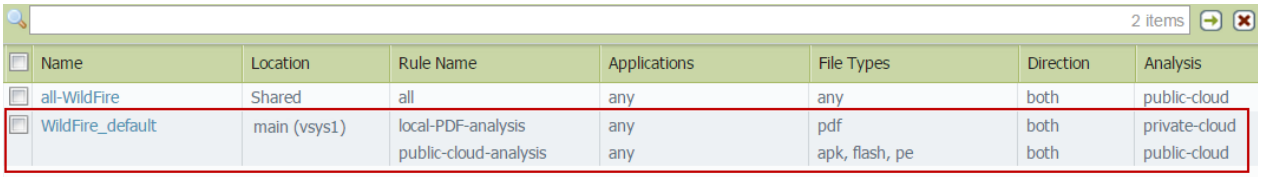

2 items

| Name | Location | Rule Name | Applications | File Types | Direction | Analysis |
|---|---|---|---|---|---|---|
| all-WildFire | Shared | all | any | any | both | public-cloud |
| WildFire_default | main (vsys1) | local-PDF-analysis | any | pdf | both | private-cloud |
| | | public-cloud-analysis | any | apk, flash, pe | both | public-cloud |

**Step 4** Attach the WildFire analysis profile rule to a security policy rule.

Traffic allowed by the security rule is then further evaluated against the WildFire analysis profile rule.

1. Select **Policies > Security** and **Add** or modify a policy.
2. Click the **Actions** tab within the policy rule.
3. In the Profile Settings section, select **Profiles** as the **Profile Type** and then select the WildFire_default profile from the **WildFire Analysis** drop-down.
4. Click **OK** and **Commit** the configuration.

**Step 5** Check that the firewall is forwarding files to the correct WildFire analysis location.

See Verify WildFire Submissions for more options to verify file forwarding in a WildFire hybrid cloud deployment, including options to:

- Confirm the firewall connection status to the WildFire public cloud and/or private cloud.
- Monitor the number of files forwarded to the WildFire public cloud and/or private cloud.
- View the number of files forwarded to the WildFire public cloud and/or private cloud according to file type.
- Check the status for a single sample forwarded by the firewall.

Wait a few minutes after performing the last step and then select **Monitor > WildFire Submission**. Check the **WildFire Cloud** field; this field displays the location to which each entry was forwarded for analysis.

Enable logging for benign and grayware files to quickly verify file forwarding, as only malware samples are logged as **WildFire Submissions** by default (**Device > Setup > WildFire**).

 © Palo Alto Networks, Inc.

| Enable WildFire Hybrid Cloud (Continued) | | |
|---|---|---|
| Step 6 | (Optional) Remove the now empty **File Blocking** profile rules that were migrated to the new **WildFire Analysis** profile.<br>It is a best practice to delete the empty file blocking rules so that the rules are not counted towards the overall security profile limit for the firewall. | Following the upgrade to PAN-OS 7.0, **File Blocking** profile rules set to **forward** or **continue and forward** display only the profile name with all other settings cleared; all other profile rule settings were migrated to the new **WildFire Analysis** profile rules during the upgrade.<br>Select **Objects > Security Profiles > File Blocking** and select and **Delete** empty File Blocking profiles rules. |

# WildFire Analysis Profile

A new WildFire Analysis profile is introduced with PAN-OS 7.0 in order to forward files and email links for WildFire analysis, replacing the need to use File Blocking profile rules to forward files for WildFire analysis. File Blocking profile rules configured to forward or to continue and forward files to WildFire are migrated to the WildFire Analysis profile during the PAN-OS 7.0 upgrade.

After upgrading, start by reviewing your WildFire settings and WildFire Analysis profile rules to ensure that your WildFire configuration is correct. Continue to modify the migrated WildFire Analysis profile rules and to add new WildFire analysis profiles rules as a part of your regular WildFire workflow.

| Create or Modify a WildFire Analysis Profile | | |
|---|---|---|
| Step 1 | Create or modify WildFire Analysis profile rules after upgrading to PAN-OS 7.0.<br>Only File Blocking profile rules with the action set to **forward** or **continue and forward** are migrated to WildFire Analysis profile rules during the PAN-OS 7.0 upgrade migration. | 1. Select **Objects > Security Profiles > WildFire Analysis**.<br>2. Review the WildFire Analysis profile rules to ensure that the **Analysis** location to which samples will be forwarded for analysis is set correctly for each profile rule (to either **public-cloud** or **private-cloud** as needed).<br>After the upgrade to PAN-OS 7.0, all WildFire Analysis profiles are set to forward files to the **public-cloud** by default. For files that you want to be forwarded to the WildFire appliance for analysis, change the default Analysis setting for the profile to **private-cloud**.<br>If the firewall was configured to forward files to a WF-500 appliance before the upgrade to PAN-OS 7.0, a check is in place to prevent any forwarding to the **WildFire Public Cloud** unless you manually change the server setting to point to wildfire.paloaltonetworks.com (**Device > Setup > WildFire**). |
| Step 2 | Verify server settings for the WildFire public and/or private cloud to which the firewall will forward samples for analysis.<br>Do not perform this step until you have reviewed your **WildFire Analysis** profile (**Objects > Security Profiles > WildFire Analysis**). Confirm that the **Analysis** location for each profile rule is correctly defined to forward samples matched to the rule to either the WildFire public cloud or a WildFire private cloud. | 1. Select **Device > Setup > WildFire**.<br>2. Update **WildFire Public Cloud** settings. For PAN-OS 7.0, enter `wildfire.paloaltonetworks.com`.<br>If the firewall was configured to forward files to a WF-500 appliance before the upgrade to PAN-OS 7.0, a check is in place to prevent the firewall from forwarding files to the public cloud, until you modify this field. Files or links matched to a WildFire analysis profile with the public-cloud setting will begin to be forwarded for WildFire public cloud analysis when this field is updated.<br>3. Review and update **WildFire Private Cloud** settings. Enter the IP address or FQDN of a WF-500 appliance. |

 © Palo Alto Networks, Inc.

| Create or Modify a WildFire Analysis Profile | |
|---|---|
| ❑ Verify that the firewall is forwarding files and email links. | Verify File Forwarding, which includes options to:<br>• Verify the status of the firewall connection to the WildFire public and/or WildFire private cloud, including the total number of files forwarded by the firewall for analysis.<br>• Verify that a specific sample was forwarded by the firewall and check that status of that sample using the troubleshooting command `debug wildfire upload-log`. (This command replaces the need in previous release versions to use the **Monitor > Data Filtering** logs to check that a single sample was forwarded). |
| ❑ View the samples that were successfully submitted for WildFire analysis. | Check the **WildFire Submissions** logs (**Monitor > WildFire Submissions**) on the firewall.<br>By default, only samples that receive malware verdicts are displayed as **WildFire Submissions** entries. To enable logging for benign and/or grayware samples, select **Device > Setup > WildFire > Report Benign Files/ Report Grayware Files**. |
| ❑ If you have enabled a WildFire hybrid cloud, confirm the WildFire location to which a sample is forwarded for analysis. | • Select **Monitor > WildFire Submissions** and check the **WildFire Cloud** column for an entry to confirm if that sample was analyzed in the WildFire public cloud or a WildFire private cloud.<br>• Verify File Forwarding to the WildFire public cloud with the CLI command debug wildfire upload-log channel public.<br>• Verify File Forwarding to the WildFire private cloud with the CLI command debug wildfire upload-log channel private. |

 © Palo Alto Networks, Inc.

# Content Inspection Features

- Configurable Drop Actions in Security Profiles
- Blocking of Encoded Content
- Negate Operator for Custom Threat Signatures

# Configurable Drop Actions in Security Profiles

When traffic matches an allow rule in security policy, the security profiles that are attached to the rule are applied for further content inspection. In PAN-OS 7.0, the Vulnerability Protection, Anti-Spyware, and Antivirus profiles provide granular actions for handling sessions when the firewall detects a threat.

▲ Actions in Security Profiles

▲ Set the Action in a Security Profile

## Actions in Security Profiles

The action specifies how the firewall responds to a threat event.

- **Default**—For each threat signature and Antivirus signature that is defined by Palo Alto Networks, a default action is specified internally. Typically the default action is an alert or a reset-both. The default action is displayed in parenthesis, for example default (alert) in the threat or antivirus signature.
- **Allow**—Permits the application traffic.
- **Alert**—Generates an alert for each application traffic flow. The alert is saved in the Threat log.
- **Drop**—Drops the application traffic.
- **Reset Client**—For TCP, resets the client-side connection. For UDP, drops the connection.
- **Reset Server**—For TCP, resets the server-side connection. For UDP, drops the connection.
- **Reset Both**—For TCP, resets the connection on both client and server ends. For UDP, drops the connection.
- **Block IP**— This action blocks traffic from either a source or a source-destination pair. It is configurable for a specified period of time.

The block action is no longer available. On upgrade from earlier PAN-OS versions, any profile that uses the block action or any decoder, threat, or virus signature with a default block action will be handled as a reset-both action in PAN-OS 7.0.

The following actions are now available by profile (the green checkmark indicates newly supported actions):

| Profiles | Actions | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Default | Allow | Alert | Drop | Reset Client | Reset Server | Reset Both | Block IP |
| **Antivirus** | | | | | | | | |
| **> Antivirus signature** | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | — |
| **> WildFire signature** | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | — |
| **> Application exception** | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | — |

 © Palo Alto Networks, Inc.

| Profiles | Actions | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Default | Allow | Alert | Drop | Reset Client | Reset Server | Reset Both | Block IP |
| **Anti-Spyware** & Vulnerability Protection | | | | | | | | |
| **> Rules** | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| **> Exceptions** | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| **Custom Objects** | | | | | | | | |
| > Spyware | – | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | – |
| **> Vulnerability Protection** | – | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | – |

The **Action on DNS Queries** in **Anti-Spyware Profiles> DNS Signatures** have not changed. The supported actions remain as default (**alert**), **block**, **allow**, **sinkhole**.

## Set the Action in a Security Profile

The following example shows how to create an anti-spyware rule with reset-client as the action for all threat signatures that are triggered on a TCP session; UDP sessions are dropped.

| Create an Anti-Spyware Profile and Set the Action | |
|---|---|
| Step 1 | Select **Objects > Security Profiles > Anti-Spyware** and **Add** a new anti-spyware profile. |
| Step 2 | Give the profile a descriptive **Name**. |
| Step 3 | If the firewall is enables for multiple virtual systems, enable the profile to be **Shared** by all virtual systems. |
| Step 4 | Click **Add**, and define a new rule. Enter a **Rule Name**; to cover all threats, set **Threat Name**, **Category**, and **Severity** to **Any**. |
| Step 5 | Set the **Action** to **Reset Client**. |
| Step 6 | Attach the profile to a security policy rule.<br>1. Select **Policies > Security**, select the desired policy rule to modify it and then click the **Actions** tab.<br>2. In **Profile Settings**, click the drop-down next to the Anti-Spyware security profile and select the profile you just created. |

# Blocking of Encoded Content

The firewall now identifies and inspects files that have been encoded or compressed up to four times by default (in previous release versions, the firewall supported two levels of decoding). As an extension of this feature, the new file type *Multi-Level-Encoding* can be used to block content that is not inspected by the firewall due to being encoded five or more times. Because multiple layers of encoding can be used as an evasion technique to circumvent security devices, use Multi-Level-Encoding in a file blocking profile to ensure that unidentified files which have not been processed for threats are not passed through the firewall.

| Block Files With Five or More Levels of Encoding | | |
|---|---|---|
| Step 1 | Create a file blocking profile. | 1. Select **Objects > Security Profiles > File Blocking** and **Add** or modify an file blocking profile.<br>2. Enter a **Name** for the file blocking profile, for example Block_High_Encoding_Levels. |
| Step 2 | Configure the file blocking options using the Multi-Level-Encoding file type. | 1. **Add** a new profile rule and give the rule a descriptive **Name**, for example, Block_Multi_ZIP.<br>2. Set **File Types** to **Multi-Level-Encoding**.<br>3. Set the **Direction** to **both**.<br>4. Set the **Action** to **block**. With this action, the firewall will block any files that it detects with five or more levels of encoding.<br>5. Click **OK** to save the profile. |
| Step 3 | Apply the file blocking profile to a security policy rule. | 1. Select **Policies > Security** and **Add** or modify a policy.<br>2. Click the **Actions** tab within the policy rule.<br>3. In the Profile Settings section, click the drop-down and select the file blocking profile you configured. In this case, the profile name is Block_High_Encoding_Levels.<br>4. **Commit** the configuration. |

 © Palo Alto Networks, Inc.

# Negate Operator for Custom Threat Signatures

A negate operator is now available when creating custom vulnerability or spyware signatures. The Negate operator can be used to ensure that the vulnerability or spyware signature is not triggered under certain conditions. When a custom signature is set up with conditions that match to traffic based on a configured data pattern, the negate operator can be used so that a signature is generated for traffic only when a pattern is not present. For example, a custom signature can be created to trigger when a Uniform Resource Identifier (URI) pattern is matched to traffic, but only when the HTTP referer field is not equal to a certain value. A custom signature must include at least one positive condition in order for a negated condition to be specified.

| Configure a Negate Condition for Custom Signatures | |
|---|---|
| Step 1 | Select **Objects > Vulnerability** and **Add** a vulnerability signature or modify an existing custom signature.<br>This example shows how to create a custom vulnerability signature with a negate operator; however, you can use the same steps to create a custom spyware signature with a negate operator (**Objects > Spyware**). |
| Step 2 | On the **Configuration** tab, enter values to set the **Threat ID**, **Name**, **Severity**, and **Direction** for the signature. |
| Step 3 | On the **Signatures** tab, **Add** a **Standard** signature.<br>Combination signatures, which are triggered depending on the number of times the signature is matched to traffic in a given period, do not support the negate operator. |
| Step 4 | Give the custom signature a descriptive **Name** and select whether to apply the signature to the current **Transaction** or the full user **Session**.<br>When the **Scope** of the signature is set to **Session**, a negate operator cannot be configured as the last condition to be matched to traffic. After creating a signature that applies to the full user session, check that a positive condition is listed as the last condition to match to traffic. |
| Step 5 | **Add Or Condition** or **Add And Condition** to the signature. You can use the negate operator for both AND and OR conditions. AND conditions match to all criteria defined for the condition. OR conditions match to any criteria defined for the condition. Use an AND condition to narrow the criteria that triggers the signature, or use an *OR* condition to broaden the criteria that triggers the signature. |
| Step 6 | Add a positive **Pattern Match Operator**, where the signature is triggered when the defined pattern is matched to traffic.<br>A custom signature cannot be configured with only Negate conditions. You must include at least one positive condition in order to then enable a Negate condition.<br>For example, select the **Context http-req-uri-path** and enter the **Pattern** `redirect`:<br> 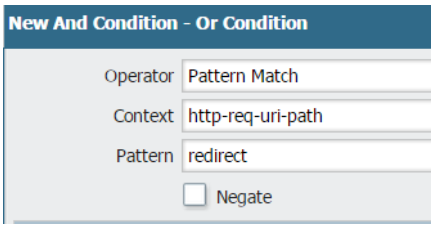 <br>The signature will be triggered for traffic when the path in a HTTP request header contains the word `redirect`.<br>Click **OK** to save the condition. |

**Configure a Negate Condition for Custom Signatures (Continued)**

Step 7 Add a second **Pattern Match Operator**, entering a **Context** and **Pattern**, and select the **Negate** check box. For example, select the **Context http-req-host-header**, enter the **Pattern** `amazon/.com` and select **Negate**:

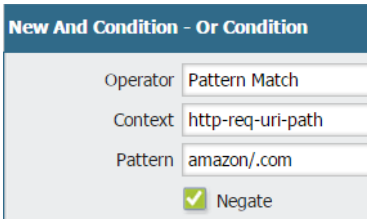


The signature will now be triggered for traffic only when the path in the HTTP request header does not contain `amazon.com`.

Click **OK** to save the condition and the signature and **Commit** your changes.

In this example, the vulnerability signature is triggered for traffic when the host field in the HTTP request header contains the word `redirect` but does not contain the string amazon.com.

| And Condition | Conditions | Operator | Context | Value | Qualifier | Negate |
|---|---|---|---|---|---|---|
| **And Condition 1** | | | | | | |
| And Condition 1 | Or Condition 1 | pattern-match | http-req-uri-path | redirect | | |
| **And Condition 2** | | | | | | |
| And Condition 2 | Or Condition 1 | pattern-match | http-req-uri-path | amazon/.com | | ☑ |

 © Palo Alto Networks, Inc.

# Authentication Features


- Authentication and Authorization Enhancements
- SSL/TLS Service Profiles
- TACACS+ Authentication
- Kerberos V5 Single Sign-On
- Suite B Cryptography Support
- Authentication Server Connectivity Testing

# Authentication and Authorization Enhancements

This release provides the following authentication and authorization enhancements:

| Enhancement | Description |
|---|---|
| Easier configuration workflow for authentication servers, profiles, and sequences | The workflow to configure server profiles, authentication profiles, and authentication sequences is simpler and more intuitive. Authentication server profiles now have only the parameters for accessing authentication services (for example, IP addresses), while authentication profiles and sequences have only the parameters for the authentication process (for example, user domain). You no longer need to set the allowed login attempts and account lockout duration for authentication sequences, only authentication profiles.<br>To enable authentication, configure a server profile (LDAP, RADIUS, Kerberos, or TACACS+) if you will use an external authentication service, configure an authentication profile (to which you assign the server profile), optionally assign multiple authentication profiles to an authentication sequence, and assign the authentication profile or sequence to administrator accounts, GlobalProtect portals or gateways, or the Captive Portal configuration. For the relevant procedures, refer to:<br>• Manage firewall administrators<br>• Configure the GlobalProtect Portal<br>• Configure GlobalProtect Gateways<br>• Map IP Addresses to User Names Using Captive Portal<br>  You can perform Authentication Server Connectivity Testing to check whether the Palo Alto Networks device can successfully communicate with the authentication server specified in an authentication profile. |
| Kerberos V5 single sign-on (SSO) support | Palo Alto Networks devices now support Kerberos V5 single sign-on, and automatic modifications to user-entered domains and usernames. This reduces the number of logins and the amount of typing required to authenticate administrators and end users. You configure these features in authentication profiles. |
| Certificate verification for LDAP connections | To improve security, you can now enable Palo Alto Network devices to verify the certificate that an LDAP server presents for SSL/TLS connections. You enable the verification in an LDAP server profile. |
| Intuitive configuration workflow for group mapping | The workflow to configure username to group mapping for User-ID is more intuitive: you set the user domain override in the group mapping configuration instead of in LDAP server profiles. |
| Delivery of VSAs from GlobalProtect clients to RADIUS servers | You can now configure the firewall send Vendor-Specific Attributes (VSAs) from GlobalProtect clients to RADIUS servers so that RADIUS administrators can perform administrative tasks based on those VSAs. For example, RADIUS administrators might use the client operating system attribute to define a policy that mandates regular password authentication for Microsoft Windows users and one-time password (OTP) authentication for Google Android users. RADIUS administrators can also log in to the firewall CLI as SSH users without first logging in to the web interface. |
| CLI commands for authentication troubleshooting | The device CLI has new commands for displaying authentication statistics, debugging authentication events, and unlocking users. |

 © Palo Alto Networks, Inc.

# SSL/TLS Service Profiles

You can now assign SSL/TLS service profiles to device services that use SSL/TLS, including Captive Portal, GlobalProtect portals and gateways, management traffic access using the web interface or XML API, URL Admin Override, and the User-ID Syslog listening service. SSL/TLS service profiles specify a certificate and the allowed protocol version or range of versions (now including TLSv1.2). By defining the protocol versions, the profiles enable you to restrict the cipher suites that are available to secure communication with the clients requesting the services. This improves network security by enabling devices to avoid SSL/TLS versions that have known weaknesses. If a service request involves a protocol version that is outside the specified range, the device downgrades or upgrades the connection to a supported version.

| Configure and Assign an SSL/TLS Service Profile | | |
|---|---|---|
| Step 1 | Configure the SSL/TLS service profile. | 1. For each desired service, Generate a Certificate.<br>Use only signed certificates for SSL/TLS services, not certificate authority (CA) certificates.<br>2. Select **Device > Certificate Management > SSL/TLS Service Profile**.<br>3. For a firewall with more than one virtual system (vsys), select the **Location** (vsys or **Shared**) where the profile is available.<br>4. Click **Add** and enter a **Name** to identify the profile.<br>5. Select the **Certificate** you just generated.<br>6. Define the range of protocols that the service can use:<br>• For the **Min Version**, select the earliest TLS version to support: **TLSv1.0**, **TLSv1.1**, or **TLSv1.2**.<br>• For the **Max Version**, select the latest TLS version to support: **TLSv1.0**, **TLSv1.1**, **TLSv1.2**, or **Max** (the latest available version.<br>7. Click **OK** and **Commit**. |
| Step 2 | Assign the profile to the desired service.<br>SSL/TLS service profiles replace certificates for all these services except the User-ID syslog listener, for which the profile is a new setting. | • Content-ID: URL Admin Override—Select **Device > Setup > Content-ID** and **Add** or edit the URL Admin Override entries.<br>• Captive Portal—Select **Device > User Identification > Captive Portal Settings** and edit the settings.<br>• GlobalProtect portals—Select **Network > GlobalProtect > Portals**, **Add** or edit a portal, select the **Portal Configuration** tab, and edit the Network Settings.<br>• GlobalProtect gateways—Select **Network > GlobalProtect > Gateways**, **Add** or edit a gateway, select the **General** tab, and edit the Network Settings.<br>• Management interface: inbound traffic—Select **Device > Setup > Management** and edit the General settings.<br>• User-ID user mapping, syslog listening service—Select **Device > User Identification > User Mapping**, edit the Palo Alto Networks User-ID Agent Setup settings, select the **Server Monitor** tab, and select a **Syslog Service Profile**. The SSL/TLS service profile secures syslog messages that the firewall parses when it performs user mapping. |

© Palo Alto Networks, Inc.

# TACACS+ Authentication

Palo Alto Networks devices now support Terminal Access Controller Access-Control System Plus (TACACS+) protocol for authenticating administrator and end users. TACACS+ provides greater security than RADIUS because it encrypts usernames and passwords (instead of just passwords), and is also more reliable (it uses TCP instead of UDP).

| Configure TACACS+ Authentication | | |
|---|---|---|
| Step 1 | Configure a TACACS+ server profile. | 1. Select **Device > Server Profiles > TACACS+** and click **Add**.<br>2. Enter a **Profile Name** to identify the server profile.<br>3. If this profile is for a firewall with multiple virtual systems capability, select a virtual system or **Shared** as the **Location** where the profile is available.<br>4. Enter a **Timeout** interval after which an authentication request times out (1-20 seconds, default 3 seconds).<br>5. Select the **Use single connection for all authentication** check box to use the same TCP session for all authentications that use this profile. This option improves performance by avoiding the need to start and end a separate TCP session for each authentication. The check box is cleared by default.<br>6. For each TACACS+ server, click **Add** and enter a **Name** (to identify the server), server IP address or FQDN (**TACACS+ Server** field), **Secret**/**Confirm Secret** (a key to encrypt usernames and passwords), and server **Port** for authentication requests (default 49).<br>7. Click **OK** and **Commit**. |
| Step 2 | Use the profile for authentication. | Assign the TACACS+ server profile to an authentication profile and assign the authentication profile to administrator accounts, GlobalProtect portals or gateways, or the Captive Portal configuration. For the specific steps, refer to:<br>• Create an administrative account<br>• Configure the GlobalProtect Portal<br>• Configure GlobalProtect Gateways<br>• Map IP Addresses to User Names Using Captive Portal |

 © Palo Alto Networks, Inc.

# Kerberos V5 Single Sign-On

Palo Alto Networks devices now support Kerberos V5 single sign-on (SSO) for authenticating administrative users to the web interface and end users to Captive Portal. A network that supports Kerberos SSO prompts a user to log in only for initial access to the network (for example, logging in to Microsoft Windows). After this initial login, the user can access any browser-based service in the network (for example, the firewall web interface) without having to log in again, until the SSO session expires. (Your Kerberos administrator sets the duration of SSO sessions.)

You enable SSO for a Palo Alto Networks device by importing a Kerberos *keytab* into an authentication profile and assigning the profile to administrator accounts or the Captive Portal configuration. A keytab is a file that contains Kerberos account information (principal name and hashed password) for the device, which is required for SSO authentication. Each authentication profile can have one keytab. If SSO authentication fails, the device prompts the user to log in manually and performs authentication of the type specified in the profile (for example, RADIUS).

- Configure Kerberos SSO for Administrator Authentication
- Configure Kerberos SSO for Captive Portal Authentication

## Configure Kerberos SSO for Administrator Authentication

Configuring Kerberos SSO for administrator authentication has the following prerequisites:

- Your network has a Kerberos infrastructure set up, including a Key Distribution Center (KDC) with an Authentication Server (AS) and Ticket Granting Service (TGS).
- The Palo Alto Networks device that will authenticate users has a Kerberos account, including a principal name and password. You will use these to create the keytab.

| Configure Kerberos SSO for Administrator Authentication | | |
|---|---|---|
| Step 1 | Create a Kerberos keytab for the Palo Alto Networks device.<br>  <br>If the device is in FIPS or CC mode, the algorithm must be `aes128-cts-hmac-sha1-96` or `aes256-cts-hmac-sha1-96`. Otherwise, you can also use `des3-cbc-sha1` or `arcfour-hmac`. To use an AES algorithm, the functional level of the KDC must be Windows Server 2008 or later and you must enable AES encryption for the device account.<br>The algorithm in the keytab must match the algorithm in the service ticket that the TGS issues to clients. Otherwise, the SSO process will fail. Your Kerberos administrator determines which algorithms the service tickets use. | 1. Log in to the KDC and open a command prompt.<br>2. Enter the following command, where <principal_name>, <password>, and <algorithm> are variables. The Kerberos principal name and password are of the device, not the user.<br>**ktpass /princ** `<principal_name>` **/pass** `<password>` **/crypto** `<algorithm>` **/ptype KRB5_NT_PRINCIPAL /out** `<file_name>`**.keytab**<br>For example:<br>**/princ host/computer.abrealm.com@ABREALM.COM /pass securepwd /crypto des3-cbc-sha1 /ptype KRB5_NT_PRINCIPAL /out device.keytab** |
| Step 2 | (External authentication only) Configure an authentication server profile to authenticate administrators if the SSO process fails. | • Configure a RADIUS Server Profile.<br>• Configure an LDAP Server Profile.<br>• Configure a Kerberos Server Profile.<br>• Configure a TACACS+ server profile. |
| Step 3 | Configure an authentication profile. | 1. Select **Device > Authentication Profile**, click **Add**, and enter a **Name** to identify the authentication profile.<br>2. (External authentication only) If the authentication **Type** is an external server, assign the **Server Profile** you just created.<br>3. (New) Enter the **Kerberos Realm** and import the **Kerberos Keytab** you just created. Usually the realm is the DNS domain name of the administrator, except that the realm name is uppercase.<br>4. In the **Advanced** tab, **Add** the administrator user, user group, or **all** to the **Allow List**, and then click **OK**. |
| Step 4 | Assign the authentication profile to the administrator account. | 1. Configure the administrator account.<br>2. Click **Commit**. |

 © Palo Alto Networks, Inc.

# Configure Kerberos SSO for Captive Portal Authentication

Configuring Kerberos SSO for Captive Portal authentication has the following prerequisites:

- Your network has a Kerberos infrastructure set up, including a Key Distribution Center (KDC) with an Authentication Server (AS) and Ticket Granting Service (TGS).
- The Palo Alto Networks device that will authenticate users has a Kerberos account, including a principal name and password. You will use these to create the keytab.
- You have set up the initial Captive Portal and policy configurations. In the following procedure, you modify the configurations by configuring redirect mode.

If you configure both Kerberos SSO and NT LAN Manager (NTLM) authentication, the firewall tries Kerberos SSO authentication first, and if that fails, falls back to NTLM authentication for Microsoft Windows clients. If NTLM authentication also fails, the firewall falls back to web form or client certificate authentication, depending on your Captive Portal configuration.

| Configure Kerberos SSO for Captive Portal Authentication | |
|---|---|
| Step 1 Create a Kerberos keytab for the redirect host you will assign to the Captive Portal configuration. The redirect host is the intranet hostname that resolves to the IP address of the firewall Layer 3 interface to which you are redirecting requests that match a Captive Portal policy.<br>  <br>If the device is in FIPS or CC mode, the algorithm must be `aes128-cts-hmac-sha1-96` or `aes256-cts-hmac-sha1-96`. Otherwise, you can also use `des3-cbc-sha1` or `arcfour-hmac`. To use an AES algorithm, the functional level of the KDC must be Windows Server 2008 or later and you must enable AES encryption for the device account.<br>The algorithm in the keytab must match the algorithm in the service ticket that the TGS issues to clients to enable single sign-on (SSO). Otherwise, the SSO process will fail. Your Kerberos administrator determines which algorithms the service tickets use. | 1. Log in to the KDC and open a command prompt.<br>2. Enter the following command, where <principal_name>, <password>, and <algorithm> are variables. The Kerberos principal name and password are of the device, not the user.<br>**ktpass /princ** `<principal_name>` **/pass** `<password>` **/crypto** `<algorithm>` **/ptype KRB5_NT_PRINCIPAL /out** `<file_name>`**.keytab**<br>For example:<br>**/princ host/computer.abrealm.com@ABREALM.COM /pass securepwd /crypto des3-cbc-sha1 /ptype KRB5_NT_PRINCIPAL /out device.keytab** |

| Configure Kerberos SSO for Captive Portal Authentication (Continued) | | |
|---|---|---|
| Step 2 | Enable Captive Portal redirect requests on a Layer 3 interface. | 1. Create an Interface Management profile so the interface can display Captive Portal response pages:<br>a. Select **Network > Network Profiles > Interface Mgmt** and click **Add**.<br>b. Enter a **Name** for the profile, select **Response Pages**, and click **OK**.<br>2. Assign the Interface Management profile to a Layer 3 interface:<br>a. Select **Network > Interfaces > Ethernet** and click the Name of a Layer 3 interface.<br>b. Select **Advanced > Other**, select the **Management Profile** you just created, and click **OK**.<br>3. Create DNS address (A) and pointer (PTR) records that map the IP address on the Layer 3 interface to the redirect host. |
| Step 3 | To transparently redirect users without displaying certificate errors, install a certificate that matches the IP address of the interface in Step 2. You can generate and import a self-signed certificate (as described in this example) or import a certificate that an external certificate authority (CA) signed. | To use a self-signed certificate, create a root CA certificate and use it to sign the certificate you will use for Captive Portal:<br>1. Select **Device > Certificate Management > Certificates > Device Certificates**.<br>2. Generate a root CA certificate. Be sure to select the **Certificate Authority** check box.<br>3. Generate a certificate to use for Captive Portal. Be sure to configure the following fields:<br>• **Common Name**—Enter the DNS name of the intranet host for the Layer 3 interface.<br>• **Signed By**—Select the CA you created in the previous step.<br>• Certificate Attributes—Click **Add**, for the **Type** select **IP**, and for the **Value** enter the IP address of the Layer 3 interface.<br>4. Configure clients to trust the certificate:<br>a. Select the CA certificate you just created and click **Export**.<br>b. Select a **File Format** and click **OK**.<br>c. Select **Save File** and click **OK** to download the certificate to your default download folder.<br>d. Import the certificate as a trusted root CA into all client browsers, either by manually configuring the browser or by adding the certificate to the trusted roots in an Active Directory Group Policy Object (GPO). |
| Step 4 | (External authentication only) Configure an authentication server profile to authenticate users if the SSO process fails. | • Configure a RADIUS Server Profile.<br>• Configure an LDAP Server Profile.<br>• Configure a Kerberos Server Profile.<br>• Configure a TACACS+ server profile. |

 © Palo Alto Networks, Inc.

| Configure Kerberos SSO for Captive Portal Authentication (Continued) | | |
|---|---|---|
| Step 5 | Configure an authentication profile for Captive Portal. | 1. Select **Device > Authentication Profile**, click **Add**, and enter a **Name** to identify the authentication profile.<br>2. (External authentication only) If the authentication **Type** is an external server, assign an authentication **Server Profile**.<br>3. (New) Enter the **Kerberos Realm** and import the **Kerberos Keytab** you just created. Usually the realm is the DNS domain name of the users, except that the realm name is uppercase.<br>4. In the **Advanced** tab, **Add** the users, user groups, or **all** to the **Allow List**, then click **OK**. |
| Step 6 | Assign the authentication profile to the Captive Portal configuration. | 1. Select **Device > User Identification > Captive Portal Settings** and edit the settings.<br>2. Select the **Authentication Profile** you just created.<br>3. Set the **Mode** to **Redirect** and enter the FQDN of the **Redirect Host**.<br>4. Click **OK** to save the configuration. |
| Step 7 | Configure the Captive Portal policy rules for Kerberos SSO. | 1. Perform the following steps for each Captive Portal policy rule that will use Kerberos SSO.<br>a. Select **Policies > Captive Portal** and click the Name of the rule.<br>b. Select the **Actions** tab, select **browser-challenge** in the drop-down, and click **OK**.<br>2. Click **Commit**. |

# Suite B Cryptography Support

You can now use Suite B ciphers to authenticate administrators and to secure site-to-site VPN, remote access VPN, and Large Scale VPN (LSVPN). For the VPN tunnels between GlobalProtect gateways and clients, the ciphers are available in a new GlobalProtect IPSec Crypto profile. Suite B ciphers enable you to meet U.S. federal network security standards. The following topics describe the ciphers and how to implement them:

- Suite B Ciphers
- Generate and Assign ECDSA Certificates
- Configure a GlobalProtect IPSec Crypto Profile

## Suite B Ciphers

| Cipher | Purpose/Benefits | Where Implemented |
|---|---|---|
| Elliptic Curve Digital Signature Algorithm (DSA) algorithm | Use ECDSA to generate certificate keys. See the list of features for which you can Generate and Assign ECDSA Certificates.<br>ECDSA uses smaller key sizes than the RSA algorithm, and therefore provides a performance enhancement for processing SSL/TLS connections. ECDSA also provides equal or greater security than RSA.<br>The **Device > Certificate Management > Certificates** page has a new Algorithm column to indicate the key type for existing certificates. | Certificate generation<br>ECDSA is recommended for client browsers and operating systems that support it. Otherwise, select RSA for compatibility with legacy browsers and operating systems. |
| Elliptic Curve Diffie–Hellman (DH) Group 19 (256-bit) and Group 20 (384-bit) | Use these Internet Key Exchange (IKE) options to secure VPN tunnels based on IPSec. | IKE Crypto profiles<br>IPSec Crypto profiles |
| Advanced Encryption Standard (AES) Galois/Counter Mode (GCM) algorithms: 128-bit (aes-128-gcm) and 256-bit (aes-256-gcm) | Use these algorithms to secure traffic in remote access VPN, site-to-site VPN, and Large Scale VPN deployments.<br>The names of new and existing AES algorithms now indicate their mode:<br>• gcm—Galois/Counter Mode<br>• cbc—Cipher-Block Chaining<br>• ccm—Counter with CBC-MAC | IPSec Crypto profiles<br>GlobalProtect IPSec Crypto profiles |


© Palo Alto Networks, Inc.

# Generate and Assign ECDSA Certificates

Palo Alto Networks devices now support certificates that use **Elliptical Curve DSA (ECDSA)**, a Suite B key generation **Algorithm**. Previously, the only available **Algorithm** was **RSA**.

OCSP responders can verify the revocation status of both ECDSA and RSA certificates.
You can't use ECDSA keys in outbound or inbound SSL/TLS decryption certificates, or for public key authentication for administrators.

| Generate and Assign Certificates | | |
|---|---|---|
| Step 1 | Generate a Certificate. | 1. Select **Device > Certificate Management > Certificates** and click **Generate**.<br>2. Enter a **Certificate Name**.<br>3. For the **Common Name**, enter an IP address or FQDN.<br>4. In the **Signed By** drop-down, select the certificate that will authenticate the certificate you are creating or select **External Authority** to generate a certificate signing request (CSR).<br>5. (New) For the key generation **Algorithm**, select **RSA** (the default option) or **Elliptical Curve DSA** (recommended for client browsers and operating systems that support ECDSA).<br>6. For the **Number of Bits**, select the certificate key length. Higher numbers are more secure but require more processing time. The available options depend on the key generation **Algorithm**:<br>• **RSA**—**512**, **1024**, **2048** (default), or **3072**<br>• **Elliptical Curve DSA**—**256** or **384**<br>7. Select the **Digest** algorithm. From most to least secure, the options are: **sha512**, **sha384**, **sha256** (default), **sha1**, **md5**.<br>8. For the **Expiration**, enter the number of days (default 365) for which the certificate is valid.<br>9. Click **Generate**, **OK**, and **Commit**. |
| Step 2 | Assign the certificate to a certificate profile if you will assign certificate profiles to device services. | If the certificate profile will have multiple certificates, repeat Step 1 for each certificate before creating the profile. |

| Generate and Assign Certificates (Continued) | |
|---|---|
| Step 3 Assign the certificate or certificate profile to the desired device services. | • Administrator authentication—Configure Certificate-based Authentication for the WebUI.<br>• Captive Portal authentication—Select **Device > User Identification > Captive Portal Settings**, edit the settings, and assign a **Certificate Profile**.<br>• (New) SSL/TLS services—See SSL/TLS Service Profiles.<br>• Site-to-site VPN:<br>  • IKE gateways—Select **Network > Network Profiles > IKE Gateways**, **Add** or edit an IKE gateway, and in the **General** tab select either or both of the following:<br>  – Select **Certificate** as the **Authentication** type, then select a **Local Certificate**.<br>  – **Certificate Profile**<br>  • IPSec tunnels—Select **Network > IPSec Tunnels**, **Add** or edit an IPSec tunnel, set the **Type** to **GlobalProtect Satellite**, select the **Advanced** tab, select the **External Certificate Authority** check box, and select a **Local Certificate** and/or **Certificate Profile**.<br>• Large Scale VPN (LSVPN) or remote access VPN:<br>  • GlobalProtect portals—Select **Network > GlobalProtect > Portals**, **Add** or edit a portal, and select either or both of the following:<br>  – **Portal Configuration** tab—Select a **Client Certificate** or **Certificate Profile** for the portal server.<br>  – **Agent Configuration** tab—Add a **Trusted Root CA**.<br>  – **Satellite Configuration** tab (LSVPN only)—Add a **Trusted Root CA** and/or select an **Issuing Certificate**. Besides using an issuing certificate, you can also manually import certificates into satellite clients.<br>  • GlobalProtect gateway—Select **Network > GlobalProtect > Portals**, **Add** or edit a gateway, select the **General** tab, and select a **Certificate Profile** for the gateway server. |

 © Palo Alto Networks, Inc.

## Configure a GlobalProtect IPSec Crypto Profile

Use GlobalProtect IPSec Crypto profiles to secure the VPN tunnels between GlobalProtect clients and gateways. To use Suite B ciphers, the clients must run GlobalProtect agent/app v2.2 or later releases. The following procedure assumes you already set up your GlobalProtect infrastructure and will assign a new GlobalProtect IPSec Crypto profile to an existing gateway configuration. The firewall has a default (predefined) profile that uses **aes-128-cbc** encryption and **sha1** authentication.

| Configure an GlobalProtect IPSec Crypto Profile | | |
|---|---|---|
| Step 1 | Create the GlobalProtect IPSec Crypto profile. | 1. Select **Network > Network Profiles > GlobalProtect IPSec Crypto**, click **Add** and enter a **Name** to identify the profile.<br>2. To specify the Encryption and Authentication algorithms that the VPN peers use to negotiate the keys for securing data traversing the tunnel, click **Add** in the corresponding sections and select from the drop-downs, then click **OK** to save the profile.<br>If you are not certain of what the VPN peers support, you can add multiple encryption algorithms in the order of most-to-least secure as follows: **aes-256-gcm**, **aes-128-gcm**, **aes-128-cbc**. The Suite B options are **aes-256-gcm** and **aes-128-gcm**. The peers will use the strongest supported algorithm to establish the tunnel. |
| Step 2 | Assign the profile to the client configuration of a GlobalProtect gateway. | 1. Select **Network > GlobalProtect > Gateways**, and click the Name of a gateway to edit it.<br>2. Select the **Client Configuration** tab and select the **GlobalProtect IPSec Crypto** profile you just created.<br>3. Click **OK** and **Commit**. |

# Authentication Server Connectivity Testing

You can now test an authentication profile to determine if your firewall or Panorama management server can communicate with a backend authentication server and if the authentication request was successful. You can test administrator authentication for Panorama and for PAN-OS. You can additionally test authentication profiles used for GlobalProtect and Captive Portal authentication. You can perform authentication tests on the candidate configuration, so that you know the configuration is correct before committing.

Authentication server connectivity testing is supported for local database, RADIUS, TACACS+, LDAP, and Kerberos authentication.

| Run the Test Authentication Command | |
|---|---|
| Step 1 | On the PAN-OS firewall or Panorama server, configure an authentication profile. You do not need to commit the authentication or server profile configuration prior to testing. |
| Step 2 | Using a terminal emulation application, such as PuTTY, launch an SSH session to the firewall. |
| Step 3 | (Firewalls with virtual systems configured) Define the target virtual system that the test command will access.<br>This is required on firewalls with multiple virtual systems (vsys) configured, so the test authentication command can locate the user (GlobalProtect or Captive Portal, for example) in the correct vsys.<br>To define the target vsys:<br>`admin@PA-3060>` **set system setting target-vsys <vsys-name>**<br>For example, if the user is defined in vsys2, run the following command:<br>`admin@PA-3060>` **set system setting target-vsys vsys2**<br>The `target-vsys` command is per-login session, so the system clears the option when you log off. |
| Step 4 | Test an authentication profile by entering the following command:<br>`admin@PA-3060>` **test authentication authentication-profile <authentication-profile-name> username <username> password**<br>For example, to test an authentication profile named my-profile for a user named bsimpson, run the following command:<br>`admin@PA-3060> test authentication authentication-profile my-profile username bsimpson password`<br>When entering authentication profile names and server profile names in the test command, the names are case sensitive. Also, if the authentication profile has a username modifier defined, you must enter the modifier with the username. For example, if you add the username modifier %USERINPUT%@%USERDOMAIN% for a user named bsimpson and the domain name is mydomain.com, enter bsimpson@mydomain.com as the username. This will ensure that the correct credentials are sent to the authentication server. In this example, mydomain.com is the domain that you define in the **User Domain** field in the Authentication profile. |

 © Palo Alto Networks, Inc.

**Run the Test Authentication Command**

Step 5 View the output of the test results.

If the authentication profile is configured correctly, the output displays `Authentication succeeded`. If there is a configuration issue, the output displays information to help you troubleshoot the configuration.

The output results vary based on several factors related to the authentication type that you are testing as well as the type of issue. For example, RADIUS and TACACS+ use different underlying libraries, so the same issue that exists for both of these types will produce different errors. Also, if there is a network problem, such as using an incorrect port or IP address in the authentication server profile, the output error is not specific. This is because the test command cannot perform the initial handshake between the firewall and the authentication server to determine details about the issue.

 © Palo Alto Networks, Inc.

# Decryption Features

▲ SSL Decryption Enhancements

# SSL Decryption Enhancements

When using SSL decryption to inspect and enforce security rules for connections between clients and destination servers, the following new decryption policy and decryption profile options are available as increased security measures:

### New Decryption Profile Options

- Enforce the use of strong cipher suites for decrypted traffic. This includes support to specifically enforce the use of the Suite B ciphers aes-128-gcm and aes-256-gcm.
- Enforce the use of minimum and maximum protocol versions.
- Enforce certificate validation on a per-policy basis (where previously, certificate validation was performed at the device level).

### New Decryption Policy Option

- Define traffic to be decrypted based on TCP port numbers. This enables you to apply different decryption policies to a single server's traffic; traffic being transmitted using different protocols can receive different treatment.

The following steps highlight how decrypt to traffic based on TCP port number and how to enforce certificate validation, protocol version, and the use of strong cipher suites for decrypted traffic.

| Enable SSL Decryption Enhancements | |
|---|---|
| Before you get started, confirm that you have completed the preliminary steps for configuring decryption. Depending on the type of decryption you plan to enable, this might include setting up the certificates that the firewall uses to perform decryption. | The following steps are useful before setting up a decryption profile:<br>❑ Ensure that the appropriate interfaces are configured as either virtual wire, Layer 2, or Layer 3 interfaces. Decryption can only be performed on these types of interfaces (**Network > Interfaces > Ethernet**).<br>❑ To perform SSL Forward Proxy, you must first configure a forward trust certificate and a forward untrust certificate to be presented to clients.<br>❑ To perform SSL Inbound Inspection, first check that the targeted server certificate is installed on the firewall. |
| Step 1 Create a decryption profile. | Select **Objects > Decryption Profile** and **Add** a decryption profile to block and control specific aspects of decrypted traffic. |


© Palo Alto Networks, Inc.

| Enable SSL Decryption Enhancements (Continued) | |
|---|---|
| • (New) Enable certificate validation for decrypted SSL traffic. | Select **SSL Decryption > SSL Forward Proxy** and choose from the following options to control server certificates:<br>• **Block sessions with expired certificates**—Terminate the SSL connection if the server certificate is expired. This will prevent a user from being able to accept an expired certificate and continuing with an SSL session.<br>• **Block sessions with untrusted issuers**—Terminate the SSL session if the server certificate issuer is untrusted.<br>• (New) **Block sessions with unknown certificate status**—Terminate the SSL session if a server returns a certificate revocation status of `unknown`. Certificate revocation status indicates if trust for the certificate has been or has not been revoked.<br>• (New) **Block sessions on the certificate status check timeout**—Terminate the SSL session if the certificate status cannot be retrieved within the amount of time that the firewall is configured to stop waiting for a response from a certificate status service. You can configure **Certificate Status Timeout** value when creating or modifying a certificate profile (**Device > Certificate Management > Certificate Profile**).<br>• **Restrict certificate extensions**—Limits the certificate extensions used in the dynamic server certificate to key usage and extended key usage. |
| • (New) Enable certificate validation for traffic that is not decrypted (traffic that is matched to a decryption policy with a **No Decrypt** action). | Control server certificates for traffic that is not decrypted:<br>• **Block sessions with expired certificates**—Terminate the SSL connection if the server certificate is expired. This will prevent a user from being able to accept an expired certificate and continuing with an SSL session.<br>• **Block sessions with untrusted issuers**—Terminate the SSL session if the server certificate issuer is untrusted. |
| • (New) Enforce protocol versions for decrypted SSL traffic. | Select **SSL Decryption > SSL Protocol Settings** and select a minimum or a maximum protocol version to enforce for decrypted traffic:<br>• **Min Version**—Set the minimum protocol version that can be used to establish the SSL connection.<br>• **Max Version**—Set the maximum protocol version that can be used to establish the SSL connection. You can choose **Max** so that no maximum version is specified; in this case, protocol versions that are equivalent to or are a later version than the selected minimum version are supported. |
| • (New) Enforce the use of selected encryption and authentication algorithms for an SSL session. This includes support to specifically enforce the use of the Suite B Ciphers aes-128-gcm and aes-256-gcm for a decrypted session. | Continue on the **SSL Decryption > SSL Protocol Settings** tab and select specific encryption and authentication algorithms to enforce for SSL traffic. |

| Enable SSL Decryption Enhancements (Continued) | | |
|---|---|---|
| Step 2 | Configure a decryption policy to identify traffic to be decrypted.<br>This includes attaching the profile you just created to the policy, so that traffic matched to the policy is enforced according to the profile settings. | 1. Select **Policies > Decryption** and **Add** or modify a policy.<br>2. Give the policy a descriptive **Name** and use the **Source** and **Destination** tabs to identify traffic to be decrypted.<br>3. Continue to use the **Service/URL Category** tab to decrypt traffic based on **URL Category** or the **Service** used.<br>(New) **Add** or select a **Service** to decrypt traffic based on specific TCP port numbers.<br>4. On the **Options** tab, select **Decrypt** and select the **Type** of decryption to perform.<br>5. Select the **Decryption Profile** you created to enforce cipher suites, protocol versions, and certificate validation, attaching the profile to the policy.<br>6. Click **OK** to save. |
| Step 3 | Save your changes. | **Commit** the configuration. |

 © Palo Alto Networks, Inc.

# User-ID Features

- User Attribution Based on X-Forwarded-For Headers
- Custom Groups Based on LDAP Filters

# User Attribution Based on X-Forwarded-For Headers

You can now configure User-ID to read the IPv4 or IPv6 addresses of users from the X-Forwarded-For (XFF) header in client requests for web services when the firewall is deployed between the Internet and a proxy server that would otherwise hide the user IP addresses. User-ID matches the IP addresses with usernames that your policies reference so that those policies can control and log access for the associated users and groups. XFF user attribution applies only to HTTP traffic, and only if the proxy server supports the XFF header. If the header has an invalid IP address, User-ID uses that IP address as a username for group mapping references in policies. If the header has multiple IP addresses, User-ID uses the first entry from the left.

To use XFF header values, you must also configure user mapping and the desired policies.

| Configure User Attribution Based on X-Forwarded-For Headers | |
|---|---|
| Step 1 | Select **Device > Setup > Content-ID** and edit the X-Forwarded-For Headers settings. |
| Step 2 | Select the **X-Forwarded-For Header in User-ID** check box.<br>Selecting the **Strip-X-Forwarded-For Header** check box doesn't disable the use of XFF headers for user attribution; the firewall zeros out the XFF value only after using it for user attribution. |
| Step 3 | Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

# Custom Groups Based on LDAP Filters

You can now define custom groups based on LDAP filters so that you can base firewall policies on user attributes that don't match existing user groups in an LDAP-based service such as Active Directory (AD). Defining custom groups can be quicker than creating new groups or changing existing ones on an LDAP server, and doesn't require an LDAP administrator to intervene. User-ID maps all the LDAP directory users who match the filter to the custom group. For example, you might want a security policy rule that allows contractors in the Marketing Department to access social networking sites. If no AD group exists for that department, you can configure an LDAP filter that matches users for whom the LDAP attribute Department is set to Marketing. Log queries and reports that are based on user groups will include custom groups. You can add custom groups to the Allow List of authentication profiles.

| Configure a Custom Group Based on an LDAP Filter | |
|---|---|
| Step 1 | Configure an LDAP server profile if you haven't already. |
| Step 2 | Select **Device > User Identification > Group Mapping Settings** and click **Add**. |
| Step 3 | Enter a **Name** to identify the group mapping configuration. |
| Step 4 | In the **Server Profile** tab, select the LDAP **Server Profile** and select the **Enabled** check box to enable this group mapping configuration. |
| Step 5 | Select the **Custom Group** tab and click **Add**. |
| Step 6 | Enter a group **Name** that is unique in the group mapping configuration for the current firewall or virtual system. If the **Name** has the same value as the Distinguished Name (DN) of an existing AD group domain, the firewall uses the custom group in all references to that name (for example, in policies and logs). |
| Step 7 | Specify an **LDAP Filter** of up to 2,048 UTF-8 characters, then click **OK**. The firewall doesn't validate LDAP filters.<br>  To optimize LDAP searches and minimize the performance impact on the LDAP directory server, use only indexed attributes in the filter. |
| Step 8 | Click **OK** and **Commit**. A commit is necessary for the group to be available in policies and objects. |

 © Palo Alto Networks, Inc.

# Virtualization Features

The following features were introduced for the VM-Series firewall in PAN-OS 7.0:

- Support for High Availability on the VM-Series Firewall
- High Availability for VM-Series in AWS
- Support for Jumbo Frames
- Support for Hypervisor Assigned MAC Addresses

# Support for High Availability on the VM-Series Firewall

High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration is synchronized to prevent a single point of failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up the firewalls in a two-device cluster provides redundancy and allows you to ensure business continuity. In an HA configuration on the VM-Series firewalls, both peers must be deployed on the same type of hypervisor, have identical hardware resources (such as CPU cores/network interfaces) assigned to them, and have the set same of licenses/subscriptions. For general information about HA on Palo Alto Networks firewalls, see High Availability.

The VM-Series firewalls support stateful active/passive or active/active high availability with session and configuration synchronization. The only exceptions are the following:

- The VM-Series firewall in the Amazon Web Services (AWS) cloud supports active/passive HA only. For details, see High Availability for VM-Series in AWS.
- HA is not relevant for the VM-Series NSX Edition firewall.

The active/active deployment is supported in virtual wire and Layer 3 deployments.

| Features/ Links Supported | ESX | KVM | Xen | AWS | NetX |
|---|---|---|---|---|---|
| Active/ Passive HA | Yes | Yes | Yes | Yes | No |
| Active/ Active HA | Yes | Yes | Yes | No | No |
| HA 1 | Yes | Yes | Yes | Yes | No |
| HA2—(session synchronization and keepalive) | Yes | Yes | Yes | Yes | No |
| HA3 | Yes | Yes | Yes | No | No |

## HA Timers on the VM-Series Firewalls

The following table describes each timer included in the HA Timers profiles and the current preset values for VM-Series firewalls:

| HA Timers | Default values for Recommended/Aggressive profiles |
|---|---|
| Promotion hold time | 2000/500 ms |
| Hello interval | 8000/8000 ms |
| Heartbeat interval | 1000/1000 ms |
| Max number of flaps | 3/3 |
| Preemption hold time | 1/1 min |

 © Palo Alto Networks, Inc.

| HA Timers | Default values for Recommended/Aggressive profiles |
|---|---|
| Monitor fail hold up time | 0/0 ms |
| Additional master hold up time | 500/500 ms |

For instructions on configuring the VM-Series firewall as an HA pair, see Configure Active/Passive HA and Configure Active/Active HA.

# High Availability for VM-Series in AWS

The VM-Series firewall is a core part of your network security solution in the Amazon Web Services Virtual Private Cloud (AWS-VPC). To ensure redundancy, you can deploy the VM-Series firewalls in AWS in an active/passive High Availability (HA) configuration. The active peer continuously synchronizes its configuration and session information with the identically configured passive peer. A heartbeat connection between the two devices ensures seamless failover if the active device goes down. When the passive peer detects this failure it becomes active and triggers API calls to the AWS infrastructure to move all the dataplane interfaces (ENIs) from the failed peer to itself. The failover time can vary from 20 seconds to over a minute depending on the responsiveness from the AWS infrastructure.

Because the dataplane interfaces are moved from the active firewall on failover, the VM-Series firewall in AWS must have permissions to initiate API actions for detaching and attaching network interfaces from the active peer to the passive peer. Therefore, you must create a role in the AWS Identity and Access Management (IAM) service and assign it to a user or group. The role must have permissions for the following operations (at a minimum):

- AttachNetworkInterface—For permission to attach an ENI to an instance.
- DescribeNetworkInterface—For fetching the ENI parameters in order to attach an interface to the instance.
- DetachNetworkInterface—For permission to detach the ENI from the EC2 instance.
- DescribeInstances—For permission to obtain information on the EC2 instances in the VPC.

For detailed steps in the AWS console, refer to the AWS documentation. The following screenshot shows the access management settings for the IAM role described above:


At the time you deploy the VM-Series firewalls on an Elastic Compute Cloud (EC2) instance, you must attach this IAM role to each peer in the HA pair.

The devices in an HA pair use HA links to synchronize data and maintain state information. In AWS, the VM-Series firewall uses the management port as the control link (HA1 link) for exchanging hellos, heartbeats, and HA state, and User-ID information. This link is also used to synchronize configuration changes on either

 © Palo Alto Networks, Inc.

the active or passive device with its peer. Ethernet1/1 must be assigned as the data link (HA2 link). The data link is used to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between devices in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active device to the passive device.

The VM-Series on AWS does not support backup links for HA1 or HA2.

# Support for Jumbo Frames

A jumbo frame is an Ethernet frame with a Maximum Transmission Unit (MTU) value greater than 1500 bytes. With the exception of the VM-Series NSX edition firewall, all the other deployments of the VM-Series firewall can be configured to handle jumbo frames.

When jumbo frames are enabled globally, the default MTU size for all Layer 3 interfaces is set to 9192 bytes. Based on your other network infrastructure, you can modify/ override the MTU value for a specific interface to a value that ranges between 512 to 9216 bytes.

To enable jumbo frames, see Increase Jumbo Frame Size.

 © Palo Alto Networks, Inc.

# Support for Hypervisor Assigned MAC Addresses

The VM-Series firewall supports the ability to detect the MAC address assigned to the physical interface by the host/hypervisor and use that MAC address on the VM-Series firewall. This capability allows non-learning switches, such as the VMware vSwitch to forward traffic to the dataplane interface on the firewall without requiring that promiscuous mode be enabled on the vSwitch.

Until PAN-OS 7.0, to configure a vSwitch to forward frames to the VM-Series firewall, you had to enable promiscuous mode on the switch. This was required because the VM-Series firewall assigned a unique MAC address for each dataplane interface from its own pool, and as some vSwitches could not learn the custom MAC address, the host would drop the frame when the destination MAC address for an interface was not the same as the host-assigned MAC address. Enabling promiscuous mode on the vSwitch was the only way to forward traffic to the firewall.

Now, to allow the VM-Series firewall to use the interface MAC addresses provided by the host/hypervisor, select **Use Hypervisor Assigned MAC Address** in **Device > Management > Setup**. When the MAC address change occurs, the firewall generates a system log to record this transition and the interface generates a gratuitous ARP. VM-Series firewalls with hypervisor assigned MAC address in a high-availability configuration behave differently than the hardware appliances with respect to MAC addressing. Hardware firewalls use self-generated floating MAC addresses between devices in an HA pair. In contrast, VM-Series firewalls in an HA configuration, will use the hypervisor-assigned MAC address (if so configured). When a failover is triggered, the now active VM-Series firewall will send a gratuitous ARP so that neighboring devices can learn the updated MAC/IP address pairing.

 © Palo Alto Networks, Inc.

# Networking Features

- ECMP
- DHCP Options
- Granular Actions for Blocking Traffic in Security Policy
- Session-Based DSCP Classification
- Per-Virtual System Service Routes
- LLDP
- Network Prefix Translation (NPTv6)
- TCP Split Handshake Drop

# ECMP

Equal Cost Multi Path (ECMP) processing is a networking feature that enables the firewall to use up to four equal-cost routes to the same destination. Without this feature, if there are multiple equal-cost routes to the same destination, the virtual router chooses one of those routes from the routing table and adds it to its forwarding table; it will not use any of the other routes unless there is an outage in the chosen route.

Enabling ECMP on a virtual router and selecting one of the ECMP Load-Balancing Algorithms allows the firewall have up to four equal-cost paths to a destination in its forwarding table, allowing the firewall to:

- Load balance flows (sessions) to the same destination over multiple equal-cost links.
- Make use of the available bandwidth on links to the same destination rather than leave some links unused.
- Dynamically shift traffic to another ECMP member to the same destination if a link fails, rather than having to wait for the routing protocol or RIB table to elect an alternative path/route. This can help reduce down time when links fail.

▲ ECMP Platform, Interface, and IP Routing Support

▲ Configure ECMP on a Virtual Router

## ECMP Platform, Interface, and IP Routing Support

ECMP is supported on all Palo Alto Networks firewall platforms, with hardware forwarding support on the PA-7000 Series, PA-5000 Series, PA-3060 firewalls, and PA-3050 firewalls. PA-3020 firewalls, PA-500 firewalls, PA-200 firewalls, and VM-Series firewalls support ECMP through software only. Performance is affected for sessions that cannot be hardware offloaded.

ECMP is supported on Layer 3, Layer 3 subinterface, VLAN, tunnel, and Aggregated Ethernet interfaces.

ECMP can be configured for static routes and any of the dynamic routing protocols the firewall supports.

ECMP affects the route table capacity because the capacity is based on the number of paths, so an ECMP route with four paths will consume four entries of route table capacity. ECMP implementation might slightly decrease the route table capacity because more memory is being used by session-based tags to map traffic flows to particular interfaces.

If you are using HA, consider the HA Active/Active Failover Behavior with ECMP.

## Configure ECMP on a Virtual Router

Use the following procedure to enable ECMP on a virtual router. This task assumes you have specified the interfaces that belong to a virtual router (**Network > Virtual Routers > Router Settings > General**) and the IP routing protocol.

Enabling, disabling or changing ECMP for an existing virtual router causes the system to restart the virtual router, which might cause sessions to be terminated.

 © Palo Alto Networks, Inc.

| Configure ECMP on a Virtual Router | | |
|---|---|---|
| Step 1 | Enable ECMP for a virtual router. | 1. Select **Network > Virtual Routers** and select the virtual router on which to enable ECMP.<br>2. Select **Router Settings > ECMP** and select the **Enable** check box. |
| Step 2 | (Optional) Enable symmetric return of packets from server to client. | Optionally select the **Symmetric Return** check box to cause return packets to egress out the same interface on which the associated ingress packets arrived. That is, the firewall will use the ingress interface on which to send return packets, rather than use the ECMP interface, so the **Symmetric Return** setting overrides load balancing. This behavior occurs only for traffic flows from the server to the client. |
| Step 3 | Specify the maximum number of equal-cost paths (to a destination network) that can be copied from the Routing Information Base (RIB) to the Forwarding Information Base (FIB). | For **Max Path** allowed, enter 2, 3, or 4. Default: 2. |
| Step 4 | Select one of the ECMP Load-Balancing Algorithms for the virtual router. | For **Load Balance**, select one of the following options from the **Method** drop-down:<br>• **IP Modulo**—By default, the virtual router load balances sessions using this option, which uses a hash of the source and destination IP addresses in the packet header to determine which ECMP route to use.<br>• **IP Hash**—Uses a hash of the source and destination IP addresses and optionally the source and destination port numbers in the packet header to determine which ECMP route to use. Specify options in Step 5 below.<br>• **Balanced Round Robin**—Uses round robin among the ECMP paths and re-balances paths when the number of paths changes.<br>• **Weighted Round Robin**—Uses round robin and a relative weight to select from among ECMP paths. Specify the weights in Step 6 below. |
| Step 5 | (Optional) (**IP Hash** only) Configure IP Hash options. | 1. Select the **Use Source/Destination Ports** check box if you want to use source or destination port numbers in the **IP Hash** calculation.<br>2. Enter a **Hash Seed** value (an integer with a maximum of 9 digits). Specify a **Hash Seed** value to further randomize load balancing. Specifying a hash seed value is useful if you have a large number of sessions with the same tuple information. |

| Configure ECMP on a Virtual Router (Continued) | |
|---|---|
| Step 6 (**Weighted Round Robin** only) Define a weight for each interface in the ECMP group. | If you selected **Weighted Round Robin** as the **Method**, define a weight for each of the interfaces that are the egress points for traffic to be routed to the same destinations (that is, interfaces that are part of an ECMP group, such as the interfaces that provide redundant links to your ISP or interfaces to the core business applications on your corporate network).<br>The higher the weight, the more often that equal-cost path will be selected for a new session.<br>A higher speed link should be given a higher weight than a slower link so that more of the ECMP traffic goes over the faster link.<br>1. Create an ECMP group by clicking **Add** and selecting an **Interface** from the drop-down.<br>2. Repeat to add all the interfaces in the ECMP group.<br>3. Click on **Weight** and specify the relative weight for each interface (range is 1-255; default is 100). |
| Step 7 Save the configuration. | 1. Click **OK**.<br>2. Click **Yes** to restart the virtual router. Restarting the virtual router might cause sessions to be terminated.<br> 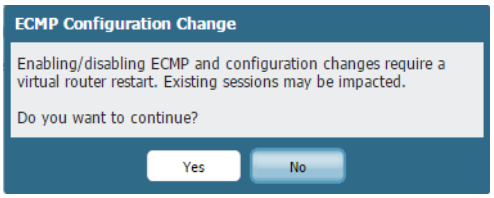 <br>This message displays only if you are modifying an existing virtual router with ECMP. |
| Step 8 Save the configuration. | **Commit** the configuration. |
| Step 9 (Optional) For BGP routing, Enable ECMP for Multiple BGP Autonomous Systems. | — |
| Step 10 Verify that some routes are equal-cost multiple paths.<br>A virtual router configured for ECMP indicates in the Forwarding Information Base (FIB) which routes are ECMP routes. An ECMP flag (E) for a route indicates that it is participating in ECMP for the egress interface to the next hop for that route. | 1. Select **Network > Virtual Routers**.<br>2. In the row of the virtual router for which you enabled ECMP, click **More Runtime Stats**.<br>3. Select **Routing > Forwarding Table** to see the FIB. In the table, multiple routes to the same Destination (out a different Interface) have the E flag.<br>An asterisk [*] denotes the preferred path for the ECMP group. |

 © Palo Alto Networks, Inc.

# DHCP Options

Prior to PAN-OS 7.0, Palo Alto Networks firewalls supported predefined DHCP options in the DHCP server implementation. Beginning with PAN-OS 7.0, the firewalls also support user-defined DHCP options in the DHCP server implementation. You can configure DHCP options that are defined in RFC 2132, DHCP Options and BOOTP Vendor Extensions, and your own vendor-specific and customized options.

The DHCP options support a wide variety of office equipment, such as IP phones and wireless infrastructure devices. Each option code supports multiple values, which can be IP address, ASCII, or hexadecimal format. With the firewall enhanced DCHP option support, branch offices do not need to purchase and manage their own DHCP servers in order to provide vendor-specific and customized options to DHCP clients, such as DHCP Options 43, 55, and 60 and other Customized Options. You can enter Multiple Values for a DHCP Option.

Prior to configuring your DHCP server options, complete the following tasks:

- Select an interface to be a DHCP Server in Configure an Interface as a DHCP Server.
- Collect the DHCP options, values, and Vendor Class Identifiers you plan to configure.

Perform the following task to configure the DHCP options that the server provides to clients.

| Configure DHCP Server Options | | |
|---|---|---|
| Step 1 | Access the DHCP options. | Select **Network > DHCP > DHCP Server**, click **Add** and select the **Options** tab. |
| Step 2 | Configure the predefined DHCP options that the server sends to its clients. | See Step 2 in Configure an Interface as a DHCP Server. |
| Step 3 | (Optional) Configure a vendor-specific or custom DHCP option that the DHCP server sends to its clients. | See Step 3 in Configure an Interface as a DHCP Server. |
| Step 4 | (Optional) Add another vendor-specific or custom DHCP option. | See Step 4 in Configure an Interface as a DHCP Server. |
| Step 5 | Save the configuration. | Click **OK** and **Commit** the changes. |

# Granular Actions for Blocking Traffic in Security Policy

Security policy allows you to either allow or block traffic. When you configure the firewall to block traffic, it either resets the connection or silently drops packets. Because silently dropping packets causes some applications to break and appear unresponsive to the user, you now have new actions to gracefully block traffic and provide a better user experience. These actions are written to the traffic logs and can be used for log queries.

For traffic that matches the criteria defined in a security policy, you can apply the following actions:

- **Allow**—(default action) Allows the traffic.
- **Deny**—Blocks traffic, and enforces the default *Deny Action* defined for the application that is being denied. To view the deny action defined by default for an application, view the application details in **Objects > Applications** or check the application details in Applipedia.
- **Drop**—Silently drops the traffic; for an application, it overrides the default deny action. A TCP reset is not sent to the host/application.

  For Layer 3 interfaces, to optionally send an ICMP unreachable response to the client, set Action: **Drop** and enable the **Send ICMP Unreachable** checkbox. When enabled, the firewall sends the ICMP code for *communication with the destination is administratively prohibited*— ICMPv4: Type 3, Code 13; ICMPv6: Type 1, Code 1.
- **Reset client**—Sends a TCP reset to the client-side device.
- **Reset server**—Sends a TCP reset to the server-side device.
- **Reset both**—Sends a TCP reset to both the client-side and server-side devices.

A reset is sent only after a session is formed. If the session is blocked before a 3-way handshake is completed, the firewall will not send the reset.

For a TCP session with a reset action, the firewall does not send an ICMP Unreachable response.

For a UDP session with a drop or reset action, if the **ICMP Unreachable** checkbox is selected, the firewall sends an ICMP message to the client.

Because you can now enable the firewall to send ICMP unreachable responses for a better user experience, the following configuration settings are also available:

- **Send ICMP Unreachable**—When you configure security policy to drop traffic or to reset the connection, the traffic does not reach the destination host. On Layer 3 interfaces, for UDP traffic that is dropped or reset and for TCP traffic that is dropped, you can enable the firewall to send an ICMP Unreachable response to the source IP address from where the traffic originated. Enabling this setting allows the source to gracefully close/clear the session and prevents applications from breaking.
- **ICMP Unreachable Packet Rate (per sec)**—Sets the maximum number of ICMP unreachable messages that the firewall can send in a second (range is 1-65535; default is 200). This value applies for ICMP unreachable messages sent to IPv4 and IPv6 addresses, and across all sessions and source IP addresses.

 © Palo Alto Networks, Inc.

| Create a Security Policy to Drop Traffic and Send an ICMP Unreachable Response | | |
|---|---|---|
| Step 1 | Create a security policy to drop an application and send an ICMP message to the client. | 1. Select **Policies > Security** and click **Add**.<br>2. Give the rule a descriptive name in the General tab.<br>3. In the Source tab, set the Source Zone; in the Destination tab, Set the Destination Zone.<br>4. In the Application tab, **Add** the application that you want to drop.<br>5. In the Actions tab, complete these tasks:<br>a. Set the **Action** to Drop and select the checkbox for **Send ICMP Unreachable**.<br> 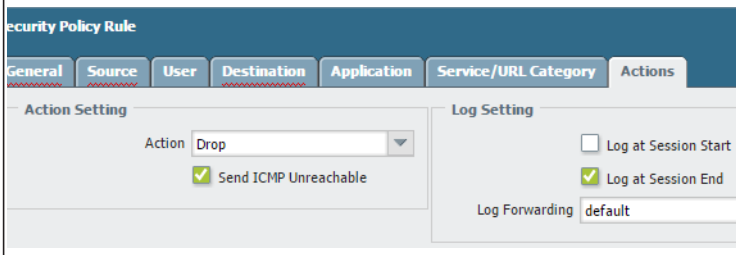 <br>b. Verify that logging is enabled, **Log at Session End** in the Log Setting section. When a policy match occurs, the traffic log will record the drop action for the application.<br>6. Click **OK** and **Commit** the changes. |
| Step 2 | (Optional) View or change the ICMP Unreachable packet rate defined on the firewall. | 1. On the firewall web interface, select **Device > Setup > Session**.<br>2. Edit the Session Settings section, and view the value for the maximum number of ICMP packets that the firewall can send in a second in **ICMP Unreachable Packet Rate (per sec)**.<br>3. If you edited the settings, click **OK** and **Commit** the changes. |

# Session-Based DSCP Classification

A Differentiated Services Code Point (DSCP) is a packet header value that can be used to indicate the level of service requested for traffic, such as high priority or best effort delivery. Session-based DSCP classification allows you to both honor the service class requested for traffic and to mark a session to receive continued QoS treatment. Session-based DSCP extends the power of Quality of Service (QoS), which polices traffic as it passes through the firewall, by allowing all network devices between the firewall and the client to also police traffic. All inbound and outbound traffic for a session can receive continuous QoS/DSCP treatment as it flows through your network. For example, inbound return traffic from an external server can now be treated with the same QoS/DSCP priority that the firewall initially enforced for the outbound flow. Network devices intermediate to the firewall and end user will also then enforce the same priority for the return traffic (and any other outbound or inbound traffic for the session).

At a glance, the following new options are provided to support session-based DSCP classification:

- Use DSCP/ToS values as matching criteria when configuring a QoS policy rule (**Policies > QoS > DSCP/ToS**).

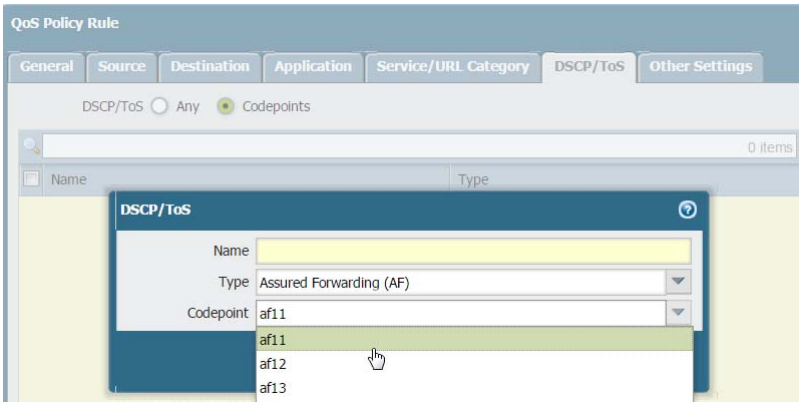


- The new **Follow Client to Server Flow** setting enables the firewall to detect the DSCP/ToS value for inbound traffic matched to a security rule; this allows the firewall to then apply the same DSCP/ToS value to the return traffic matching the security rule (**Policies > Security > QoS Marking**).You then have the option to configure a QoS policy rule to police the return traffic an deprioritize it in the same way as the outbound traffic.

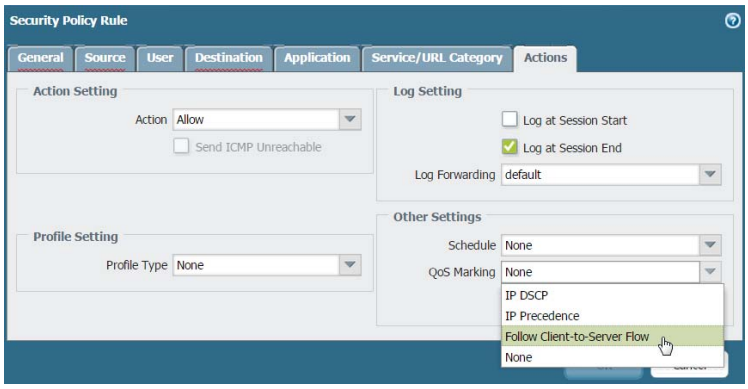


Use the following steps to enable session-based DSCP classification. Start by configuring QoS based on DSCP marking detected at the beginning of a session. You can then continue to enable the firewall to mark the return flow for a session with the same DSCP value used to enforce QoS for the initial outbound flow.

 © Palo Alto Networks, Inc.

## Provide QoS Based on DSCP/ToS Marking

Before you set up or modify a QoS policy to control traffic based on DSCP/ToS code point values, make sure that you have performed the preliminary steps for configuring QoS:

- ❑ Identify the traffic to which to apply QoS.
- ❑ Identify the egress interface for traffic to be enforced with QoS.

**Step 1** Apply QoS to traffic based on the DSCP value detected at the beginning of the session.

1. Select **Policies > QoS** and **Add** or modify an existing QoS rule and populate required fields.
2. (NEW) **Add** a **DSCP/ToS** rule and give the rule a descriptive **Name**. You can choose to add multiple DSCP/ToS rules to a single QoS rule to enforce the same QoS priority for sessions with different DSCP values.
3. (NEW) Select the **Type** of DSCP/ToS marking for the QoS rule to match to traffic:

   

   It is a best practice to use a single DSCP type to manage and prioritize your network traffic.

   - **Expedited Forwarding (EF)**: Use to request low loss, low latency and guaranteed bandwidth for traffic. Packets with EF codepoint values are typically guaranteed highest priority delivery.
   - **Assured Forwarding (AF)**: Use to provide reliable delivery for applications. Packets with AF codepoint indicate a request for the traffic to receive higher priority treatment than best effort service provides (though packets with an EF codepoint will continue to take precedence over those with an AF codepoint).
   - **Class Selector (CS)**: Use to provide backward compatibility with network devices that use the IP precedence field to mark priority traffic.
   - **IP Precedence (ToS)**: Use with legacy network devices to mark priority traffic (the IP Precedence header field was used to indicate the priority for a packet before the introduction of the DSCP classification).
   - **Custom Codepoint**: Create a custom codepoint to match to traffic by entering a **Codepoint Name** and **Binary Value**.

   For example, select **Assured Forwarding (AF)** to ensure traffic marked with an **AF** codepoint value has higher priority for reliable delivery over applications marked to receive lower priority.
4. (NEW) Match the QoS policy to traffic on a more granular scale by specifying the **Codepoint** value. For example, with Assured Forwarding (AF) selected as the **Type** of DSCP value for the policy to match, further specify an AF **Codepoint** value such as AF11.

   

   When Expedited Forwarding (EF) is selected as the **Type** of DSCP marking, a granular **Codepoint** value cannot be specified. The QoS policy will match to traffic marked with any EF codepoint value.
5. Select **Other Settings** and a **QoS Class** to assign to traffic matched to the QoS rule. For example, assign Class 1 to sessions where a DSCP marking of AF11 is detected for the first packet in the session.
6. Click **OK** to save the QoS rule.

| | | |
|---|---|---|
| **Step 2** | Define the QoS priority for traffic to receive when it is matched to a QoS rule based the DSCP marking detected at the beginning of a session. | 1. Select **Network > Network Profiles > QoS Profile** and **Add** or modify an existing QoS profile. For details on profile options to set priority and bandwidth for traffic, see QoS Concepts and Configure QoS.<br>2. **Add** or modify a profile class. For example, because Step 1 showed steps to classify AF11 traffic as Class 1 traffic, you could add or modify a **class1** entry.<br>3. Select a **Priority** for the class of traffic, such as **high**.<br>4. Click **OK** to save the QoS Profile. |

| Provide QoS Based on DSCP/ToS Marking (Continued) | | |
|---|---|---|
| Step 3 | Enable QoS on an interface. | Select **Network > QoS** and **Add** or modify an existing interface and **Turn on QoS feature on this interface**.<br>In this example, traffic with an AF11 DSCP marking is matched to the QoS rule and assigned Class 1. The QoS profile enabled on the interface enforces high priority treatment for Class 1 traffic as it egresses the firewall (the session *outbound* traffic). |
| Step 4 | (NEW) Enable DSCP Marking.<br>Mark return traffic with a DSCP value, enabling the inbound flow for a session to be marked with the same DSCP value detected for the outbound flow. | 1. Select **Policies > Security** and **Add** or modify a security policy.<br>2. Select **Actions** and in the **QoS Marking** drop-down, choose **Follow-Client-to-Server-Flow**.<br>3. Click **OK** to save your changes.<br>Completing this step enables the firewall to mark traffic with the same DSCP value that was detected at the beginning of a session (in this example, the firewall would mark return traffic with the DSCP AF11 value). While configuring QoS allows you to shape traffic as it egresses the firewall, enabling this option in a security rule allows the other network devices intermediate to the firewall and the client to continue to enforce priority for DSCP marked traffic. |
| Step 5 | Save the configuration. | **Commit** your changes. |

 © Palo Alto Networks, Inc.

# Per-Virtual System Service Routes

The firewall uses the MGT interface (by default) to access external services, such as DNS servers, software updates, and software licenses. An alternative to using the MGT interface is to configure a data port (a regular interface) to access these services. The path from the interface to the service on a server is known as a *service route*. Each service allows redirection of management services to the respective virtual system owner through one of the interfaces associated with that virtual system.

Starting with PAN-OS 7.0, the source interface and source IP address for service routes can be configured for individual virtual systems, in addition to the global service route. This feature provides the flexibility to customize service routes for numerous tenants or departments on a single firewall. The service packets exit the firewall on a port that is assigned to a specific virtual system, and the server sends its response to the configured source interface and source IP address. Any virtual system that does not have a service route configured for a particular service inherits the interface and IP address that are set globally for that service.

Prior to PAN-OS 7.0, each service route to a service was configured globally and applied to the entire firewall.

A DNS Proxy object is where you configure the settings that determine how the firewall functions as a DNS Proxy. The introduction of the Per-Virtual System Service Route feature enables you to assign a DNS Proxy object to a single virtual system. In this case you can specify a DNS Server Profile, which specifies the primary and secondary DNS server addresses, along with other information.

Consider the Use Cases for Service Routes for a Virtual System and Multi-Tenant DNS Deployments.

## Customize Service Routes for a Virtual System

- Prior to performing this task, in order to see the Global and Virtual Systems tabs, **Multi Virtual System Capability** must be enabled.

If **Multi Virtual System Capability** is enabled, any virtual system that does not have specific service routes configured inherits the device's global service and service route settings.

The following example shows how to configure service routes for a virtual system.

| Customize Service Routes to Services Per Virtual System | |
|---|---|
| Step 1 Customize service routes for a virtual system. | 1. Select **Device > Setup > Services > Virtual Systems**, and select the virtual system you want to configure.<br>2. Click the **Service Route Configuration** link.<br>3. Select one of the radio buttons:<br>• **Inherit Global Service Route Configuration**—Causes the virtual system to inherit the global service route settings relevant to a virtual system. If you choose this option, skip down to Step 7.<br>• **Customize**—Allows you to specify a source interface and source address for each service.<br>4. If you chose **Customize**, select the **IPv4** or **IPv6** tab, depending on what type of addressing the server offering the service uses. You can specify both IPv4 and IPv6 addresses for a service. Click the check box(es) for the services for which you want to specify the same source information. (Only services that are relevant to a virtual system are available.) Click **Set Selected Service Routes**.<br>• For **Source Interface**, select **Any**, **Inherit Global Setting**, or an interface from the drop-down to specify the source interface that will be used in packets sent to the external service(s). Hence, the server's response will be sent to that source interface. In our example deployment, you would set the source interface to be the tenant's subinterface.<br>• **Source Address** will indicate **Inherited** if you selected **Inherit Global Setting** for the **Source Interface**. Or it will indicate the source address of the **Source Interface** you selected. If you selected **Any** for **Source Interface**, select an IP address from the drop-down, or enter an IP address (using the IPv4 or IPv6 format that matches the tab you chose) to specify the source address that will be used in packets sent to the external service.<br>• If you modify an address object and the IP family type (IPv4/IPv6) changes, a **Commit** is required to update the service route family to use.<br>5. Click **OK**.<br>6. Repeat Step 4 and Step 5 to configure source addresses for other external services.<br>7. Click **OK**. |
| Step 2 Save the configuration. | Click **Commit** and **OK**. |
| **Next Steps...** | • If you are configuring per-virtual system service routes for logging services for a PA-7000 Series firewall, perform the task to Configure a PA-7000 Series Firewall for Logging Per Virtual System.<br>• Configure a DNS Proxy Object<br>• Configure a DNS Server Profile<br>• Configure Administrative Access Per Virtual System or Device |

 © Palo Alto Networks, Inc.

# LLDP

Palo Alto Networks firewalls now support Link Layer Discovery Protocol (LLDP), which functions at the link layer to discover neighboring devices and their capabilities. LLDP allows the firewall and other network devices to send and receive LLDP data units (LLDPDUs) to and from neighbors. The receiving device stores the information in a MIB, which the Simple Network Management Protocol (SNMP) can access. LLDP makes troubleshooting easier, especially for virtual wire deployments where the firewall would typically go undetected by a ping or traceroute.

Due to differences in how Aggregate Ethernet interfaces function, the PA-2000 Series platform is not supported. Panorama, the GlobalProtect Mobile Security Manager, and the WildFire appliance are also not supported.

Interface types that do not support LLDP are TAP, High Availability (HA), Decrypt Mirror, virtual wire/vlan/L3 subinterfaces, and PA-7000 Series firewall Log Processing Card (LPC) interfaces.

An LLDP Ethernet frame and the TLV structure within the frame are illustrated in the LLDP Overview. The Supported TLVs in LLDP include mandatory and optional TLVs.

## Configure LLDP

To configure LLDP, and create an LLDP profile, you must be a superuser or device administrator (deviceadmin). A firewall interface supports a maximum of five LLDP peers.

| Configure LLDP | | |
|---|---|---|
| Step 1 | Enable LLDP on the firewall. | Select **Network > LLDP** and edit the LLDP General section; select the **Enable** check box. |
| Step 2 | (Optional) Change LLDP global settings. | 1. For **Transmit Interval (sec)**, specify the interval (in seconds) at which LLDPDUs are transmitted (range is 1-3600; default is 30).<br>2. For **Transmit Delay (sec)**, specify the delay time (in seconds) between LLDP transmissions sent after a change is made in a TLV element (range is 1-600; default is 2).<br>The delay helps to prevent flooding the segment with LLDPDUs if many network changes spike the number of LLDP changes, or if the interface flaps. The **Transmit Delay** must be less than the **Transmit Interval**.<br>3. For **Hold Time Multiple**, specify a value that is multiplied by the **Transmit Interval** to determine the total TTL Hold Time (range is 1-100; default is 4). The maximum TTL Hold Time is 65535 seconds, regardless of the multiplier value.<br>4. For **Notification Interval**, specify the interval (in seconds) at which LLDP Syslog Messages and SNMP Traps are transmitted when MIB changes occur (range is 1-3600; default is 5).<br>5. Click **OK**. |

| Configure LLDP (Continued) | |
|---|---|
| Step 3 Create an LLDP profile.<br>For descriptions of the optional TLVs, see Supported TLVs in LLDP. | 1. Select **Network > Network Profiles > LLDP Profile** and click **Add**.<br>2. Enter a **Name** for the LLDP profile.<br>3. For **Mode**, select **transmit-receive** (the default), **transmit-only**, or **receive-only**.<br>4. Click the **SNMP Syslog Notification** check box to enable SNMP notifications and syslog messages. If enabled, the global **Notification Interval** is used. The firewall will send both an SNMP trap and a Syslog event as configured in the **Device > Log Settings > System > SNMP Trap Profile** and **Syslog Profile**.<br>5. For **Optional TLV**s, select the TLVs you want transmitted:<br>• **Port Description**<br>• **System Name**<br>• **System Description**<br>• **System Capabilities**<br>6. Specifying a **Management Address** is optional. To add one or more, select the check box and **Add** a **Name**.<br>7. Select the **Interface** from which to obtain the management address. At least one management address is required if **Management Address** TLV is enabled. If no management IP address is configured, the system uses the MAC address of the transmitting interface as the management address TLV.<br>8. Select **IPv4** or **IPv6**, and in the adjacent field, select an IP address from the drop-down (which lists the addresses configured on the selected interface), or enter an address.<br>9. Click **OK**.<br>10. Up to four management addresses are allowed. If you specify more than one **Management Address**, they will be sent in the order they are specified, starting at the top of the list. To change the order of the addresses, select an address and use the **Move Up** or **Move Down** buttons.<br>11. Click **OK**. |
| Step 4 Assign an LLDP profile to an interface. | 1. Select **Network > Interfaces** and select the interface where you will assign an LLDP profile.<br>2. Select **Advanced > LLDP**.<br>3. Select the **Enable LLDP** check box to assign an LLDP profile to the interface.<br>4. For **Profile**, select the profile you created. Selecting **None** enables LLDP with basic functionality, such as sending the three mandatory TLVs and enabling **transmit-receive** mode.<br>If you want to create a new profile, click **LLDP Profile** and follow the instructions. in the prior step.<br>5. Click **OK**. |
| Step 5 Save the configuration. | Click **Commit**. |

 © Palo Alto Networks, Inc.

| Configure LLDP (Continued) | |
|---|---|
| **Next Steps...** | • View LLDP Settings and Status<br>• Clear LLDP Statistics |

# Network Prefix Translation (NPTv6)

IPv6-to-IPv6 Network Prefix Translation (NPTv6) is now supported. NPTv6 performs a stateless, static translation of one IPv6 prefix to another IPv6 prefix (port numbers are not changed). One benefit of NPTv6 is the prevention of asymmetric routing problems that result from Provider Independent addresses being advertised from multiple data centers. NPTv6 allows more specific routes to be advertised so that return traffic arrives at the same firewall that transmitted the traffic. Another benefit is the independence of private and public addresses; you can change one without affecting the other. A third benefit of NPTv6 is the ability to translate Unique Local Addresses to globally routable addresses.

NPTv6 is supported on the following platforms (NPTv6 with hardware lookup but packets go through the CPU): PA-7000 Series, PA-5000 Series, PA-4000 Series, PA-3060, PA-3050, and PA-2000 Series firewalls. Platforms supported with no ability to have hardware perform a session look-up: PA-3020 firewall, PA 500 firewall, PA-200 firewall, and VM-Series.

The NPTv6 Overview explains that NPTv6 does not provide security. NPTv6 will translate Unique Local Addresses (ULAs) to globally routable addresses, which you might want to do if your networks use ULAs and you need to communicate beyond the limited area for which they were intended. Following the overview is an explanation of How NPTv6 Works and why it is necessary to configure NDP Proxy. A firewall acting as NDP proxy can send Neighbor Discovery (ND) advertisements and respond to ND solicitations from peers that are asking for MAC addresses of IPv6 prefixes assigned to devices behind the firewall. The NPTv6 and NDP Proxy Example illustrates how NPTv6 and NDP proxy function together.

## Create an NPTv6 Policy

Perform this task when you want to configure an NPTv6 policy to translate one IPv6 prefix to another IPv6 prefix. The prerequisites for this task are:

- Enable IPv6. Select **Device > Setup > Session**. Click **Edit** and select **IPv6 Firewalling**.
- Configure a Layer 3 Ethernet interface with a valid IPv6 address, and with IPv6 enabled. Select **Network > Interfaces > Ethernet**, select an interface, and on the **IPv6** tab, select **Enable IPv6 on the interface**.
- Create network security policies, because NPTv6 does not provide security.
- Decide whether you want source translation, destination translation, or both.
- Know the zones to which you want to apply the NPTv6 policy.
- Know your original and translated IPv6 prefixes.

| Configure an NPTv6 Policy | |
|---|---|
| Step 1 Create a new NPTv6 policy. | 1. Select **Policies > NAT** and click **Add**.<br>2. On the **General** tab, enter a descriptive **Name** for the NPTv6 policy rule.<br>3. (Optional) Enter a **Description** and **Tag**.<br>4. For **NAT Type**, select **NPTv6**. |

 © Palo Alto Networks, Inc.

| Configure an NPTv6 Policy (Continued) | | |
|---|---|---|
| Step 2 | Specify the match criteria for incoming packets; packets that match all of the criteria are subject to the NPTv6 translation.<br>Zones are required for both types of translation. | 1. On the **Original Packet** tab, for **Source Zone**, leave **Any** or click **Add** to enter the source zone to which the policy applies.<br>2. Enter the **Destination Zone** to which the policy applies.<br>3. (Optional) Select a **Destination Interface**.<br>4. (Optional) Select a **Service** to restrict what type of packets are translated.<br>5. If you are doing source translation, enter a **Source Address** or select **Any**. The address could be an address object. The following constraints apply to **Source Address** and **Destination Address**:<br>• Prefixes of **Source Address** and **Destination Address** for the **Original Packet** and **Translated Packet** must be in the format xxxx:xxxx::/yy, although leading zeros in the prefix can be dropped.<br>• The IPv6 address cannot have an interface identifier (host) portion defined.<br>• The range of supported prefix lengths is /32 to /64.<br>• The **Source Address** and **Destination Address** cannot both be set to **Any**.<br>6. If you are doing source translation, you can optionally enter a **Destination Address**. If you are doing destination translation, the **Destination Address** is required. See the constraints listed in the prior step. |
| Step 3 | Specify the translated packet. | 1. On the **Translated Packet** tab, if you want to do source translation, in the Source Address Translation section, for **Translation Type**, select **Static IP**. If you do not want to do source translation, select **None**.<br>2. If you chose **Static IP**, the **Translated Address** field appears. Enter the translated IPv6 prefix or address object. See the constraints listed in Step 2.<br>t is a best practice to configure your **Translated Address** to be the prefix of your firewall's untrust interface address. For example, if your untrust interface has the address 2001:1a:1b:1::99/64, make your **Translated Address** 2001:1a:1b:1::0/64.<br>3. (Optional) Select **Bi-directional** if you want the firewall to create a corresponding NPTv6 translation in the opposite direction of the translation you configure.<br>If you enable **Bi-directional** translation, it is very important to make sure you have security policies in place to control the traffic in both directions. Without such policies, the **Bi-directional** feature will allow packets to be automatically translated in both directions, which you might not want.<br>4. If you want to do destination translation, select **Destination Address Translation**. In the **Translated Address** field, choose an address object from the drop-down or enter your internal destination address.<br>5. Click **OK**. |

| Configure an NPTv6 Policy (Continued) | | |
|---|---|---|
| Step 4 | Configure NDP Proxy.<br>When you configure the firewall to act as an NDP Proxy for addresses, it allows the firewall to send Neighbor Discovery (ND) advertisements and respond to ND solicitations from peers that are asking for MAC addresses of IPv6 prefixes assigned to devices behind the firewall. | 1. Select **Network > Interfaces > Ethernet** and select an interface.<br>2. On the **Advanced > NDP Proxy** tab, select **Enable NDP Proxy** and click **Add**.<br>3. Enter the **IP Address(es)** for which NDP Proxy is enabled. It can be an address, a range of addresses, or a prefix and prefix length. The order of IP addresses does not matter. These addresses are ideally the same as the Translated Addresses that you configured in an NPTv6 policy.<br>  If the address is a subnet, the NDP Proxy will respond to all addresses in the subnet, so you should list the neighbors in that subnet with **Negate** selected, as described in the next step.<br>4. (Optional) Enter one or more addresses for which you do *not* want NDP Proxy enabled, and select **Negate**. For example, from an IP address range or prefix range configured in the prior step, you could negate a smaller subset of addresses. It is recommended that you negate the addresses of the firewall's neighbors. |
| Step 5 | Save the configuration. | Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

# TCP Split Handshake Drop

You can configure a TCP split handshake drop in a Zone Protection profile to prevent a TCP session from being established if the session establishment procedure does not use the well-known three-way handshake, but instead uses a variation, such as a four-way or five-way split handshake or a simultaneous open.

Perform the following optional task to configure a Zone Protection profile that prevents TCP sessions from being established unless they use the standard three-way handshake. This task assumes that the interface where this feature is desired is already assigned to a security zone.

| Configure a Zone Protection Profile to Prevent TCP Split Handshake Sessions | | |
|---|---|---|
| Step 1 | Configure a zone protection profile to prevent TCP sessions that use anything other than a three-way handshake to establish a session. | 1. Select **Network > Network Profiles > Zone Protection** and click **Add** to create a new profile, or select an existing profile.<br>2. If creating a new profile, enter a **Name** for the profile and an optional **Description**.<br>3. Select **Packet Based Attack Protection > TCP Drop** and select the **Split Handshake** check box.<br>4. Click **OK**. |
| Step 2 | Apply the profile to one or more security zones. | 1. Select **Network > Zones** and select the zone where you want to assign the Zone Protection profile.<br>2. Select the profile you configured in Step 1 from the **Zone Protection Profile** drop-down.<br>Alternatively, you could start creating a new profile here by clicking **Zone Protection Profile**, in which case you would continue accordingly.<br>3. Click **OK**.<br>4. (Optional) Repeat steps 1-3 to apply the profile to additional zones. |
| Step 3 | Save the configuration. | Click **OK** and **Commit**. |

 © Palo Alto Networks, Inc.

# VPN Features

- IKEv2 Support for VPN Tunnels
- IPSec VPN Enhancements

# IKEv2 Support for VPN Tunnels

An IPSec VPN gateway uses IKEv1 or IKEv2 to negotiate the IKE security association (SA) and IPSec tunnel. IPSec VPN tunnels support Internet Key Exchange Protocol Version 2 (IKEv2) in addition to IKEv1. IKEv2 is defined in RFC 5996.

Unlike IKEv1, which uses Phase 1 SA and Phase 2 SA, IKEv2 uses a child SA for Encapsulating Security Payload (ESP) or Authentication Header (AH), which is set up with an IKE SA. IKEv2 offers cookie validation protection, traffic selectors, and a certificate exchange method using Hash and URL. The following table provides the procedure for configuring an IKEv2 gateway and several optional tasks.

| Configure IKEv2 Gateway and Options | | |
|---|---|---|
| Step 1 | Configure an IKEv2 gateway. | Use an IKEv2 gateway to communicate with another IKEv2 gateway and thereby communicate with an IPSec VPN. Configuring an IPv2 gateway includes authenticating the peer at the opposite end of the tunnel. See Set Up an IKE Gateway. |
| Step 2 | (Optional) Import a certificate for IKEv2 gateway authentication. | Import a certificate if you are authenticating a peer for an IKEv2 gateway and you do not use a local certificate on the firewall. See Import a Certificate for IKEv2 Gateway Authentication. |
| Step 3 | (Optional) Export a certificate for a peer to access using Hash and URL. | IKEv2 supports Hash and URL Certificate Exchange as a method of having the peer at the remote end of the tunnel fetch the certificate from a server where you have exported the certificate. Perform this task to export your certificate to that server. See Export a Certificate for a Peer to Access Using Hash and URL. |
| Step 4 | (Optional) Change the key lifetime or authentication interval. | Two IKE crypto profile values, **Key Lifetime** and **IKEv2 Authentication Multiple**, control the establishment of IKEv2 IKE SAs. The key lifetime is the length of time that a negotiated IKE SA key is effective. See Change the Key Lifetime or Authentication Interval for IKEv2. |
| Step 5 | (Optional) Change the cookie activation threshold. | Cookie validation is always enabled for IKEv2; it helps protect against half-SA DoS attacks. You can configure the global threshold number of half-open SAs that will trigger cookie validation. You can also configure individual IKE gateways to enforce cookie validation for every new IKEv2 SA. See Change the Cookie Activation Threshold for IKEv2. |
| Step 6 | (Optional) Configure IKEv2 traffic selectors. | You can configure IKEv2 traffic selectors, which are components of network traffic used during IKE negotiation. Traffic selectors are used during the CHILD_SA (tunnel creation) Phase 2 to set up the tunnel and to determine what traffic is allowed through the tunnel. See Configure IKEv2 Traffic Selectors. |

 © Palo Alto Networks, Inc.

# IPSec VPN Enhancements

You can now enable, disable, refresh or restart an IKE gateway or VPN tunnel to make troubleshooting easier.

- Refresh and Restart Behavior for IKE Gateway and IPSec Tunnel
- Enable or Disable an IKE Gateway or IPSec Tunnel
- Refresh or Restart an IKE Gateway or IPSec Tunnel

## Refresh and Restart Behavior for IKE Gateway and IPSec Tunnel

| | Refresh | Restart |
|---|---|---|
| IKE Gateway (IKE Phase 1) | Updates the onscreen statistics for the selected IKE gateway.<br>Equivalent to issuing a second `show` command in the CLI (after an initial `show` command). | Restarts the selected IKE gateway.<br>**IKEv2**: Also restarts any associated child IPSec security associations (SAs).<br>**IKEv1**: Does not restart the associated IPSec SAs.<br>A restart is disruptive to all existing sessions.<br>Equivalent to issuing a `clear, test, show` command sequence in the CLI. |
| IPSec Tunnel (IKE Phase 2) | Updates the onscreen statistics for the selected IPSec tunnel.<br>Equivalent to issuing a second `show` command in the CLI (after an initial `show` command). | Restarts the IPSec tunnel.<br>A restart is disruptive to all existing sessions.<br>Equivalent to issuing a `clear, test, show` command sequence in the CLI. |

## Enable or Disable an IKE Gateway or IPSec Tunnel

| Enable or Disable an IKE Gateway or IPSec Tunnel | |
|---|---|
| Enable or disable an IKE gateway. | 1. Select **Network > Network Profiles > IKE Gateways** and select the gateway you want to enable or disable.<br>2. At the bottom of the screen, click **Enable** or **Disable**. |
| Enable or disable an IPSec tunnel. | 1. Select **Network > IPSec Tunnels** and select the tunnel you want to enable or disable.<br>2. At the bottom of the screen, click **Enable** or **Disable**. |

# Refresh or Restart an IKE Gateway or IPSec Tunnel

Restarting an IKEv2 gateway has a result different from restarting an IKEv1 gateway. See Refresh and Restart Behavior for IKE Gateway and IPSec Tunnel.

| Refresh or Restart an IKE Gateway or IPSec Tunnel | |
|---|---|
| Refresh or restart an IKE gateway. | 1. Select **Network > IPSec Tunnels** and select the tunnel for the gateway you want to refresh or restart.<br>2. In the row for that tunnel, under the Status column, click **IKE Info**.<br>3. At the bottom of the IKE Info screen, click the action you want:<br>• **Refresh**—Updates the statistics on the screen.<br>• **Restart**—Clears the SAs, so traffic is dropped until the IKE negotiation starts over and the tunnel is created again. |
| Refresh or restart an IPSec tunnel.<br>You might determine that the tunnel needs to be refreshed or restarted because you use the tunnel monitor to monitor the tunnel status, or you use an external network monitor to monitor network connectivity through the IPSec tunnel. | 1. Select **Network > IPSec Tunnels** and select the tunnel you want to refresh or restart.<br>2. In the row for that tunnel, under the Status column, click **Tunnel Info**.<br>3. At the bottom of the Tunnel Info screen, click the action you want:<br>• **Refresh**—Updates the statistics on the screen.<br>• **Restart**—Clears the SAs, so traffic is dropped until the IKE negotiation starts over and the tunnel is created again. |

 © Palo Alto Networks, Inc.

# GlobalProtect Features


- Disable Direct Access to Local Networks
- Static IP Address Allocation
- Apply a Gateway Configuration to Users, Groups, and/or Operating Systems
- Welcome Page Management
- RDP Connection to a Remote Client
- Simplified GlobalProtect License Structure
- SSL/TLS Service Profiles for GlobalProtect Portals and Gateways
- GlobalProtect IPSec Crypto Profiles for GlobalProtect Client Configurations


© Palo Alto Networks, Inc.

# Disable Direct Access to Local Networks

You can now disable direct access to local networks so that users cannot send traffic to proxies or local resources while connected to a GlobalProtect VPN. For example, if a user establishes a GlobalProtect VPN tunnel while connected to a public hotspot or hotel Wi-Fi, and this feature is enabled, all traffic is routed through the tunnel and is subject to policy enforcement by the firewall.

By default, access to local networks is allowed.

You can select the new option **No direct access to local networks** when configuring or modifying the GlobalProtect client configuration of a GlobalProtect gateway.

| Disable Direct Access to Local Networks | | |
|---|---|---|
| Step 1 | Add or modify the network settings of a a gateway client configuration.<br>To configure network settings, you must have enabled **Tunnel Mode** and defined a **Tunnel Interface** on the **Tunnel Settings** tab. | After defining the settings for a tunnel interface:<br>1. Select **Network > Gateways**. Select the name of an existing tunnel configuration, or **Add** a new configuration.<br>2. Select **Client Configuration > Network Settings**. Select the name of an existing configuration or **Add** a new configuration. |
| Step 2 | Disable direct access to local networks. | Select the **Network Settings** tab and then enable **No direct access to local network**. |
| Step 3 | Save your settings. | Click **OK** and **Commit** the changes. |

 © Palo Alto Networks, Inc.

# Static IP Address Allocation

You can now configure a GlobalProtect gateway to assign IP addresses in the following ways:

- **Static IP address allocation**—When you configure an IP address pool, the GlobalProtect gateway can now maintain an index of clients and IP addresses so that the endpoint automatically receives the same IP address for all subsequent GlobalProtect VPN connections. The gateway continues to issue IP addresses in a round-robin fashion until all IP addresses are exhausted. To ensure that an endpoint receives the same address and to avoid IP address conflicts, create an IP address pool large enough to accommodate the number of endpoints.
- **Fixed IP address allocation**—For situations that require a more permanent allocation of IP addresses, a new option on the client configuration enables the GlobalProtect gateway to assign fixed IP addresses using an external authentication server. This is useful when downstream resources such as printers, servers, and applications require a fixed source IP address/IP address pool.

  When a user logs into the GlobalProtect gateway, the external authentication server allocates an IP address for the GlobalProtect client using the Framed-IP-Address attribute. The GlobalProtect client saves this IP address as its preferred IP address and, on subsequent logins, sends this saved preferred IP as a hint to the external authentication server.

  If the authentication server cannot return an IP address or the returned IP address causes a conflict, the gateway allocates an IP address from the static IP address pool instead. Each IP address range or subnet must be unique and cannot overlap.

Use the following procedure to configure the network settings and services to assign the clients' virtual network adapter when an agent establishes a tunnel with the gateway. Network settings are not required in internal gateway configurations in non-tunnel mode because, in this case, the agents use the network settings assigned to the physical network adapter.

| Configure Fixed IP Addressing | | |
|---|---|---|
| Step 1 | Add or modify the network settings of a a gateway client configuration.<br>  To configure network settings, you must have enabled **Tunnel Mode** and defined a **Tunnel Interface** on the **Tunnel Settings** tab. | 1. On the GlobalProtect Gateway dialog, select **Client Configuration > Network Settings**.<br>2. Select the name of an existing configuration or **Add** a new configuration. |
| Step 2 | (New) Configure the user or user group and/or the client operating system to which to apply the client configuration. | • To select a specific user or user group to which this configuration will apply from the list (group mapping must be configured for the list of users and groups to display), click **Add**. You can also create configurations to be deployed to agents in **pre-logon** mode (that is, before the user has logged in to the system) or configurations to be applied to **Any** user.<br>• To deploy configurations based on the specific operating system running on the end system, click **Add** in the OS section, and then select the applicable operating systems (**Android**, **iOS**, **Mac**, or **Windows**). Or leave the value in this section set to **Any** for the configurations to be deployed based on user/group only. |

| Configure Fixed IP Addressing (Continued) | | |
|---|---|---|
| Step 3 | (New) (RADIUS or LDAP authentication servers only) Specify the authentication server IP address pool to use to assign to clients that require fixed IP addresses. | 1. Select the **Network Settings** tab.<br>2. Select the **Retrieve Framed-IP-Address attribute from authentication server** check box.<br>3. In the Authentication Server IP Pool area, click **Add** to specify the subnet or IP address range to use to assign to remote users. When the tunnel is established, an interface is created on the remote user's computer with an address in this range that matches the Framed-IP attribute of the authentication server.<br>  The authentication server IP address pool must be large enough to support all concurrent connections. IP address assignment is fixed and is retained after the user disconnects. |
| Step 4 | Specify the IP address pool to use to assign client IP addresses. | In the IP pool area, click **Add** and then specify the IP address range to use. As a best practice, use a different range of IP addresses from those assigned to clients that are physically connected to your LAN to ensure proper routing back to the gateway. |
| Step 5 | (New) Disable direct access to local networks. | To disable split tunneling including direct access to local networks on Windows and Mac OS systems, select the Disable Direct Access to Local Networks check box. |
| Step 6 | Define what destination subnets to route through the tunnel. | 1. Click **Add** in the **Access Route** area and then enter the routes as follows:<br>• To route all client traffic GlobalProtect (full-tunneling), enter 0.0.0.0/0 as the access route.<br>• To route only some traffic—likely traffic destined for your LAN—to GlobalProtect (split-tunneling), specify the destination subnets that must be tunneled.<br>The firewall supports up to 100 access routes.<br>2. Click **OK**. |
| Step 7 | Specify the network services settings for the clients. | Select the **Network Services** tab and then configure the DNS settings:<br>• You can manually assign the DNS server(s) and suffix, and WINS servers by completing the corresponding fields.<br>• If the firewall has an interface that is configured as a DHCP client, you can set the **Inheritance Source** to that interface and the GlobalProtect agent will be assigned the same settings received by the DHCP client. |
| Step 8 | Save the configuration. | Click **OK** and **Commit** the changes. |

 © Palo Alto Networks, Inc.

# Apply a Gateway Configuration to Users, Groups, and/or Operating Systems

You can now specify one or more users or user groups and/or client operating systems to which to apply to a gateway client configuration. By configuring different IP address pools and access routes you can ensure that devices with different network requirements receive the correct network settings. For example, you can configure different IP address pools and access routes for Windows-based clients or for users in user groups such as Engineering.

| Configure Network Settings by User, User Group, and OS | |
|---|---|
| Step 1 | Add or modify the network settings of a a gateway client configuration.<br>1. Select **Network > GlobalProtect > Gateways**, and then select the name of an existing configuration or **Add** a new configuration.<br>2. Select **Client Configuration > Network Settings**, and then select the name of an existing configuration or **Add** a new configuration. |
| Step 2 | Configure the user or user group and/or the client operating system to which to apply the client configuration:<br>• Click **Add** to select a specific user or user group to which this configuration will apply from the list (group mapping must be configured for the list of users and groups to display). You can also create configurations to be deployed to agents in pre-logon mode (that is, before the user has logged in to the system) or configurations to be applied to any user.<br>• To deploy configurations based on the specific operating system running on the end system, click **Add** in the OS section, and then select the applicable operating systems (Android, iOS, Mac, or Windows). Or leave the value in this section set to **Any** for the configurations to be deployed based on user/group only. |
| Step 3 | Configure additional **Network Settings** that are applicable for the user, user group, and/or client operating system. |
| Step 4 | Click **OK** and **Commit** the changes. |

# Welcome Page Management

The GlobalProtect client configuration now forces the Welcome Page to display each time a user initiates a connection. This prevents the user from dismissing important information such as terms and conditions that may be required by your organization to maintain compliance.

When a user initiates a connection for the first time, the GlobalProtect Welcome Page displays by default. If the users in your organization are not required to view any information when establishing a connection, you can configure a new option to add a check box to the Welcome Page that allows a user to dismiss seeing the Welcome Page at subsequent logins. After selecting the **Do not display this page again** check box at the bottom of the Welcome Page, a user can manually launch the Welcome Page by right-clicking the GlobalProtect icon in the system tray and selecting **Welcome Page**.

To allow a user to dismiss the Welcome Page, select the **Enable "Do not display this welcome page again" checkbox** option when customizing the GlobalProtect agent.

 © Palo Alto Networks, Inc.

# RDP Connection to a Remote Client

The GlobalProtect VPN tunnel functionality has been enhanced to allow users, such as IT Help Desk, to RDP to a remote client device already connected to a GlobalProtect gateway thereby enabling troubleshooting and support for remote Windows users.

When IT Help Desk personnel RDP to the client device, the GlobalProtect app detects the new login and requires the remote user to authenticate with the gateway. If the remote user successfully authenticates with the GlobalProtect gateway before a preconfigured timeout period expires, the gateway reassigns the RDP tunnel to the remote user. This security measure prevents unauthorized access to VPN resources because policy enforcement for traffic through the RDP tunnel is now enforced and logged based on the privileges of the remote user.

Configure the new **User Switch Tunnel Rename Timeout** option when you customize the GlobalProtect agent to specify the grace period, during which the remote user must authenticate with the GlobalProtect gateway (default is zero seconds meaning the remote user is not permitted to authenticate with the gateway; range is one to 600 seconds). When the remote user fails to authenticate within the configured grace period, the GlobalProtect gateway terminates the RDP tunnel.

Changing the **User Switch Tunnel Rename Timeout** value only affects the RDP tunnel and does not rename a pre-logon tunnel when configured. After establishing a pre-logon tunnel, GlobalProtect will rename the tunnel to include the username after the user logs in to the system and successfully authenticates. However if the user fails to authenticate—either from an incorrect password or from ignoring the prompt—GlobalProtect does not rename the tunnel. In this instance, the grace period would not apply and the pre-logon tunnel would remain active until a network change causes the tunnel to disconnect or the inactivity timeout is reached.

# Simplified GlobalProtect License Structure

To simplify GlobalProtect licenses, a portal license is no longer required for any feature. This enables you to use GlobalProtect to provide a secure, remote access or virtual private network (VPN) solution via a single or multiple external gateways, without any GlobalProtect licenses. However, advanced features such as enabling HIP checks and support for the GlobalProtect mobile app for iOS and Android still require a gateway subscription.

To take advantage of the new license structure, you need to upgrade only the device running the GlobalProtect portal to PAN-OS 7.0 or later (the GlobalProtect gateway can run PAN-OS 7.0 or earlier).

 © Palo Alto Networks, Inc.

# SSL/TLS Service Profiles for GlobalProtect Portals and Gateways

GlobalProtect portals and gateways now support SSL/TLS service profiles that specify a certificate and a protocol version or range of versions (now including TLSv1.2) for services that use SSL/TLS. This improves network security by enabling devices to avoid SSL/TLS versions that have known vulnerabilities. If the service request involves a protocol version that is outside the specified range, the device downgrades or upgrades the connection to a supported version.

You can select the SSL/TLS service profile when you configure a service profile as part of an external authentication or two-factor authentication configuration or when you set up access to the GlobalProtect portal.

| Use an SSL/TLS Service Profile Instead of a Certificate | |
|---|---|
| Step 1 | For each desired service, Generate a Certificate. |
| Step 2 | Configure SSL/TLS Service Profiles. |
| Step 3 | Select SSL/TLS service profiles instead of certificates:<br>• For a portal: Select **Network > GlobalProtect > Portals**, **Add** or edit a portal, and select the **Portal Configuration** tab.<br>• For a gateway: Select **Network > GlobalProtect > Gateways**, **Add** or edit a gateway, and select the **General** tab.<br>In the Network Settings area, select the **SSL/TLS Service Profile**. |
| Step 4 | Click **OK** and **Commit** the changes. |

# GlobalProtect IPSec Crypto Profiles for GlobalProtect Client Configurations

You can now use Suite B ciphers to secure network connections between GlobalProtect portals/gateways and endpoints. You can also use Suite B ciphers, including elliptic curve (ECDSA), Elliptic Curve Diffie–Hellman (DH), and Advanced Encryption Standard (AES) GCM certificates, to authenticate administrators and end users (see Suite B Cryptography Support). This enables you to meet Federal network security standards.

After configuring the Suite B ciphers in a GlobalProtect IPSec crypto profile, attach it to the GlobalProtect client configuration of a GlobalProtect gateway.

| Assign a GlobalProtect IPSec Crypto Profile to an IPSec Tunnel | |
|---|---|
| Step 1 | Configure a GlobalProtect IPSec Crypto Profile. |
| Step 2 | Select **Network > GlobalProtect > Gateways** and **Add** or edit a configuration. |
| Step 3 | To assign the profile to a client configuration, select **Client Configuration** and then select the GlobalProtect **IPSec Crypto Profile**. |
| Step 4 | Click **OK** and **Commit** the changes. |

 © Palo Alto Networks, Inc.

# Licensing Features

- Support for Usage-Based Licensing in AWS
- Self-Service License & Subscription Management

# Support for Usage-Based Licensing in AWS

The on-demand or *usage-based* licensing in Amazon Web Services (AWS) allows you to obtain the Amazon Machine Image (AMI) for the VM-Series firewall from the AWS Marketplace and deploy the firewall for use in a Virtual Private Cloud (VPC). This option allows you to consolidate your billing of AWS resources and the usage fees—at an hourly or a yearly rate—for the VM-Series firewall. The following on-demand options are available:

- Bundle one—VM-Series capacity license for the VM-300, Threat Prevention license and a premium support entitlement.
- Bundle two—VM-Series capacity license for the VM-300, with the complete suite of licenses for Threat Prevention, GlobalProtect, WildFire, and PAN-DB URL Filtering capabilities. It includes a premium support entitlement.

As soon as you deploy the firewall, the appropriate licenses and entitlements are activated and the firewall can obtain content and software updates immediately. Because the usage-based licenses are automatically activated, you do not need to use Panorama for managing these licenses. When the firewall is stopped or terminated on the AWS console, the usage-based licenses are suspended or terminated. Therefore, the Self-Service License & Subscription Management capability is not supported on these usage-based licenses.

▲ Launch the VM-Series Firewall in the AWS-VPC

▲ Register the Usage-Based Model of the VM-Series Firewall in AWS

## Launch the VM-Series Firewall in the AWS-VPC

### Launch the VM-Series Firewall On-Demand

- Log into the AWS Marketplace and search for the **VM-Series** firewall. Click the **Launch with EC2 Console** button and follow the instructions to launch an instance of the firewall.

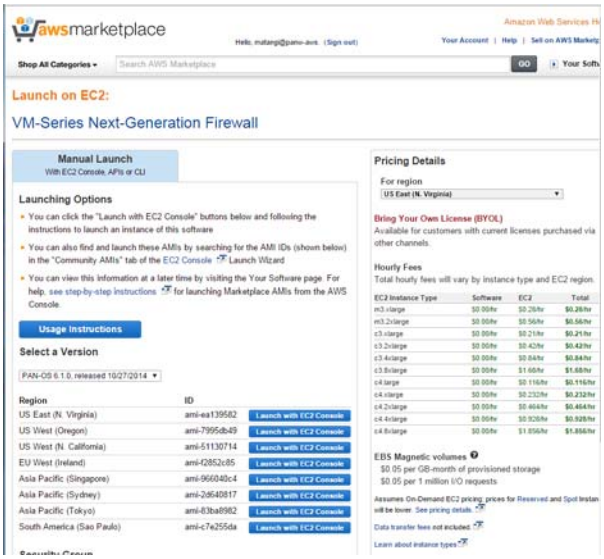


 © Palo Alto Networks, Inc.

| Launch the VM-Series Firewall On-Demand |
|---|
| • Or, Log into the EC2 Console and click the **Launch Instance** button on the Dashboard. From the **AWS Marketplace** tab of the EC2 Console, find and launch the VM-Series AMI by searching for the **VM-Series**.<br> 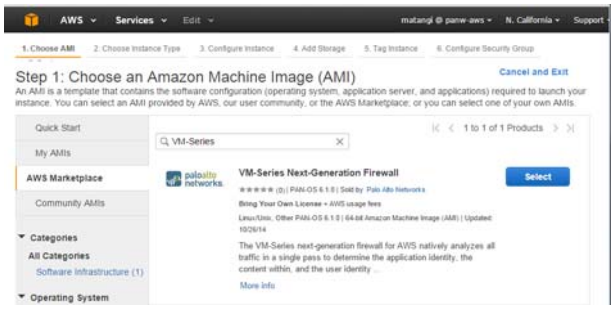 |

For instructions on launching the VM-Series firewall, see Launch the VM-Series Firewall in AWS.

## Register the Usage-Based Model of the VM-Series Firewall in AWS

In order to activate your premium support entitlement with Palo Alto Networks, you must create a support account and register the VM-Series firewall on the Palo Alto Networks Support portal.

| Register the Usage-Based Model of the VM-Series Firewall in AWS | |
|---|---|
| Step 1 Create a support account.<br>If you have a support account with Palo Alto Networks, continue to Step 2. | 1. Log in to https://support.paloaltonetworks.com.<br>2. Click **Register** and fill in the details in the user registration form. To create an account on the support portal, you must have the **AWS Instance ID**, the **AWS Product**, and the **AWS Region** in which you deployed the VM-Series firewall.<br>3. **Submit** the form. You will receive an email with a link to activate the user account; complete the steps to activate the account. After your account is verified and the registration is complete, you will be able to log in. |

| Register the Usage-Based Model of the VM-Series Firewall in AWS | |
|---|---|
| Step 2 Register the VM-Series firewall. | 1. On the **Assets** tab on the Palo Alto Networks Support portal, click the **Register New Device** link.<br> 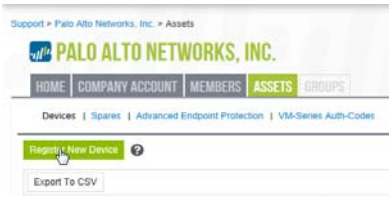 <br><br>2. Select **Register device using AWS Instance ID and Product Name**.<br>3. In the device information section, you must enter the **AWS Instance ID**, the **AWS Product**, and the **AWS Region** in which you deployed the VM-Series firewall.<br> 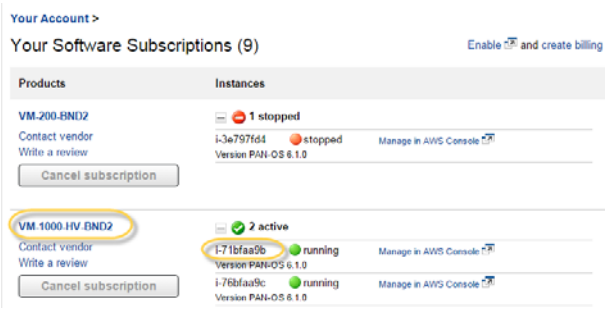 <br><br>You can view the **AWS Instance ID**, and the **AWS Product** you purchased on the **Your Account > Your Software Subscriptions** page of the AWS console. Check on the **EC2 Dashboard > Instances** page for the AWS Region in which the firewall is deployed.<br>4. Verify that the details on the licenses you purchased are displayed on the **Assets** page of the Support portal. |

 © Palo Alto Networks, Inc.

# Self-Service License & Subscription Management

Until now, when you activate the license on a firewall there was no automated method for you to reassign the license to another firewall. Now, the firewall and Panorama provide the capability to unassign or deactivate the active licenses on a firewall and assign the licenses to another firewall. This deactivation functionality allows you to:

- **Deactivate a feature license or subscription on a firewall**—If you accidentally installed a license/subscription on a firewall and need to reassign the license to another firewall, you can deactivate an individual license. This capability is supported on the Command Line Interface (CLI) of the firewall only, both hardware-based firewalls and the VM-Series firewalls. You cannot deactivate a single license/subscription from Panorama.
- **Deactivate a VM-Series firewall**—When you no longer need an instance of the VM-Series firewall, you can free up all active licenses—subscription licenses, VM-Capacity licenses, and support entitlements—using the web interface, CLI, or XML API on the firewall or Panorama. You can then use the licenses on a new instance of a VM-Series firewall, when you need it.

  You can only deactivate the licenses purchased with Palo Alto Networks in the Bring Your Own License (BYOL) model. For the VM-Series on-demand deployment or *usage-based license* in AWS, license deactivation is not supported.

You must always initiate the deactivation process from the firewall or Panorama, and not from the Palo Alto Networks Support portal. If the firewall/Panorama has Internet access and can communicate with the Palo Alto Networks Licensing servers, the license removal process completes automatically with a click of a button. If the firewall/Panorama does not have Internet access, you must complete the process manually in a two-step process. In the first step, from the firewall or Panorama, you generate and export a license token file that includes information on the deactivated keys. In the second step, you upload the token file on the Palo Alto Networks Support portal to dissociate the license keys from the firewall.

▲ Deactivate a Feature License or Subscription

▲ Deactivate VM

 © Palo Alto Networks, Inc.