



TECHDOCS

PAN-OS[®] New Features Guide

Version 10.1

Contact Information

Corporate Headquarters:

Palo Alto Networks

3000 Tannery Way

Santa Clara, CA 95054

www.paloaltonetworks.com/company/contact-support

About the Documentation

- For the most recent version of this guide or for access to related documentation, visit the Technical Documentation portal docs.paloaltonetworks.com.
- To search for a specific topic, go to our search page docs.paloaltonetworks.com/search.html.
- Have feedback or questions for us? Leave a comment on any page in the portal, or write to us at documentation@paloaltonetworks.com.

Copyright

Palo Alto Networks, Inc.

www.paloaltonetworks.com

© 2021-2022 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

Last Revised

September 12, 2022

Table of Contents

App-ID Features.....	5
App-ID Cloud Engine.....	6
SaaS Policy Rule Recommendation.....	8
Management Features.....	9
Audit Tracking for Administrator Activity.....	10
Device Certificate for Cortex Data Lake.....	11
PAN-OS OpenConfig Support.....	12
Installing the Plugin.....	13
Target Address.....	13
Persistent Uncommitted Changes on PAN-OS.....	14
Panorama Features.....	15
Authentication Key for Secure Onboarding.....	16
Optimization for Deploying Changes for Multiple Virtual Systems of the Same Firewall.....	19
Scheduled Configuration Push to Managed Firewalls.....	20
Unique Master Key for a Managed Firewall.....	22
Networking Features.....	23
LSVPN Cookie Expiry Extension.....	24
Persistent NAT for DIPP.....	26
Aggregate Group Members on Multiple Cards.....	28
Network Packet Broker.....	29
Identity Features.....	31
Cloud Identity Engine.....	32
User-ID Features.....	39
Group Mapping Centralization for Virtual System Hubs.....	40
URL Filtering Features.....	43
Enhanced Handling of SSL/TLS Handshakes for Decrypted Traffic.....	44
Advanced URL Filtering Security Subscription.....	48
SD-WAN Features.....	49
Prisma Access Hub Support.....	50
SD-WAN Support for AE and Subinterfaces.....	55
SD-WAN Support for Layer 3 Subinterfaces.....	58
GlobalProtect Features.....	61

Security Policy Enforcement for Inactive GlobalProtect Sessions.....	62
Support for Gzip Encoding in Clientless VPN.....	64
Virtualization Features.....	67
DPDK Support for Different NIC Types.....	68
CN-Series Firewall as a k8s Service.....	69
Intelligent Traffic Offload Service for VM#Series on KVM.....	72
Customize Dataplane Cores.....	73
Mobile Infrastructure Security Features.....	75
5G Multi-access Edge Computing Security.....	76

App-ID Features

- [App-ID Cloud Engine](#)
- [SaaS Policy Rule Recommendation](#)

App-ID Cloud Engine

App-ID Cloud Engine (ACE) is a new service that enables the firewall or Panorama to download App-IDs for unknown SaaS applications from the cloud. ACE converts unknown applications to known applications, vastly increases the number of known App-IDs, speeds up the availability and delivery of new App-IDs, and dramatically increases visibility into applications that previously did not have specific App-IDs.



Using ACE requires a [SaaS Security Inline](#) subscription.

Traditional, content-delivered App-ID only delivers new applications once per month and you need to analyze the new App-IDs before you install them to understand changes that they may make to Security policy rules. The monthly cadence and need for analysis slows down the adoption of new App-IDs in policy. ACE changes that scenario by providing on-demand App-IDs for SaaS applications identified as:

- ssl
- web-browsing
- unknown-tcp
- unknown-udp



Cloud-delivered App-IDs do not identify other types of public applications and do not identify private and custom applications.

Cloud-delivered App-ID provides specific identification of ssl, web-browsing, unknown-tcp, and unknown-udp applications, which enables you to understand them and control them appropriately in policy. The firewall handles cloud App-IDs differently than it handles content-delivered App-IDs. Cloud App-IDs do not force you to examine how the new App-IDs affect Security policy because the firewall uses them according to previously existing Security policy until you do one of the following:

- Create [Application Filters](#) to automate adding downloaded cloud-delivered App-IDs to Security policy.



Use Application Filters as often as possible to automate adding new cloud-delivered App-IDs to Security policy rules. When a new App-ID matches an Application Filter, it is automatically added to the filter. When you use an Application Filter in a Security policy rule, the rule automatically controls the application traffic for App-IDs that have been added to the filter. In other words, Application Filters are your “Easy Button” for securing cloud-delivered App-IDs automatically to gain maximum visibility and control with minimum effort.

- Add the App-IDs to [Application Groups](#).
- Use [Policy Optimizer](#) to add the App-IDs to a cloned rule or to an existing rule, or to an existing Application Filter or Application Group. You can also use Policy Optimizer to create new Application Filters and Application Groups directly from within the Policy Optimizer tool.

See [App-ID Cloud Engine](#) to learn how to:

- Install the SaaS Security Inline license.
- Connect to the ACE cloud and download ACE App-IDs.
- Use ACE App-IDs in Security policy to gain visibility and control over applications that were previously identified only as ssl, web-browsing, unknown-tcp, and unknown-udp traffic.

SaaS Policy Rule Recommendation

The rapid proliferation of SaaS applications makes it difficult to assign all of them specific App-IDs, gain visibility into those applications, and control them, which may introduce security risks to your network. To gain visibility into those applications and control them on the firewall, SaaS Security administrators can recommend [Security policy rules](#) with specific SaaS App-IDs provided by the App-ID Cloud Engine (ACE) to PAN-OS firewall administrators. PAN-OS administrators can import those rules on firewall's that have a [SaaS Security Inline](#) subscription.

[SaaS Security Inline for PAN-OS](#) describes the procedure for pushing Security policy rule recommendations to the firewall and [Import SaaS Policy Recommendation](#) describes how the PAN-OS administrator imports policy recommendations from the SaaS administrator. The high-level process is:

1. The SaaS Security administrator creates the new rule, adds applications, users, and groups to the rule, and sets the rule action. The rule action can be allow or block; no other actions are permitted for pushed rules.
2. The SaaS Security administrator pushes the rule to the appropriate appliances and the rule appears in the firewall interface (**Device > Policy Recommendation > SaaS**).
3. The PAN-OS administrator evaluates the recommended rule and decides whether to implement it on the firewall.
4. If the PAN-OS administrator chooses to implement the rule, the administrator imports it on the firewall and selects where to place the policy rule in the firewall rulebase. When a PAN-OS administrator imports a policy recommendation, the firewall creates the required HIP profiles, tags, and Application Groups automatically so the PAN-OS administrator doesn't have to do it.



If the SaaS Security administrator pushes Security profiles with the policy recommendation and those profiles don't exist on the firewall, the firewall import fails. If the profiles already exist on the firewall, the import succeeds.

If the SaaS Security administrator updates a policy rule recommendation, the PAN-OS administrator sees the update and imports it into the firewall. If the SaaS Security administrator deletes a policy rule recommendation, the PAN-OS administrator sees the action and deletes the rule from the firewall Security policy rulebase.

Management Features

- [Audit Tracking for Administrator Activity](#)
- [Device Certificate for Cortex Data Lake](#)
- [PAN-OS OpenConfig Support](#)
- [Persistent Uncommitted Changes on PAN-OS](#)

Audit Tracking for Administrator Activity

PAN-OS 10.1 introduces the ability to track web administrator activity in the web interface and command line interface (CLI) of firewalls, Panorama™ management server, and Log Collectors for audit purposes. By tracking administrator activity in the web interface and CLI, you can achieve real time reporting of activity across your deployment. If you have reason to believe an administrator account is compromised, you have a full history of where this administrator account navigated throughout the web interface or what operational commands they executed so you can analyze in detail and respond to all actions the compromised administrator took.

An event occurs and generates an audit log, which is forwarded to the specified syslog server each time you navigate through the web interface or when you execute an [operational command](#) in the CLI. Each navigation or command executed generates an audit log. Take for example if you want to create a new address object. You generate one audit log when you click **Objects**, and a second audit log when you then click **Addresses**. Audit logs can only be forwarded to a syslog server, cannot be forwarded to Cortex Data Lake (CDL), and are not stored locally on the firewall, Panorama, or Log Collector.

STEP 1 | [Configure a syslog server profile](#) to forward audit logs of administrator activity on the firewall.

This step is required to successfully store audit logs for tracking administrator activity on the firewall.



For Panorama managed firewalls, the syslog server profile can be configured on the Panorama web interface.

STEP 2 | Select **Panorama > Service Profiles > Syslog** and [configure a syslog server profile](#) to forward audit logs of administrator activity on Panorama and Log Collectors.

This step is required to successfully store audit logs for tracking administrator activity on Panorama.

STEP 3 | [Configure audit tracking of administrator activity](#).

Device Certificate for Cortex Data Lake

PAN-OS 10.1 enables you to connect your firewalls to Cortex Data Lake using the same [device certificate](#) that you use to authenticate to other Palo Alto Networks cloud services such as Cortex XDR, IoT Security, and Enterprise Data Loss Prevention. This simplifies the Cortex Data Lake onboarding process if you already have a device certificate installed.

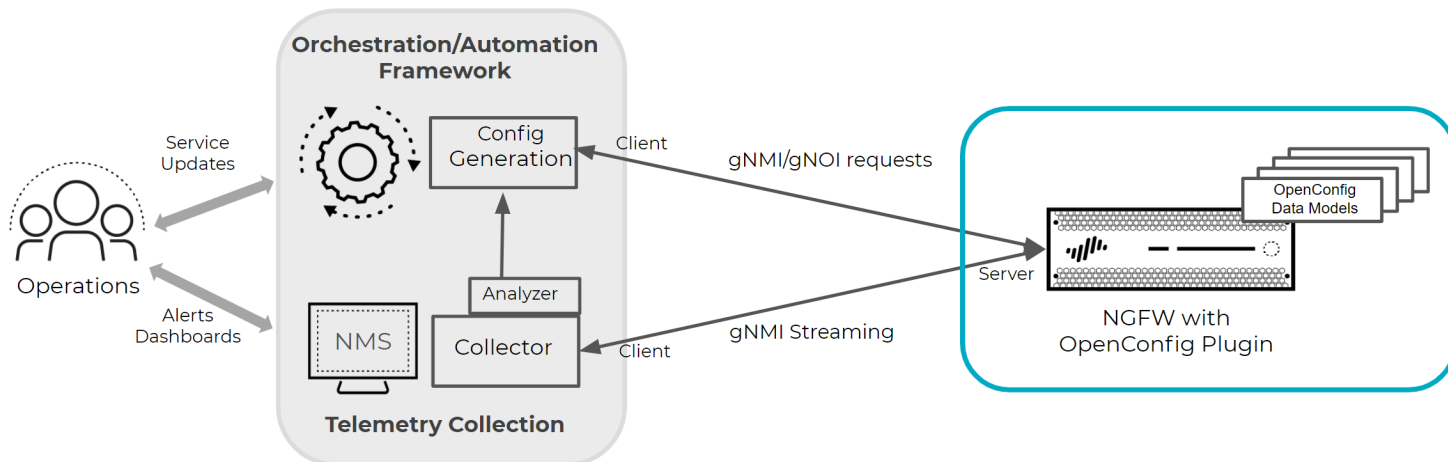
Formerly, devices used a *logging service certificate* specific to Cortex Data Lake to connect to the service. Now, firewalls and Panorama on PAN-OS 10.1 or later can use a device certificate. When onboarding to Cortex Data Lake, make sure to follow the onboarding process appropriate for your [deployment style](#).

PAN-OS OpenConfig Support

Palo Alto Networks OpenConfig plugin allows you to programmatically access the firewall based on OpenConfig data models and protocols to automate configuration and telemetry retrieval. To Learn more about OpenConfig, visit <https://www.openconfig.net>. The OpenConfig interface uses gRPC Network Management Interface (gNMI) protocol for configuration management, telemetry based on the OpenConfig data models, and gRPC Network Operations Interface (gNOI) for operational services defined by OpenConfig.

Using the plugin, you can manage configuration, generate streaming telemetry, and carry out operational services on the firewall. The OpenConfig plugin is supported on the hardware and VM-Series firewalls.

The gNMI protocol uses a client-server messaging model. The OpenConfig plugin implements a gNMI server that listens for client requests and supports all of the gNMI request types: Set, Get, Subscribe, and Capabilities. The Set request carries out transaction based edit operations whether it be single or multiple requests.



Models Supported with v1.0.0

These models are supported with the first version of the plugin:

- openconfig-bgp
- openconfig-vlan
- openconfig-platform
- openconfig-system
- openconfig-interfaces
- openconfig-local-routing
- openconfig-rib
- openconfig-lacp
- openconfig-lldp

Visit the [YANG Repository on the Palo Alto Networks Github](#) for a more comprehensive view of the models.

Installing the Plugin

STEP 1 | Download the plugin by selecting **Device > Plugins** on a PAN-OS firewall.

FILE NAME	VERSION	RELEASE DATE	SIZE	DOWNLOADED	CURRENTLY INSTALLED	ACTIONS	RELEASE NOTE URL
openconfig-1.0.0	1.0.0	2021/02/04 12:37:12	41M	✓		Install Delete	
openconfig-1.0.0	1.0.0	2021/02/22 06:51:09	40M			Install Delete	
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	40M	✓	✓	Remove Config Uninstall	
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	40M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	40M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	40M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	40M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	40M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	41M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	41M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	41M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	41M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	41M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	41M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	41M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	41M			Download	Release Notes
openconfig-1.0.0	1.0.0	2021/02/23 10:57:56	41M			Download	Release Notes

STEP 2 | Select the version of the plugin and click **Install** in the Actions column to install the plugin. PAN-OS will alert you when the plugin is complete.

Target Address

The PAN-OS OpenConfig plugin listens for requests on the management interface’s assigned IP address on port 9339. To send gNMI requests to the firewall, use the management IP address, for example: 10.0.0.1:9339.

If you want to change the IP address for gNMI requests, you should first configure the management interface for the firewall. [How to Configure the Management Interface IP](#) shows how you can set the management IP of a firewall.

Persistent Uncommitted Changes on PAN-OS

In the event of an unforeseen restart of your firewall or Panorama™ management server, PAN-OS now preserves all uncommitted configuration changes locally until a commit executes successfully. An unforeseen restart can be of the PAN-OS device itself or of a PAN-OS management process related to configuration management. This helps eliminate the need to recreate configuration changes when an unforeseen restart occurs. By default, firewalls and Panorama maintain any uncommitted configuration changes on upgrade to PAN-OS 10.1.0 or later release and does not require you to configure or enable anything to leverage this functionality.

For firewalls or Panorama in a high availability (HA) configuration, the uncommitted configuration changes do not automatically sync across the HA peers in the event of an unforeseen restart. To sync uncommitted configuration changes to the firewall or Panorama HA peer, enter the following command:

```
admin> request high availability sync-to-remote candidate-config
```

All uncommitted configuration changes on the firewall or Panorama are fully removed when you revert the running configuration. If you revert specific device group and template configuration on Panorama, only the uncommitted changes associated with the selected device groups and templates are removed.

Panorama Features

- [Authentication Key for Secure Onboarding](#)
- [Optimization for Deploying Changes for Multiple Virtual Systems of the Same Firewall](#)
- [Scheduled Configuration Push to Managed Firewalls](#)
- [Unique Master Key for a Managed Firewall](#)

Authentication Key for Secure Onboarding

To strengthen your security posture when onboarding new firewalls, Dedicated Log Collectors, and WildFire appliances to a Panorama™ management server, PAN-OS 10.1 introduces improved mutual authentication between a new device and Panorama on first connection. You can configure an authentication key to have a specific lifetime, specify the count to determine the number of times the authentication key can be used to onboard new devices, specify one or more serial numbers for which the authentication key is valid, and specify for which devices the authentication key is valid.

To securely onboard a new firewall, you must generate a unique device registration authentication key on Panorama. You then import this authentication key to the device to securely authenticate and connect to Panorama when the device is onboarded for the first time. A system log is generated each time a firewall uses the Panorama-generated authentication key is used. Additionally, the device uses the authentication key to authenticate Panorama when it delivers the device certificate that is used for all subsequent communications.



(PAN-OS 10.1 only) For devices running a PAN-OS 10.1 release, Panorama running PAN-OS 10.1.3 or later release supports onboarding devices running PAN-OS 10.1.3 or later release only. You cannot add a device running PAN-OS 10.1.2 or earlier PAN-OS 10.1 release to Panorama management if Panorama is running PAN-OS 10.1.3 or later release.

Panorama supports onboarding devices running the following releases:

- **Panorama running PAN-OS 10.1.2 or earlier PAN-OS 10.1 release**— Devices running PAN-OS 10.1.2 or earlier PAN-OS 10.1 release, and devices running PAN-OS 10.0 or earlier PAN-OS release.
- **Panorama running PAN-OS 10.1.3 or later release**— Devices running PAN-OS 10.1.3 or later release, and devices running PAN-OS 10.0 or earlier PAN-OS release.

There is no impact to devices already managed by Panorama on upgrade to PAN-OS 10.1.

STEP 1 | [Log in to the Panorama web interface.](#)

STEP 2 | Create the device registration authentication key.

1. Select **Panorama > Device Registration Auth Key** and **Add** a new authentication key.
2. Configure the authentication key.
 - **Name**—Enter a descriptive name for the authentication key.
 - **Lifetime**—Enter the key lifetime to specify how long the authentication key may be used to onboard new firewalls or Log Collectors.
 - **Count**—Specify how many times the authentication key may be used to onboard new firewalls or Log Collectors.
 - **Device Type**—Specify whether the authentication key may be used for **Firewalls, Log Collectors**, or **Any** device.
 - **(Optional) Devices**—Enter one or more device serial numbers to specify for which firewalls or Log Collectors the authentication key is valid.
3. Click **OK**.
4. **Copy Auth Key** and **Close**.

STEP 3 | On Panorama, [add a firewall as a managed device](#).

You must add the device registration authentication key when you configure the Panorama server IP address on the firewall.



The device registration authentication key is required only when manually onboarding a firewall to Panorama management. A device registration authentication key is not required when leveraging Zero Touch Provisioning (ZTP) to onboard a firewall to Panorama management because ZTP provides its own security when onboarding a firewall.

STEP 4 | Add a Dedicated Log Collector to Panorama as a managed collector.

1. On Panorama, [configure a managed collector](#).
2. [Log in to the Log Collector CLI](#).
3. Add the device registration authentication key.

```
admin> request authkey set <auth key>
```

STEP 5 | Add a WildFire appliance to manage with Panorama.

1. On Panorama, [Add Standalone WildFire Appliances to Manage with Panorama](#).
2. Log in to the WildFire CLI and add the device registration authentication key.

```
admin> request authkey set <auth key>
```

- STEP 6 |** Verify that the managed firewall, Log Collector, and WildFire appliance are connected to Panorama.
1. Select **Panorama > Managed Devices > Summary** and verify that the **Device State** for the new device shows as **Connected**.
 2. Select **Panorama > Managed Collectors** and verify that the Run Time **Status** for the Log Collector shows as **Connected**.
 3. Select **Panorama > Managed WildFire Appliances** and verify that the **Connected** status for the WildFire appliance shows as **Connected**.

Optimization for Deploying Changes for Multiple Virtual Systems of the Same Firewall

Device group configuration changes pushed manually or from a [scheduled configuration push](#) of a device groups from the Panorama™ management server to a [multi-vsys](#) firewall are automatically bundled into a single job. When a push is executed from Panorama to managed firewalls, Panorama inspects the managed firewalls associated with the device group push. If Panorama detects that multiple vsys belonging to the same multi-vsyes firewall are associated with a device group push, it bundles the commit job for each vsys into a single commit job on the managed firewall to reduce the overall commit job completion time.

If one of the bundled commit jobs fails, then the entire push fails and you need to push entire the device group configuration changes from Panorama again. Additionally, if multiple multi-vsyes firewalls are included in a push from Panorama and one push fails, then the entire push fails to all firewalls included in the push from Panorama. When you [monitor the device group push](#) locally on the firewall, a single job is displayed rather than multiple individual jobs. If any warnings or failures occur, an error description indicating the impacted vsyes is displayed.

This functionality is supported for multi-vsyes firewalls managed by Panorama running PAN-OS 10.1 and later releases by default.

Scheduled Configuration Push to Managed Firewalls

Often as you accumulate configuration changes on Panorama, you must wait until your off-business hours change management window to push configuration changes to reduce the risk of outages during business hours. To reduce the operational overhead of pushing configuration changes to managed firewalls, PANOS 10.1.0 allows you to create a scheduled configuration push to automatically push changes to your managed firewalls on a specified date and time. You can configure a scheduled configuration push to either occur once or to push on a regularly occurring schedule. This allows you to effectively push configurations made by multiple administrators to multiple firewalls without the need for involvement of any administrator.

To create a scheduled configuration push to managed firewalls, you set the schedule parameters of when and how frequently a push occurs and to which managed firewalls to push to. For a Panorama in a high availability (HA) configuration, the scheduled configuration push is synchronized across the HA peers.



If you create multiple scheduled configuration pushes, you must create them at a minimum of a 15 minute interval to allow for the Panorama management server to validate the configuration. Scheduled configuration pushes that are within 15 minutes of each other may fail due to Panorama being unable to validate the first scheduled configuration push changes.

After a successful scheduled configuration push occurs, you can view the scheduled configuration push execution history to understand when the last push for a specific schedule occurred, and how many managed firewalls were impacted. From the total number of impacted managed firewalls, you can view how many configuration pushes to managed firewalls were successful and how many failed. Of the failed pushes, you can view the total number of managed firewalls with automatically reverted configurations due to a configurations that interrupted the connection between the managed firewall and Panorama.

STEP 1 | [Log in to the Panorama web interface.](#)

STEP 2 | Create a scheduled configuration push.

1. Select **Panorama > Scheduled Config Push** and **Add** a new scheduled configuration push.
2. Configure the scheduled configuration push.
3. In the Push Scope Selection, select one or more device groups, templates, or template stacks.

You must select at least one device group, template, or template stack to successfully schedule a configuration push. All managed firewalls associated with the selected device groups, templates, or template stacks are included in the scheduled configuration push.

1. Select one or more **Device Groups** you want to schedule to push.
2. Select one or more **Templates** you want to schedule to push.



Up to 64 templates are supported for a single scheduled configuration push.

4. Click **OK**.
5. Click **Commit** and **Commit to Panorama**.

- STEP 3 |** View the execution history to verify that the scheduled configuration push for all managed firewalls was successful.
1. Select **Panorama > Scheduled Config Push** and click the Last Executed time stamp in the Status column.
 2. View the execution history for the scheduled configuration push.
 3. Click **Tasks** to view the full operation details for the latest scheduled configuration push.

Unique Master Key for a Managed Firewall

Strengthen your security posture by configuring a unique **master key** for your Panorama™ management server and for each managed firewall. By configuring unique master keys, you can ensure that a compromised master key does not compromise the configuration encryption for your entire deployment. Unique master keys are supported only for Panorama and managed firewalls. Log Collectors and WildFire appliances must share the same master key as Panorama. For Panorama or managed firewalls in a high availability (HA) configuration, you must deploy the same master key for both HA peers as the master key is not synchronized across HA peers. Panorama and managed firewalls support the deployment of unique master keys by default on upgrade to PAN-OS 10.1.

Configuring a unique master key also eases the operational burden of updating your master keys. By configuring a unique master key for a managed firewall, you can update each master key individually without the need to coordinate changing the master key across a large number of managed firewalls.

STEP 1 | Log in to the Panorama web interface.

STEP 2 | (Optional) Select **Device > Master Key and Diagnostic** and edit the **Master Key to Auto Renew With Same Master Key** for your managed firewalls.

Configure this setting to automatically renew the master key deployed on the managed firewalls associated with the selected template. Otherwise, the master key expires per the configured master key lifetime and you must deploy a new master key.

STEP 3 | Configure a unique master key for a managed firewall.

1. Select **Panorama > Managed Devices > Summary and Deploy Master Key**.
2. Select a managed firewall and **Change** the master key.



If you want to deploy a unique master key for a specific set of managed firewalls, you can select those specific managed firewalls as well.

3. Configure the master key.
4. Review the Last Master Key Push column to verify that the master key was deployed successfully to all selected managed firewalls.

A System log generates when you deploy a new master key from Panorama.

STEP 4 | Select **Panorama > Master Key and Diagnostics** and configure a unique master key for Panorama.

STEP 5 | (Optional) Select **Panorama > Master Key and Diagnostic** and edit the **Master Key** setting to configure the Panorama master key to **Auto Renew With Same Master Key**.

Configure this setting to automatically renew the master key deployed on Panorama. Otherwise, the master key expires per the configured master key lifetime and you must deploy a new master key.

STEP 6 | Select **Commit** and **Commit and Push**.

Networking Features

The networking features for PAN-OS 10.1 are documented in the 10.1 Networking Administrator's Guide. The VPN and LSVPN features are documented in PAN-OS Administrator's Guide.

- [\(PAN-OS 10.1.7 and later 10.1 Releases\)LSVPN Cookie Expiry Extension](#)
- [\(PAN-OS 10.1.6 and later 10.1 Releases\)Persistent NAT for DIPP](#)
- [Aggregate Group Members on Multiple Cards](#)
- [Network Packet Broker](#)

LSVPN Cookie Expiry Extension


The satellite administrator manually authenticates the satellite to the portal to establish the first connection. Upon successful authentication, the portal returns a satellite cookie to authenticate the satellite on subsequent connections. The satellite cookie that the portal issues has a lifetime of 6 months. The encrypted cookie stored on an LSVPN satellite expires after every 6 months. As soon as the cookie expires, the satellite administrator must re-authenticate by manually entering their credentials, and a new cookie will be issued by the portal.

This causes the VPN tunnels associated with the satellite to go down, causing an outage until the satellite is re-authenticated to the LSVPN portal or gateway and a new cookie is generated. A re-authentication every six months causes administrative overhead, affecting productivity, network stability, and resources of the company.

You can now configure the cookie expiry period from 1 to 5 years, while the default remains as 6 months (when set to 0). In other words, the cookie expiry period is now configurable up to 5 years.

While [configuration is only done on the portal](#), you must upgrade both portal and satellite versions to PAN-OS 10.1.7 or later 10.1 releases to use this feature effectively.

Use the following operational commands to update or view the cookie expiration period:

Operational Command	Execute On	Description
<pre>username@hostname> request global-protect-portal set-satellite-cookie-expiration</pre>	Portal	<p>Changes the current satellite cookie expiration time (default is 0, range is 1 to 5 years).</p> <p>For Example:</p> <p>To configure the satellite cookie expiration time to 3 years, execute:</p> <pre>username@hostname> request global-protect-portal set-satellite-cookie-expiration value 3</pre> <p> <i>To configure the cookie expiration time from 1 to 5 years, configure the value from 1 to 5. To configure the cookie expiration time for 6 months, configure the value as 0.</i></p>
<pre>username@hostname> show global-protect-portal</pre>	Portal	Displays current satellite cookie expiration time.

Operational Command	Execute On	Description
satellite-cookie-expiration		
username@hostname> show global-protect-satellite satellite	Satellite	Displays current satellite authentication cookie's generation time. The Satellite Cookie Generation Time output field shows the updated time.

On the portal, select **Monitor > System** to view the system log for the updated satellite cookie expiration time.

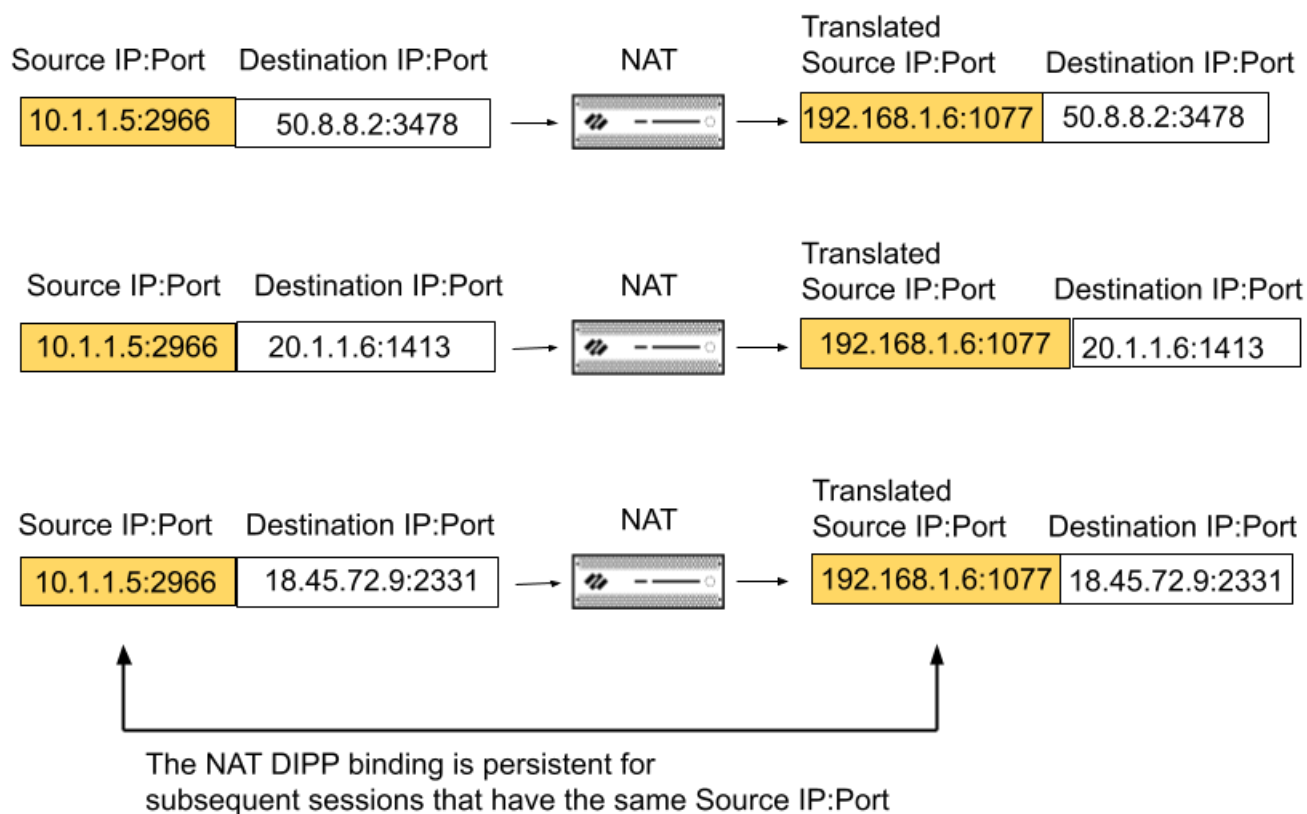
Persistent NAT for DIPP

One type of source NAT is [Dynamic IP and Port \(DIPP\)](#), which allows multiple hosts to have their source IP address translated to a single public IP address with different port numbers.

VoIP, video, cloud-based video conferencing, audio conferencing, and other applications often use DIPP and may require the Session Traversal Utilities for NAT (STUN) protocol. DIPP NAT uses symmetric NAT, which may have compatibility issues with applications that use STUN. To alleviate these issues, persistent NAT for DIPP provides additional support for connectivity with such applications.

Beginning with PAN-OS 10.1.6, persistent NAT for DIPP is available on VM-Series firewalls and single-dataplane firewalls. Beginning with PAN-OS 10.1.7, it is available on all firewalls.

When persistent NAT for DIPP is enabled, the binding of a private source IP address/port pair to a specific public (translated) source IP address/port pair persists for subsequent sessions that come in with the same original source IP address/port pair. The following example shows three sessions:



In this example, original source IP address/port 10.1.1.5:2966 is bound to the translated source IP address/port 192.168.1.6:1077 in Session 1. That binding is persistent in Session 2 and Session 3, which have the same original source IP address/port as Session 1, but different destination addresses. The persistence of the binding ends after all of the sessions for that source IP address/port pair have ended.

In Session 1 of the example, the Destination port is 3478, the default STUN port.

When persistent NAT for DIPP is enabled, it applies to all NAT and NAT64 rules subsequently configured; it is a global setting. Management plane or dataplane logs will indicate NAT DIPP/STUN support has been enabled.

The persistent NAT for DIPP setting (enabled or disabled) survives across firewall reboots.

Perform this task to enable persistent NAT for DIPP.

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

STEP 2 | `>set system setting persistent-dipp enable yes`



*Use the following command to disable persistent NAT for DIPP: **set system setting persistent-dipp enable no.***

STEP 3 | `>request restart system`

STEP 4 | If you have HA configured, repeat this procedure on the other HA peer.

Aggregate Group Members on Multiple Cards

If you have a PA-7050 or PA-7080 firewall that has an [aggregate interface group](#) of interfaces that you located on different line cards, it is a best practice to configure the firewall so that it can handle fragmented IP packets it receives on multiple interfaces of the AE group that are spread over multiple cards. To do so, use the following CLI operational command.

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

STEP 2 | Use the following operational command: **set ae-frag redistribution-policy hash**

Network Packet Broker

The new [Network Packet Broker](#) feature replaces Decryption Broker and expands its capabilities to filter and forward not only decrypted TLS traffic, but also non-decrypted TLS and non-TLS traffic, to one or more third-party appliances (a security chain). The ability to filter and forward all traffic to a security chain eliminates complications from dedicated decryption devices and security chain management devices, thus simplifying your network and reducing capital and operating costs. Network Packet Broker checks path health to and from the security chain and filters traffic based on applications, users, devices, IP addresses, and zones. These features are especially valuable in very high security environments such as financial and government institutions that require offloading traffic to external security chains.

To get started with Network Packet Broker:

1. [Install a free Network Packet Broker license and enable the App-ID cache](#). Without the free license, you can't access the Packet Broker policy and profile configuration.
2. Identify the traffic that you want to forward to one or more security chains.
3. The firewall must have at least two available Layer 3 Ethernet interfaces to use as dedicated packet broker forwarding interfaces to connect to the first and last devices in a security chain. You can configure multiple pairs of packet broker forwarding interfaces to connect to different security chains. Decide which pairs of firewall interfaces to use as dedicated Network Packet Broker forwarding interfaces.



Network Packet Broker supports routed Layer 3 security chains and Transparent Bridge Layer 1 security chains. For routed Layer 3 chains, one pair of packet broker forwarding interfaces can connect to multiple Layer 3 security chains using a properly configured switch, router, or other device to perform the required Layer 3 routing between the firewall and the security chains.

4. Configure a [Transparent Bridge](#) security chain or a [routed layer 3](#) security chain on the firewall using Packet Broker profiles and Network Packet Broker policy rules.



None of the devices in the security chain can modify the source or destination IP address, source or destination port, or protocol of the original session because the firewall would be unable to match the modified session to the original session and therefore would drop the traffic.



You can use [Policy Optimizer](#) to review and tighten Network Packet Broker policy rules.

Network Packet Broker supports:

- Decrypted TLS, non-decrypted TLS, and non-TLS traffic.
- SSL Forward Proxy and SSL Inbound Inspection traffic.
- Routed Layer 3 security chains.
- Transparent Bridge Layer 1 security chains.



You can configure both routed Layer 3 and Layer 1 Transparent Bridge security chains on the same firewall but you must use different pairs of forwarding interfaces for each type.

- Unidirectional traffic flow through the chain: all traffic to the chain egresses the firewall on one dedicated firewall interface and returns to the firewall on another dedicated firewall interface, so all traffic flows in the same direction.



Both firewall forwarding interfaces must be in the same zone.

- Bidirectional traffic flow through the security chain:
 - Client-to-server (c2s) traffic egresses the firewall on one dedicated firewall broker interface and returns to the firewall on another dedicated firewall broker interface.
 - Server-to-client (s2c) traffic uses the same two dedicated firewall broker interfaces as c2s traffic, but the traffic flows in the opposite direction through the security chain. The firewall broker interface on which the s2c traffic goes to the chain is the same interface on which the c2s traffic returns from the chain to the firewall. The firewall broker interface on which the s2c traffic returns to the firewall is the same interface on which the c2s traffic egresses to the chain.



Network Packet Broker does not support multicast, broadcast, or SSH traffic.

Identity Features

- [Cloud Identity Engine](#)

Cloud Identity Engine

The Cloud Identity Engine consists of two components: Directory Sync, which provides user information, and the Cloud Authentication Service, which authenticates users. For a more comprehensive identity solution, Palo Alto Networks recommends using both components, but you can configure the components independently.

The Cloud Authentication Service uses a cloud-based service to provide user authentication using SAML 2.0-based Identity Providers (IdPs). When the user attempts to authenticate, the authentication request is redirected to the Cloud Authentication Service, which redirects the request to the IdP. After the IdP authenticates the user, the firewall maps the user and applies the security policy.


By using a cloud-based solution, you can reallocate the resources required for authentication from the firewall or Panorama to the cloud. The Cloud Authentication Service also allows you to configure the authentication source once instead of for each authentication method you use (for example, Authentication Portal or administrator authentication).

You can now also sync directory changes to the Cloud Identity Engine, which quickly syncs only the recent changes to your directory and takes much less time than a full sync.

- STEP 1 |** Prepare to deploy the Cloud Identity Engine so that it can provide user mappings to the firewall.
1. If you have not already done so, install the [device certificate](#) for your firewall or [Panorama](#).
 2. [Activate](#) the Cloud Identity Engine app.
- STEP 2 |** Configure [Azure Active Directory](#), an [on-premises Active Directory](#), or another [cloud-based directory](#) as your identity source in the Cloud Identity Engine app.

STEP 3 | Configure a Cloud Identity Engine profile on the firewall.

The Cloud Identity Engine retrieves the information for your instance based on your device certificate and uses the Palo Alto Networks Services service route.

1. On the firewall, select **Device > User Identification > Cloud Identity Engine** and **Add** a profile.
2. For the **Instance**, specify each of the following:
 - **Region**—Select the regional endpoint for your instance.
 -  *The region you select must match the region you select when you [activate your Cloud Identity Engine instance](#).*
 - **Cloud Identity Engine Instance**—If you have more than one instance, select the instance you want to use.
 - **Domain**—Select the domain that contains the directories you want to use.
 - **Update Interval (min)**—Enter the number of minutes that you want the firewall to wait between updates. The default is 60 minutes and the range is 1–1440.

Cloud Identity Engine ?

Name

Instance | User Attributes | Group Attributes | Device Attributes

Region

Cloud Identity Engine Instance

Domain

Update Interval (min)

Enabled

3. Verify that the profile is **Enabled**.

Cloud Identity Engine
?

Name

Instance
User Attributes
Group Attributes
Device Attributes

Region ▼

Cloud Identity Engine Instance ▼

Domain ▼

Update Interval (min)

Enabled

4. For the **User Attributes**, select the format for the **Primary Username**. You can optionally select the formats for the **E-Mail** and an **Alternate Username**. You can configure up to three alternate username formats if your users log in using multiple username formats.

Cloud Identity Engine
?

Name

Instance
User Attributes
Group Attributes
Device Attributes

NAME	DIRECTORY ATTRIBUTE
Primary Username	<input style="width: 100%; border: none; border-bottom: 1px solid #ccc;" type="text" value="Common-Name"/> ▼
E-Mail	None
Alternate Username 1	None
Alternate Username 2	None
Alternate Username 3	None

5. For the **Group Attributes**, select the format for the **Group Name**.

The screenshot shows the 'Cloud Identity Engine' configuration window. The 'Name' field is set to 'TechDocs_Example'. The 'Group Attributes' tab is selected, showing a table with two columns: 'NAME' and 'DIRECTORY ATTRIBUTE'. The first row has 'Group Name' in the 'NAME' column and a dropdown menu with 'Name' selected in the 'DIRECTORY ATTRIBUTE' column. The second row has 'E-Mail' in the 'NAME' column and 'None' in the 'DIRECTORY ATTRIBUTE' column. At the bottom right, there are 'OK' and 'Cancel' buttons.

NAME	DIRECTORY ATTRIBUTE
Group Name	Name
E-Mail	None

- For the **Device Attributes**, select the **Endpoint Serial Number**.

If you are using GlobalProtect and you have enabled Serial Number Check, select the Endpoint Serial Number option to allow the Cloud Identity Engine to collect serial numbers from managed endpoints. This information is used by the GlobalProtect portal to check if the serial number exists in the directory for verification that the endpoint is managed by GlobalProtect.

The screenshot shows the 'Cloud Identity Engine' configuration window. The 'Name' field is set to 'TechDocs_Example'. The 'Device Attributes' tab is selected, showing a dropdown menu with 'Endpoint Serial Number' selected and 'Serial Number' in the adjacent field. At the bottom right, there are 'OK' and 'Cancel' buttons.

- Click **OK** then **Commit** your changes.

STEP 4 | Configure [security policy rules](#) for your users (for example, by specifying one or more users or groups that the firewall retrieves from the Cloud Identity Engine as the **Source User**).

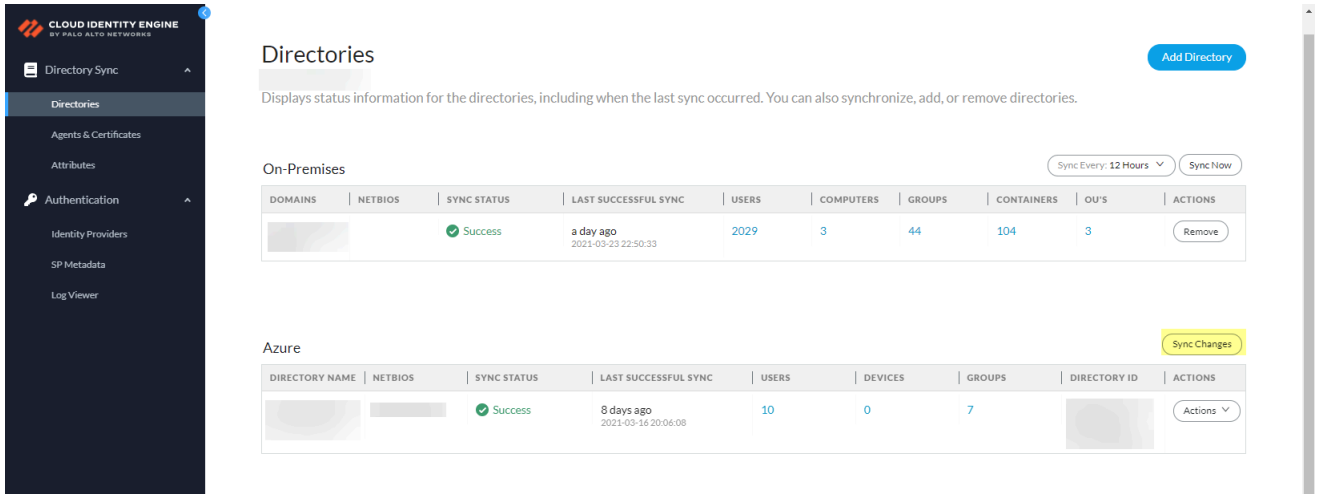
The firewall collects attributes only for the users and groups that you use in security policy rules, not all users and groups in the directory.

STEP 5 | Verify that the firewall has the mapping information from the Cloud Identity Engine.

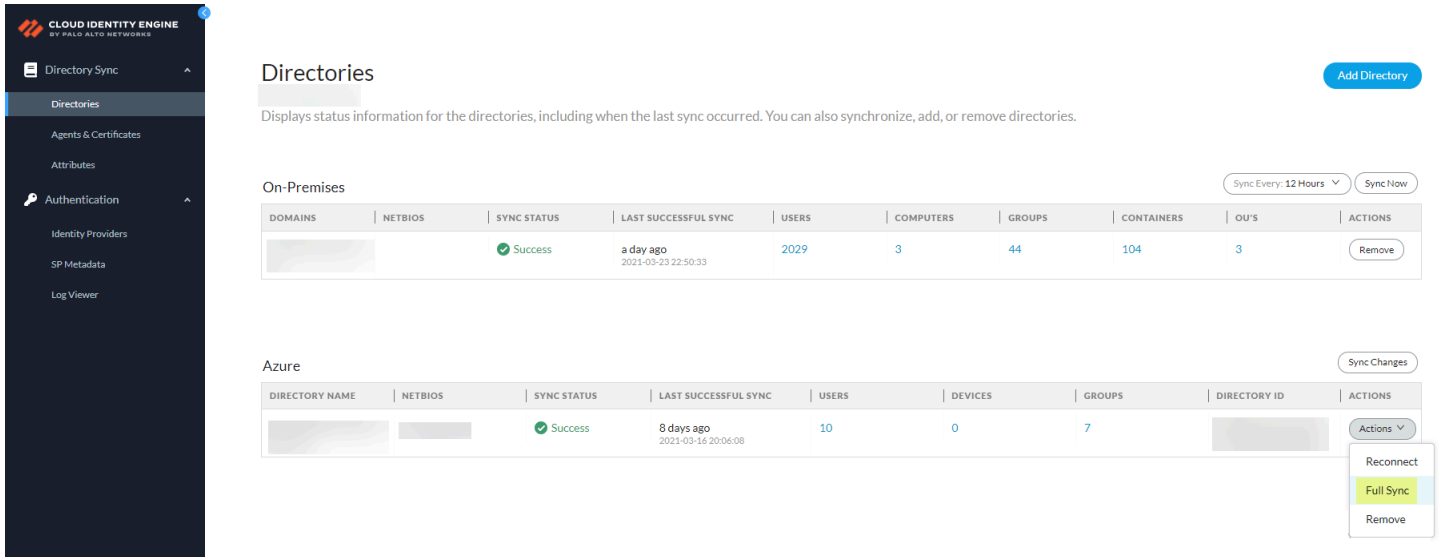
- On the client device, use the browser to access a web page that requires authentication.
- Enter your credentials to log in.
- On the firewall, use the `show user ip-user-mapping all` command to verify that the mapping information is available to the firewall.

STEP 6 | If you make changes to the directory that you configure in the Cloud Identity Engine, **Sync Changes** for your directory.

By default, the Cloud Identity Engine syncs changes every five minutes. If you want to instantly sync your directory updates, you can sync just the changes to your Azure Active Directory (Azure AD) or on-premises AD, which is much faster than a full sync of your directory.



*Palo Alto Networks recommends a full sync of your directory if you lose connectivity or are experiencing issues. To sync the entire directory, select **Actions > Full Sync**. If a full sync is in progress, you cannot sync changes. After a full sync in the Cloud Identity Engine app, the firewall must also complete a full sync.*



STEP 7 | Configure an identity provider (IdP) for the Cloud Identity Engine for user authentication.

STEP 8 | Configure an **Authentication profile** to use the Cloud Authentication Service.

STEP 9 | Configure an **Authentication policy** that uses this Authentication profile.

STEP 10 | Verify that the firewall redirects authentication requests to the Cloud Authentication Service.

1. On the client device, use the browser to access a web page that requires authentication.
2. Enter your credentials to log in.
3. Confirm that the access request redirects to the Cloud Authentication Service.



For more information on the Cloud Identity Engine, refer to the [Cloud Identity Engine Getting Started](#) guide and the [Cloud Identity Engine release notes](#).

User-ID Features

- [Group Mapping Centralization for Virtual System Hubs](#)


Group Mapping Centralization for Virtual System Hubs

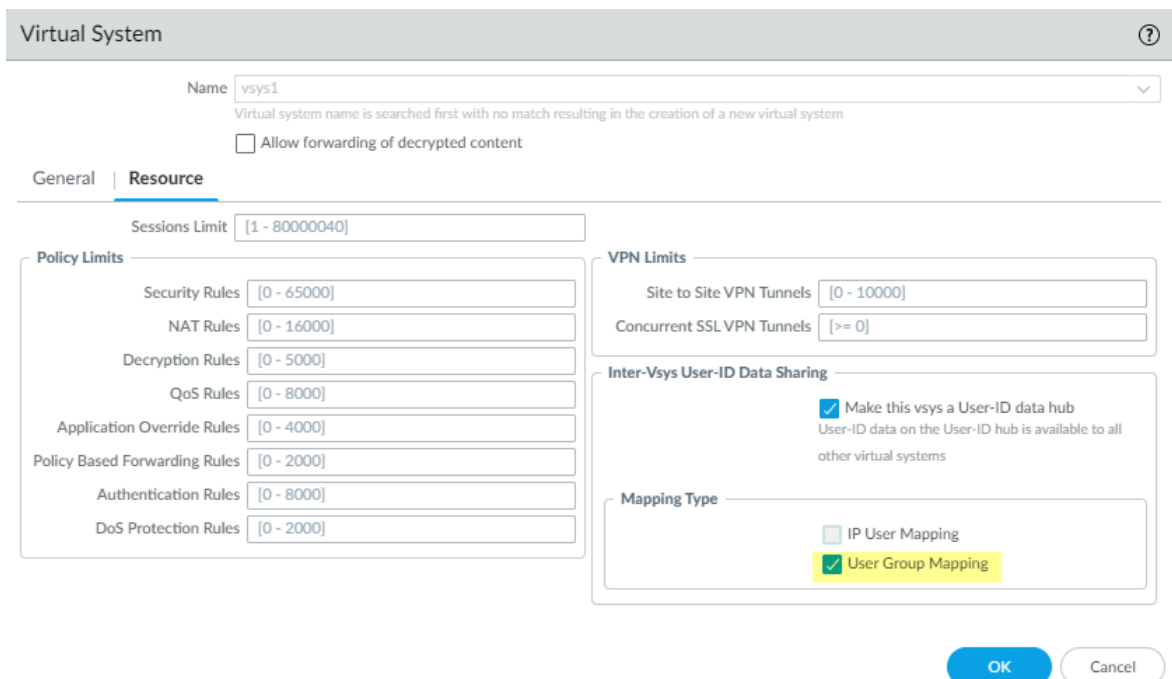
To simplify group-based policy configuration and enforcement, you can now [share group mappings](#) across virtual systems. When you configure a virtual system as a hub, other virtual systems can refer to the hub for mappings when they need to identify groups instead of each virtual system collecting the information independently.

-  *If the same group mapping on the local firewall differs from the group mapping on the virtual system hub, the firewall uses the local mapping.*
-  *Use the same format for the Primary Username across virtual systems and firewalls.*

STEP 1 | Assign the [virtual system](#) as a User-ID hub.

STEP 2 | Confirm **User Group Mapping** as the **Mapping Type** that you want to share then click **OK**.

-  *You must select at least one mapping type.*



STEP 3 | Follow the best practices to [consolidate your User-ID sources](#) on the hub and then remove the duplicate sources from the existing virtual systems.

STEP 4 | **Commit** your changes to enable the User-ID hub and begin collecting mappings for the consolidated sources.

If the group mapping on a firewall differs from the group mapping on the hub, the group mapping on the firewall overrides the group mapping on the hub.

STEP 5 | Confirm the User-ID hub is mapping the groups by entering the following commands:

- `show user group-mapping statistics`
- `show user group-mapping state all`
- `show user group list`
- `show user group name <group-name>`

URL Filtering Features

- [Enhanced Handling of SSL/TLS Handshakes for Decrypted Traffic](#)
- [Advanced URL Filtering Security Subscription](#)

Enhanced Handling of SSL/TLS Handshakes for Decrypted Traffic

The firewall now [inspects the SSL/TLS handshakes](#) of web traffic marked for decryption to block potential threats as early as possible. Specifically, the Content and Threat Detection (CTD) engine on the firewall inspects the Server Name Indication (SNI) field, an extension to the TLS protocol found in the Client Hello message. The SNI field contains the hostname for the website requested by the client. The firewall can use the hostname (if available) to classify the HTTPS traffic, determine its destination, and enforce the matching Security policy rules. For example, the firewall blocks a web session immediately if the domain in the SNI field belongs to a malicious URL category, provided that you have enabled your firewalls to decrypt traffic in malicious URL categories, block malicious domains, and inspect SSL/TLS handshake messages. The inspection also addresses concerns that malicious actors may exploit fields in the handshake to evade Security policy and exfiltrate data.

To take advantage of this capability, you must have an active URL Filtering license, enable SSL/TLS decryption of web traffic, and block URL categories in Security policy rules. You must enable this feature in your SSL decryption settings.

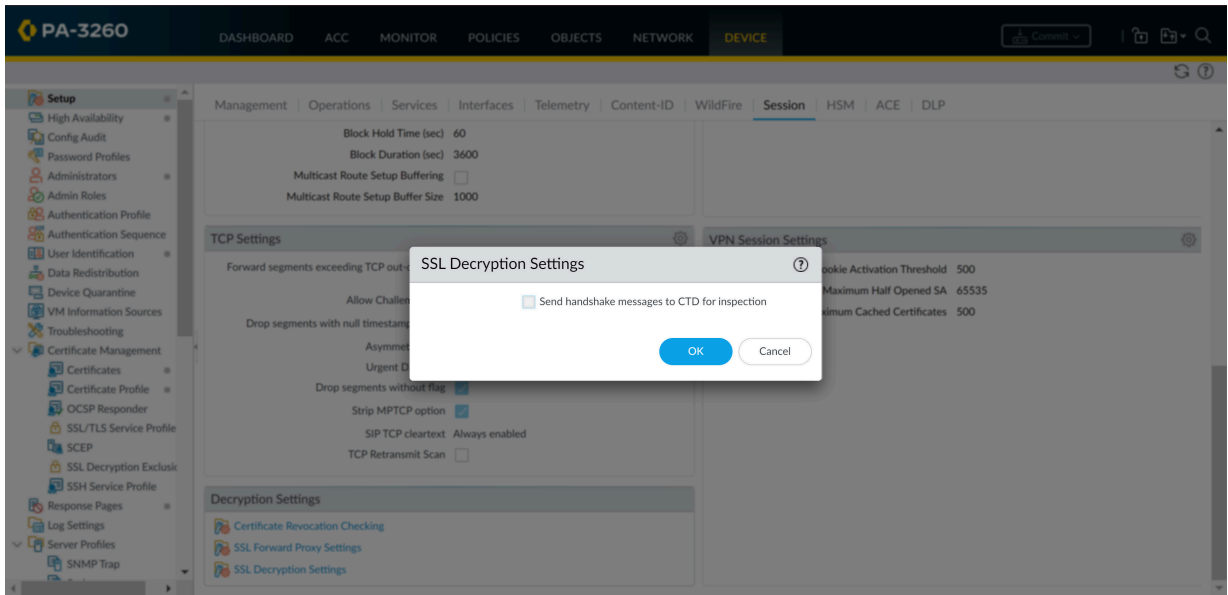
STEP 1 | Select **Device > Licenses** to confirm that you have an [active URL Filtering license](#).

STEP 2 | Set up decryption on your web traffic.

1. Configure [SSL Forward Proxy](#) or [SSL Inbound Inspection](#).

STEP 3 | Enable inspection of SSL/TLS handshakes by CTD.
By default, the option is disabled.

URL Filtering Features



1. Select **Device > Setup > Session > Decryption Settings > SSL Decryption Settings**.
2. **Send handshake messages to CTD for inspection.**

Alternatively, you can use the **set deviceconfig setting ssl-decrypt scan-handshake <yes|no>** CLI command.

STEP 4 | Commit your changes.

Advanced URL Filtering Security Subscription

Palo Alto Networks [Advanced URL Filtering](#) subscription provides real-time URL analysis and malware prevention to generate a more accurate analysis of URLs than possible with traditional web database filtering techniques alone. This subscription service is available on firewalls operating PAN-OS 9.0 and later, with the installation of content release 8390-6607 and later.

Malicious URLs can be updated or introduced before URL filtering databases have an opportunity to analyze the content; this lag time gives attackers an open period from which they can launch precision attack campaigns. Advanced URL filtering compensates for the coverage gaps inherent in database solutions by providing real-time URL analysis per request. When a user visits a URL designated as risky, the firewall submits the URL to the advanced URL filtering service for machine learning analysis and searches PAN-DB for the site's category (information for recently visited websites is cached for fast retrieval). The analysis data is used to generate a verdict that the firewall retrieves to enforce the web-access rules in your policy configuration. If there is a verdict mismatch while the data is being analyzed in the cloud, the more severe categorization takes precedence.

Advanced URL filtering is enabled in a URL Filtering profile and uses the same configuration settings. If you already have an operational URL filtering deployment, no additional configuration is necessary to take advantage of advanced URL filtering—all web requests designated as risky are automatically forwarded for analysis. URLs analyzed using advanced URL filtering are displayed in the logs with the category `real-time-detection`, in addition to the threat type.



The Advanced URL Filtering security subscription is not available on CN-Series firewalls.

STEP 1 | [Install the advanced URL filtering license and verify the installation.](#)

STEP 2 | [Download and install the latest PAN-OS content release.](#)



Follow the [Best Practices for Applications and Threats Content Updates](#) when updating to the latest content release version.

STEP 3 | Verify that you have an active URL Filtering profile. If none is configured, create and [configure a URL Filtering profile](#).

STEP 4 | [Verify that URLs are being analyzed and categorized using the advanced URL Filtering service.](#)

Next Steps:

- [Read the advanced URL filtering blog post.](#)

SD-WAN Features

- [Prisma Access Hub Support](#)
- [SD-WAN Support for AE and Subinterfaces](#)
- [SD-WAN Support for Layer 3 Subinterfaces](#)

Prisma Access Hub Support

With SD-WAN plugin 2.2 and later releases, PAN-OS Secure SD-WAN provides you with [Prisma Access hub support](#) to give you full control of how and where applications are secured. Prisma Access Hub support allows PAN-OS firewalls to connect to Prisma Access compute nodes (CNs) to achieve cloud-based security in an SD-WAN hub-and-spoke topology. This support enables a seamless link failover from on-premises security to Prisma Access and the ability to mix both to meet your security needs.

In a mixed topology with both SD-WAN firewalls and Prisma Access hubs, the SD-WAN hubs are Prisma Access CNs (IPSec Termination Nodes) and the SD-WAN branches are PAN-OS firewalls. SD-WAN automatically creates IKE and IPSec tunnels that connect the branch to the hub. Using Traffic Distribution profiles, you can create SD-WAN policies to match specific internet applications and redirect them to a PAN-OS firewall or Prisma Access deployment of your choice. With Prisma Access hub support, on-premises and cloud security platforms work together to provide a complete solution with consistent security policies managed by Panorama.

The minimum [PAN-OS and SD-WAN plugin versions](#) required for Prisma Access Hub support are:


	Minimum Release
PAN-OS	10.0.8
Prisma Access Compute Node	10.0.7
Prisma Access Cloud Configuration Plugin	2.1
SD-WAN Plugin	2.2
Panorama	10.1.0

Before you connect SD-WAN to Prisma Access, you must have a branch firewall with an interface that has SD-WAN enabled. You must also have performed the [Prisma Access prerequisites](#).


STEP 1 | [Log in to the Panorama Web Interface.](#)

STEP 2 | Specify the BGP local address pool for loopback addresses.

1. Select **Panorama > SD-WAN > VPN Clusters.**
2. Select **BGP Prisma Address Pool.**



3. **Add** an unused private subnet (prefix and netmask) for the local BGP addresses for Prisma Access.



4. Click **OK.**
5. **Commit.**

STEP 3 | Select the SD-WAN branch firewall to connect to the Prisma Access hub and configure the connection.

1. Select **Panorama > SD-WAN > Devices**.
2. Select the branch firewall on which you enabled SD-WAN, whose name then populates the **Name** field.
3. Select the **Type** of device as **Branch**.
4. Select the **Virtual Router Name**.
5. Enter the **Site**.



All SD-WAN devices must have a unique Site name.

6. Select **Prisma Access Onboarding** and **Add**.

BGP												
INTERFACES	TENANT NAME	REGIONS	IPSEC TERMINAT... NODES	LINK TAG	BGP	ADVERTISE DEFAULT ROUTE	SUMMARI... MOBILE USER ROUTES BEFORE ADVERTISI...	DON'T ADVERTISE PRISMA ACCESS ROUTES	PRISMA AS NUMBER	TUNNEL MONITOR IP	SERVICE IP	COMMENT
<input type="checkbox"/>	ethernet1/1	SDWAN_...	us-west-2 us- northwest- longan	Prisma-DIS- VIF	true	false	false	false				

7. Select a local, SD-WAN-enabled **Interface** on the firewall to connect to the Prisma Access hub.
 8. Select a Prisma Access **Tenant** (select **default** for a single tenant environment).
- All SD-WAN interfaces on a branch firewall must use the same Prisma Access tenant.

9. **Add** a compute node to a **Region** by selecting the region where the CN (Prisma Access hub) is located.

There can be multiple regions per interface.

10. Select an **IPSec Termination Node** (GP gateway) from the list of nodes; the list is based on the nodes that Prisma Access spun up for the region earlier. You are choosing the hub to which this branch connects. SD-WAN Auto VPN configuration builds IKE and IPsec relationships and tunnels with this node.
11. **Enable** BGP for communication between the branch and hub (Enable is the default).
12. [Complete the configuration](#) for the connection.
13. Click **OK**.

STEP 4 | Commit and Push the configuration to the cloud, where Prisma Access spins up the correct number of IPsec Termination Nodes based on requested bandwidth.

STEP 5 | [Verify that onboarding is complete.](#)

STEP 6 | [Synchronize the branch firewall to Prisma Access](#) to retrieve the service IP address(es) of the CNs.

STEP 7 | **Commit** to Panorama.

STEP 8 | **Push to Devices** to push to the local branch firewall. **Edit Selections** to select the Push Scope Selection. Select the correct **Template** and **Device Group**.

STEP 9 | On the branch firewall, select **Network > Interfaces > SD-WAN** and see the new interface; [Verify the IPSec tunnel and IKE gateway are up.](#)

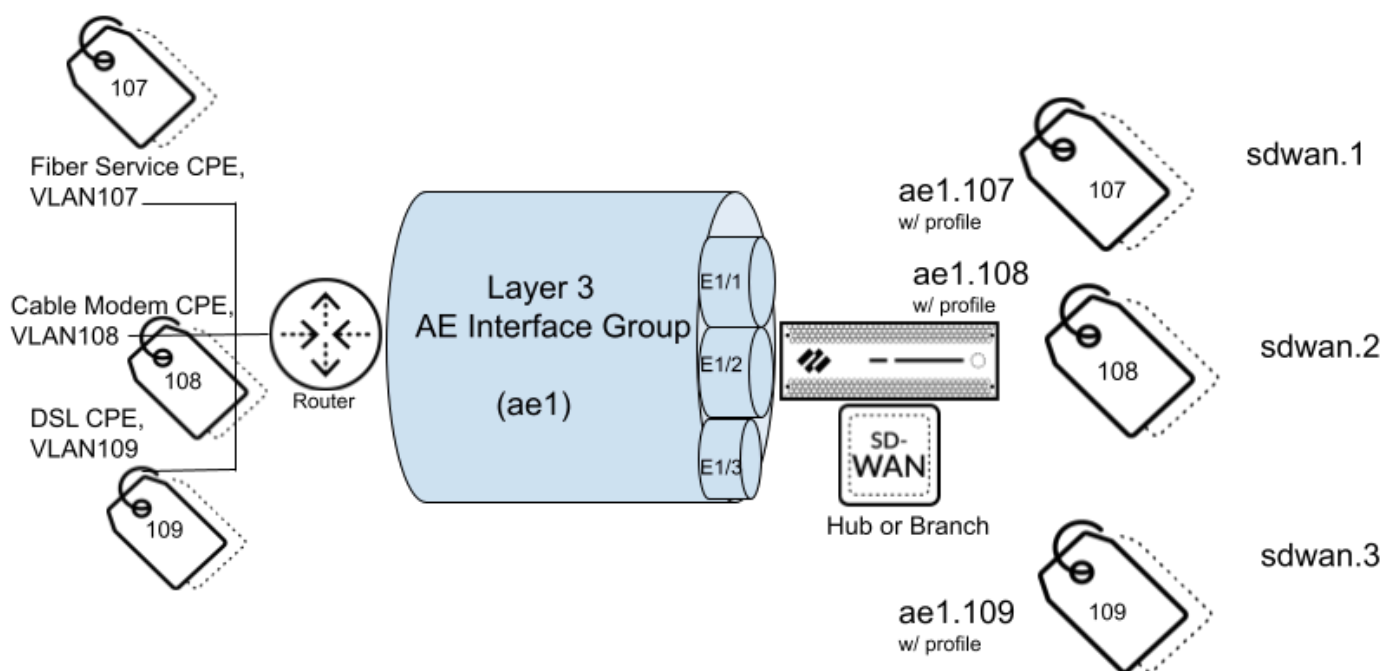
STEP 10 | [Create an SD-WAN policy rule](#) to generate monitoring data.

STEP 11 | **Commit** and **Commit and Push** to branch firewalls.

STEP 12 | [Monitor Prisma Access Hubs.](#)

SD-WAN Support for AE and Subinterfaces

Physical firewalls running PAN-OS 10.1 and SD-WAN Plugin 2.1.0 support SD-WAN on aggregated Ethernet (AE) interfaces so that an SD-WAN firewall in a data center, for example, can have an aggregate interface group (bundle) of physical Ethernet interfaces that provide link redundancy. SD-WAN supports AE interfaces with or without subinterfaces. You can create an AE interface with subinterfaces that you can tag for different ISP services in order to provide end-to-end traffic segmentation. Thus, your ISP services can reach multiple labs or buildings without needing a dedicated pair of fibers for each connection. A Layer 3 AE interface group connects to a router:



VM-Series firewalls do not support AE interfaces. An SD-WAN hub or branch firewall that has an AE interface should not belong to the same VPN cluster as a VM-Series SD-WAN hub or branch firewall because AE interfaces are not supported on VM-Series firewalls.

The following task illustrates how to create an AE interface group, select its member Layer 3 interfaces, create a subinterface for each ISP (using a static IP address or DHCP), assign a VLAN tag to each subinterface, and enable SD-WAN on each subinterface. Create an SD-WAN interface profile to define each ISP connection and assign the profile to the corresponding subinterface (a virtual SD-WAN interface).

STEP 1 | [Log in to the Panorama Web Interface.](#)

STEP 2 | [Create an SD-WAN Interface Profile](#) for each ISP connection (subinterface) in the AE interface group.

STEP 3 | Create a Layer 3 AE interface group.

STEP 4 | Assign physical interfaces to the aggregate group.

STEP 5 | For the aggregate group, create a subinterface that uses a static IP address.

1. Select **Network > Interfaces > Ethernet**, highlight the aggregate interface, such as ae1, and click **Add Subinterface** at the bottom of the screen.
2. Configure the subinterface.

Layer3 Aggregate Subinterface

Interface Name: ae1 . 107

Comment:

Tag: 107

Netflow Profile: None

Config | **IPv4** | IPv6 | SD-WAN | Advanced

Enable SD-WAN

Enable Bonjour Reflector

Type: Static DHCP Client

IP	NEXT HOP GATEWAY
<input checked="" type="checkbox"/> 10.1.1.100/24	10.1.1.1

+ Add - Delete ↑ Move Up ↓ Move Down

IP address/netmask. Ex. 192.168.2.254/24

OK Cancel

STEP 6 | Alternatively, for the aggregate group, create a subinterface that uses DHCP to get its address.

1. Select **Network > Interfaces > Ethernet** and in the **Template** field, select a Template Stack.
2. Highlight the aggregate interface, such as ae1, and click **Add Subinterface** at the bottom of the screen.
3. Highlight the subinterface and click **Override**.
4. Continue to configure the subinterface, selecting the DDNS vendor as **Palo Alto Networks DDNS**.

Layer3 Aggregate Subinterface

Interface Name: ae16.1
 Comment: as1
 Tag: 1
 Netflow Profile: None

Config | IPv4 | IPv6 | SD-WAN | **Advanced**

Other Info | ARP Entries | ND Entries | NDP Proxy | **DDNS**

Settings

Enable
 Certificate Profile: None
 Update Interval (days): 1
 Hostname: ae16-1
 Vendor: Palo Alto Networks DDNS

NAME	VALUE
TTL (sec)	30 [5 - 300]

IP ^
 DHCP

STEP 7 | Apply an SD-WAN Interface Profile to the subinterface.

STEP 8 | Repeat the prior steps to create additional Layer3 subinterfaces for the aggregate interface group and apply an SD-WAN Interface Profile to each subinterface.

STEP 9 | Commit.

SD-WAN Support for Layer 3 Subinterfaces

Firewalls running PAN-OS 10.1 and SD-WAN Plugin 2.1.0 support SD-WAN on Layer 3 subinterfaces so that the firewall can segment traffic using VLAN tags. The following task shows how to create a Layer3 subinterface that uses a static IP address and how to create one that uses DHCP to get its address. It shows how to assign a VLAN tag to the subinterface and enable SD-WAN on the subinterface. Create an SD-WAN interface profile to define each ISP connection and assign the profile to the corresponding subinterface (a virtual SD-WAN interface).



If you configure SD-WAN Layer 3 subinterfaces on VM-Series firewalls, the VMware configuration must have respective portgroups attached to those interfaces that allow all VLANs.

STEP 1 | [Log in to the Panorama Web Interface.](#)

STEP 2 | [Create an SD-WAN Interface Profile](#) for each ISP connection (subinterface).

STEP 3 | [Configure a Layer 3 subinterface](#) that uses a static IP address.

1. Select **Network > Interfaces > Ethernet** and in the **Template** field, select a template.
2. Select an interface.
3. For **Interface Type**, select **Layer3** and click **OK**.
4. Highlight the interface and click **Add Subinterface** at the bottom of the screen.
5. Continue to configure the subinterface.

Layer3 Subinterface

Interface Name: ethernet1/1 | 104

Comment:

Tag: 104

Netflow Profile: None

Config | **IPv4** | IPv6 | SD-WAN | Advanced

Enable SD-WAN
 Enable Bonjour Reflector

Type: Static DHCP Client

IP	NEXT HOP GATEWAY
<input type="checkbox"/> 192.168.16.1/24	192.168.16.2

+ Add - Delete ↑ Move Up ↓ Move Down

IP address/netmask. Ex. 192.168.2.254/24

OK Cancel

- STEP 4 |** Alternatively, [Configure a Layer 3 subinterface](#) that uses DHCP to get its address.
1. Select **Network > Interfaces > Ethernet** and in the **Template** field, select a template stack (not a template).
 2. Select an interface.
 3. For **Interface Type**, select **Layer3** and click **OK**.
 4. Highlight the interface and click **Add Subinterfaces** at the bottom of the screen.
 5. Highlight the subinterface and click **Override**.
 6. Continue to configure the subinterface, selecting the DDNS vendor as **Palo Alto Networks DDNS**.

Layer3 Subinterface

Interface Name: ethernet1/1

Comment:

Tag: 1

Netflow Profile: None

Config | IPv4 | IPv6 | SD-WAN | **Advanced**

Other Info | ARP Entries | ND Entries | NDP Proxy | **DDNS**

Settings

Enable

Update Interval (days): 1

Certificate Profile: None

Hostname: 1_1-1

Vendor: Palo Alto Networks DDNS

IPv4 | IPv6

NAME	VALUE
TTL (sec)	30 [5 - 300]

+ Add - Delete

OK Cancel

STEP 5 | Apply an SD-WAN Interface Profile to the subinterface.

STEP 6 | Repeat the prior steps to add more subinterfaces to the interface.

STEP 7 | Commit.

GlobalProtect Features

- [Security Policy Enforcement for Inactive GlobalProtect Sessions](#)
- [Support for Gzip Encoding in Clientless VPN](#)

Security Policy Enforcement for Inactive GlobalProtect Sessions

You can now enforce a security policy rule to track traffic from endpoints while end users are connected to GlobalProtect and to quickly log out inactive GlobalProtect sessions. You can now enforce a shorter inactivity logout period. If a GlobalProtect session remains inactive during the configured time period, the session is automatically logged out and the VPN tunnel is terminated. By enforcing a security policy, you can quickly gain visibility into active user sessions, and better utilize the gateway resources so that the tunnel IP address and memory assigned to sessions are quickly available for reuse. When you configure an internal gateway in non-tunnel mode, GlobalProtect will continue to enforce the **Inactivity Logout** based on several missing HIP reports because the gateway may not be in accordance with identifying active traffic per user session.

STEP 1 | Specify a shorter amount of time after which idle users are logged out of GlobalProtect.

1. [Launch the Web Interface.](#)
2. Select **Network > GlobalProtect > Gateways > <gateway-config> > Agent > Connection Settings.**
3. Specify the amount of time after which idle users are logged out of GlobalProtect (range is 5 to 43200 minutes; default is 180 minutes).

Users are logged out of GlobalProtect if the GlobalProtect app has not routed traffic through the VPN tunnel or if the gateway does not receive a HIP check from the endpoint within the configured time period.

You must specify the **Inactivity Logout** period to be greater than the **Automatic Restoration of VPN Connection Timeout** to allow GlobalProtect to attempt to reestablish the connection after the tunnel is disconnected (range is 0 to 180 minutes; default is 30 minutes). When you configure an internal gateway in non-tunnel mode, the **Inactivity Logout** period must be greater than the current HIP check interval value that the GlobalProtect app waits before it sends the HIP report.

GlobalProtect Gateway Configuration ?

General | Tunnel Settings | Client Settings | Client IP Pool | Network Services | **Connection Settings** | Video Traffic | HIP Notif

Authentication

Agent

Satellite

Timeout Configuration

Login Lifetime: Hours

Inactivity Logout:
Users are logged out of GlobalProtect when the GlobalProtect app has not sent traffic through the VPN tunnel in the specified amount of minutes.

Authentication Cookie Usage Restrictions

Disable Automatic Restoration of SSL VPN
If the Automatic Restoration of VPN Connection setting is enabled in the GlobalProtect Portal, this setting can be used to disable it for this gateway.

Restrict Authentication Cookie Usage(for Automatic Restoration of VPN tunnel or Authentication Override) to:

The original Source IP for which the authentication cookie was issued The original Source IP network range
Specify using a netmask, the range of source IP addresses from which the authentication cookie can be used.


STEP 2 | Click **OK** twice.

STEP 3 | Commit the configuration.

Support for Gzip Encoding in Clientless VPN

With Palo Alto Networks next-generation firewall deployments, you can now allow Clientless VPN users to access Gzip-compressed websites to use both internal and SaaS applications. Support for Gzip encoding ensures that the Gzip encoding request within the HTTP header is accepted by the Clientless VPN portal. This ensures that the content from the Gzip-compressed web pages is rendered correctly when accessed through the Clientless VPN portal.

The following diagram illustrates the extended support to allow users to access internal and SaaS applications through Clientless VPN.



The Clientless VPN can determine whether to use Gzip encoding based on the HTTP request from the client and the corresponding response from the app. The **gzip** value must be included as one of the Accept-Encoding header values so that it is accepted by the Clientless VPN.

For example, consider the following scenarios when the Clientless VPN uses Gzip encoding:

1. The browser sends an HTTP request to the website with the Accept-Encoding header values set to **gzip**, **deflate**, and **br**, as shown in the following example.

```

Hypertext Transfer Protocol
[Expert Info (Warning/Security): Unencrypted HTTP protocol detected over encrypted port, could indicate a dangerous misconfiguration.]
GET /http/www.facebook.com/ HTTP/1.1\r\n
Host: 10.6.16.51\r\n
Connection: keep-alive\r\n
sec-ch-ua: " Not;A Brand";v="99", "Google Chrome";v="91", "Chromium";v="91"\r\n
sec-ch-ua-mobile: ?0\r\n
Upgrade-Insecure-Requests: 1\r\n
User-Agent: Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/91.0.4472.124 Safari/537.36\r\n
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,image/avif,image/webp,image/apng,*/*;q=0.8,application/signed-exchange;v=b3;q=0.9\r\n
Sec-Fetch-Site: same-origin\r\n
Sec-Fetch-Mode: navigate\r\n
Sec-Fetch-User: ?1\r\n
Sec-Fetch-Dest: document\r\n
Referer: https://10.6.16.51/global-protect/portal/portal.esp?consent=true\r\n
Accept-Encoding: gzip, deflate, br\r\n
Accept-Language: en-US,en;q=0.9\r\n
    
```

2. The Clientless VPN portal parses the incoming HTTP request from the browser and sets the Accept-Encoding header value to **gzip** that indicates support for Gzip encoding, as shown in the following example.

```

Hypertext Transfer Protocol
[Expert Info (Warning/Security): Unencrypted HTTP protocol detected over encrypted port, could indicate a dangerous misconfiguration.]
GET / HTTP/1.1\r\n
[Expert Info (Chat/Sequence): GET / HTTP/1.1\r\n
Request Method: GET
Request URI: /
Request Version: HTTP/1.1
Host: www.facebook.com\r\n
Connection: keep-alive\r\n
Upgrade-Insecure-Requests: 1\r\n
User-Agent: Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/91.0.4472.124 Safari/537.36\r\n
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,image/avif,image/webp,image/apng,*/*;q=0.8,application/signed-exchange;v=b3;q=0.9\r\n
Sec-Fetch-Site: same-origin\r\n
Sec-Fetch-Mode: navigate\r\n
Sec-Fetch-User: ?1\r\n
Sec-Fetch-Dest: document\r\n
sec-ch-ua: " Not;A Brand";v="99", "Google Chrome";v="91", "Chromium";v="91"\r\n
sec-ch-ua-mobile: ?0\r\n
Accept-Language: en-US,en;q=0.9\r\n
Accept-Encoding: gzip\r\n
\r\n
    
```

3. If the website supports Gzip encoding in the HTTP response, the website sends the Content-Encoding header as **gzip** that indicates the content is in Gzip format, as shown in the following example.

```

Hypertext Transfer Protocol
[Expert Info (Warning/Security): Unencrypted HTTP protocol detected over encrypted port, could indicate a dangerous misconfiguration.]
HTTP/1.1 200 OK\r\n
[Expert Info (Chat/Sequence): HTTP/1.1 200 OK\r\n
Response Version: HTTP/1.1
Status Code: 200
[Status Code Description: OK]
Response Phrase: OK
Vary: Accept-Encoding\r\n
Content-Encoding: gzip\r\n
Set-Cookie: fr=1dHw7x4ENvtCPg..Bg7xJL96.AAA.0.0.Bg7xJL.AMqwa2dRRs; expires=Tue, 12-Oct-2021 16:35:22 GMT; Max-Age=775999; path=/; domain=.facebook.com; secure; httpOnly; SameSite=None\r\n
Set-Cookie: sb=5xLV2a2a8I85y_105VHfpg; expires=Fri, 14-Jul-2023 16:35:23 GMT; Max-Age=63072000; path=/; domain=.facebook.com; secure; httpOnly; SameSite=None\r\n
report-to: {"max_age":2592000,"endpoints":[{"url":"https://www.facebook.com/browser_reporting/"}],"group":"coep_report","include_subdomains":true}\r\n
x-fb-liftr: 0\r\n
cross-origin-opener-policy-report-only: same-origin-allow-popups;report-to="coep_report"\r\n
Pragma: no-cache\r\n
Cache-Control: private, no-cache, no-store, must-revalidate\r\n
Expires: Sat, 01 Jan 2000 00:00:00 GMT\r\n
X-Content-Type-Options: nosniff\r\n
X-XSS-Protection: 0\r\n
[truncated]content-security-policy: default-src * data: blob: 'self';script-src *.facebook.com *.fbcdn.net *.facebook.net *.google-analytics.com *.google.com 127.0.0.1; 'unsafe-inline' 'unsafe-eval' blob: 'self';style-src data: b\r\n
X-Frame-Options: DENY\r\n
Strict-Transport-Security: max-age=15552000; preload\r\n
Content-Type: text/html; charset=utf-8\r\n
X-FB-Debug: P087L5a9Z9HHD3415Jms59677TK16x8Roap/218IDKPN2309Acuw5zJmsC0TIUGBYwCC02u3g8mcg=r\n
Date: Wed, 14 Jul 2021 16:35:23 GMT\r\n
Alt-Svc: h3="29";443"; ma=3600,h3-27="443"; ma=3600\r\n
Transfer-Encoding: chunked\r\n
Connection: keep-alive\r\n
\r\n
HTTP response 1/1
[Time since request: 0.542151000 seconds]
[Request in frame: 3]
[Request URI: http://www.facebook.com/]
HTTP chunked response
Content-Encoding: gzip body (gzip): 40315 bytes => 203002 bytes
    
```

4. The Clientless VPN forwards the response received from the website to the web browser in the same format, as shown in the following example.

```

Hypertext Transfer Protocol
[Expert Info (Warning/Security): Unencrypted HTTP protocol detected over encrypted port, could indicate a dangerous misconfiguration.]
HTTP/1.1 200 OK\r\n
[Expert Info (Chat/Sequence): HTTP/1.1 200 OK\r\n
Response Version: HTTP/1.1
Status Code: 200
[Status Code Description: OK]
Response Phrase: OK
Vary: Accept-Encoding\r\n
Content-Encoding: gzip\r\n
Set-Cookie: frsld0h7vAEDvvtCPg..Bg7x3L.96.AAA.0.0.Bg7x3L.AM0waZKR8; expires=Tue, 12-Oct-2021 16:35:22 GMT; Max-Age=7775999; path=/; domain=facebook.com; secure; httponly; SameSite=None\r\n
Set-Cookie: sb=SxLVYA2qaB1BSy_t85VHFpg; expires=Fri, 14-Jul-2023 16:35:23 GMT; Max-Age=63072000; path=/; domain=facebook.com; secure; httponly; SameSite=None\r\n
report-to: {"max_age":2592000,"endpoints":[{"url":"https://www.facebook.com/browser_reporting/"}],"group":"coep_report","include_subdomains":true}\r\n
x-fb-lafrr: 0\r\n
cross-origin-opener-policy-report-only: same-origin-allow-popups;report-to="coep_report"\r\n
Pragma: no-cache\r\n
Cache-Control: private, no-cache, no-store, must-revalidate\r\n
Expires: Sat, 01 Jun 2000 00:00:00 GMT\r\n
X-Content-Type-Options: nosniff\r\n
X-XSS-Protection: 0\r\n
[truncated]content-security-policy: default-src * data: blob: 'self';script-src *.facebook.com *.fbcdn.net *.facebook.net *.google-analytics.com *.google.com 127.0.0.1; 'unsafe-inline' 'unsafe-eval' blob: data: 'self';style-src data: b
X-Frame-Options: DENY\r\n
Strict-Transport-Security: max-age=15552000; preload\r\n
Content-Type: text/html; charset=utf-8\r\n
X-FB-Debug: PROFL5dyZyW4H034X55/mo159077K16x8Roap/2181DXPM2309Aocuw5zJmsC0TIUG8YwCC0Q2U3g8mcg=-\r\n
Date: Wed, 14 Jul 2021 16:35:23 GMT\r\n
Alt-Svc: h3-29=":443"; ma=3600,h3-27=":443"; ma=3600\r\n
Transfer-Encoding: chunked\r\n
Connection: keep-alive\r\n
\r\n
[HTTP response 1/1]
HTTP chunked response
Content-encoded entity body (gzip): 40915 bytes -> 209002 bytes
File Data: 209002 bytes

```

If the HTTP request received by the Clientless VPN does not include **gzip** as one of the encoding methods, the Clientless VPN does not accept Gzip encoding either.

Virtualization Features

- [DPDK Support for Different NIC Types](#)
- [CN-Series Firewall as a k8s Service](#)
- [Intelligent Traffic Offload Service for VM-Series on KVM](#)
- [Customize Dataplane Cores](#)

DPDK Support for Different NIC Types

VM-Series firewalls now support multiple NIC types and multiple queues. You can configure both SR-IOV and DPDK for all hypervisors on cloud platforms that support multiple NIC types. In addition, a single NIC type with variable queues (available on some cloud platforms) is also supported.

For example, you can now use both SR-IOV and DPDK with the Cisco ENCS uCPE appliance, which has both IGB (WAN) and IXGBE (LAN) drivers.

Each NIC can support a different number of queues. PAN-OS queries each DPDK interface to find the number of queues, calculates the total number of queues, then assigns them to the available data plane cores in a round-robin fashion.

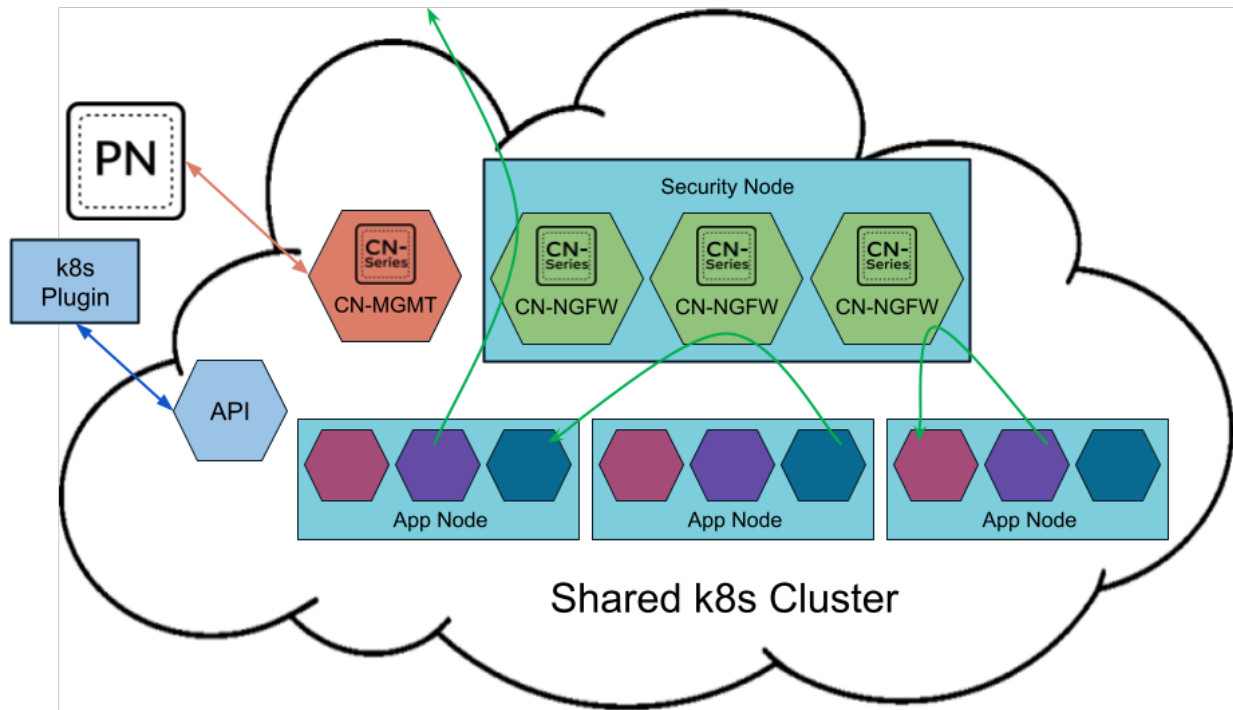
Please contact Technical Support if you want to configure this feature.

CN-Series Firewall as a k8s Service

You can now deploy the Palo Alto Networks Container Native Firewalls (CN-Series) as a service in your Kubernetes environment. By deploying the CN-Series firewall as a service, you are no longer required to deploy a CN-NGFW instance on each node in your environment. Instead, you can deploy the CN-Series anywhere in your cluster and any traffic in your environment is redirected to the CN-NGFW pods.

This is a new deployment mode for the CN-Series firewall that augments the previously released CN-Series-as-a-DaemonSet deployment mode.

The CN-Series firewall as a service requires Kubernetes 1.18 or later and kernel version 4.18 or later.



The CN-Series firewall as a service also supported the horizontal pod autoscaler. The HPA is a Kubernetes resource available in all cloud environments that automatically scales the number of CN-MGMT and CN-NGFW pods in a deployment based on monitored metrics. HPA uses two standard metrics across all cloud environments—CPU and memory utilization—as well as custom metrics specific to each cloud environment. As such, each cloud requires specific yaml files to enable HPA in AKS, EKS, and GKE.

HPA uses a cloud-specific metric adapter to retrieve metrics data from a monitoring adapter in the cloud environment, such as CloudWatch in EKS, to determine when to scale up or down based on the thresholds you define. You must modify the necessary yaml files to set the minimum and maximum number of replicas, the thresholds for each metric, and which metric are used in autoscaling your firewalls.

Cloud Environment	Metrics
AKS, EKS, and GKE	<p>CN-MGMT metrics</p> <ul style="list-style-type: none">• pan-logging-rate• pan-dataplane-slots <p>CN-NFGW metrics</p> <ul style="list-style-type: none">• dataplane-cpu-utilization-pct• dataplane-packet-buffer-utilization• pan-session-active• pan-session-utilization• pan-session-ssl-proxy-utilization• pan-throughput• pan-packet-rate• pan-connections-per-second

Intelligent Traffic Offload Service for VM-Series on KVM

With the new Intelligent Traffic Offload (ITO) service, VM-Series virtual NGFWs eliminate the tradeoff between network performance, security, and cost. The ITO service integrates with the industry's leading SmartNICs to improve virtual firewall performance by 5X by offloading traffic that does not benefit from security inspection from the firewall to the [BlueField-2 DPU](#).

For each new flow on the network, the ITO Service determines whether or not the flow can benefit from security inspection. The first few packets of the flow are routed to the firewall for inspection by the ITO service, which determines whether the rest of the packets in the flow should be inspected or offloaded. This determination is based on policy or on the flow's inability to be inspected (for example, encrypted traffic can't be inspected). By only inspecting flows that can benefit from security inspection, the overall load on the firewall is greatly reduced and performance increases without sacrificing the security posture.

The VM-Series firewall and the [BlueField-2 DPU](#) must be installed on an x86 physical host running Ubuntu 18.04, with kernel version 4.15.0-20. The VM-Series firewall must be deployed in [virtual wire](#) mode.

ITO benefits service provider networks where traffic is predominantly "elephant" flows. Elephant flows are typically media flows that do not benefit from advanced security inspection (YouTube streams, Zoom sessions, NetFlix streams, gaming traffic, etcetera), or encrypted SSL or IPsec flows without a corresponding decryption profile on the firewall.

The VM-Series firewall uses an [open API interface](#) based on [gRPC](#) to communicate with the BlueField-2 DPU, which handles offload processing and maintains the offload flow table.

The current BlueField-2 DPU scalability limitations are as follows:

- Session table capacity: 500,000 sessions
- Session table update rate: 7000 sessions/second
- Offload hairpin rate: ~90 Gbps for 1500 byte packets

Active/Passive HA is supported for the VM-Series firewalls running on physical hosts with identical configurations.

Customize Dataplane Cores

When a firewall is deployed with [Software NGFW Credits](#), the memory profile and the total number of vCPUs determine how many cores are automatically assigned to the management plane and the dataplane. The default configurations perform well in most cases.

Customize dataplane cores is an optional feature that allows you to customize the number of dataplane cores in two ways:

- During the initial deployment, use the `init-cfg.txt` file bootstrap parameter `plugin-op-commands=set-dp-cores:<#-cores>`. See [init-cfg.txt File Components](#).
- From a deployed firewall, use the VM-Series CLI command `request plugins vm_series dp-cores <#-cores>`. This procedure is the task below.

Typically you increase the number of dataplane cores (which decreases the number of management plane cores) to improve performance. Dataplane core customization does not require a change to the deployment profile or additional credits because the total number of vCPUs remains the same.

- Dataplane core customization is supported on firewalls licensed with a Software NGFW credit pool for 10.0.4 and above, and running PAN-OS 10.1 or later.
- Dataplane core customization is not supported for:
 - NSX-T
 - Intelligent Traffic Offload

Follow these steps to customize the dataplane cores on the VM-Series firewall.

STEP 1 | Log in to the VM-Series firewall and view the number of cores.

```
admin@PA-VM(active)>show plugins vm_series dp-cores
Device current DP cores: 13 (Total cores: 18)
```

STEP 2 | Change the number of dataplane cores.



Note that you must have at least one management plane core, and having too few cores also affects performance.

In this example we increase the dataplanes to 14.

```
admin@PA-VM(active)>request plugins vm_series dp-cores 14
Device current DP cores: 14 (Total cores: 18)
```

STEP 3 | Reboot the VM-Series firewall.

Select **Device > Setup > Operations** and **Reboot Device**.

STEP 4 | Use `show plugins vm_series dp-cores` to verify that the number of DP cores has changed.

Mobile Infrastructure Security Features

- [5G Multi-access Edge Computing Security](#)

5G Multi-access Edge Computing Security

5G Multi-access Edge Computing Security provides granular visibility and control for Packet Forwarding Control Protocol (PFCP) traffic by extracting information (such as subscriber ID or equipment ID) at the mobile edge, so you can apply security policy by subscriber, by equipment, or by network slice. It provides security at the protocol level through stateful inspection for PFCP traffic on 4G/LTE and 5G networks, in addition to reduced latency and higher bandwidth.

By providing context-based visibility into threats, 5G Multi-access Edge Computing Security protects networks from potential threats such as vulnerabilities, malware, and viruses. It secures devices and user that connect to multi-access edge computing (MEC), as well as applications hosted on MEC from attacks such as Denial of Service (DoS) and spoofing.

The following platforms support 5G Multi-Edge Security:

- PA-5200 Series
- VM-300 Series
- VM-500 Series
- VM-700 Series
- CN-Series on [OpenShift](#)

The following log events have been added for 5G Multi-access Edge Computing Security:

5G Multi-Edge Security Log Events

PFCP session message not matching existing PFCP association

PFCP association message sequence number mismatch

PFCP session message sequence number mismatch

PFCP association message is out of order

PFCP session message is out of order

PFCP association message

PFCP session message

PFCP association start

PFCP association end

PFCP session start

PFCP session end

STEP 1 | Enable GTP Security, commit your changes, and reboot.



If you enable stateful inspection for PFCP traffic, the following options are not available:

- IMSI/APN/RAT filtering
- GTP-U tunnel limiting
- GTPv1-C stateful inspection
- GTPv2-C stateful inspection
- 5G-HTTP2 for 5G-C
- End User IP Address Spoofing for GTP-U

Similarly, if you enable GTPv1-C stateful inspection, GTPv2-C stateful inspection, or 5G-HTTP2 for 5G-C, PFCP stateful inspection is not available.

STEP 2 | Create a Mobile Network Protection Profile and enable 5G Multi-access Edge Computing Security.

1. Select **Objects > Security Profiles > Mobile Network Protection**.
2. **Add** a profile and enter a **Name**, such as **5G Multi-access Edge Computing Security**.
3. On the **PFCP** tab, enable **Stateful Inspection**.

Mobile Network Protection Profile ?

Name

Description

GTP Inspection | Filtering Options | GTP Tunnel Limit | Overbilling Protection | Other Log Settings

GTP-C | GTP-U | 5G-C | **PFCP**

Stateful Inspection

Check Association Messages

Check Session Messages

Check Sequence Number

Log at PFCP association start

Log at PFCP association end

Log at PFCP session start

Log at PFCP session end

OK
Cancel

STEP 3 | Select which state checks you want the firewall to perform on PFCP traffic and the action you want the firewall to take if a state check is not successful.

1. Determine the state checks you want to use.
 - **Check Association Messages**—Checks for any PFCP association messages that are out of order or that have been rejected.
 - **Check Session Messages**—Checks for any PFCP session messages that are out of order or that have been rejected.
 - **Check Sequence Number**—Confirms that the sequence number in the PFCP matches the sequence number in the PFCP request message.
2. Select the action you want the firewall to take if a state check is not successful.
 - **allow**—Allow the traffic and do not generate a log entry in the GTP log.
 - **block**—Block the traffic and generate a high-severity log entry in the GTP log.
 - **alert**—(Default) Allow the traffic and generate a high-severity log entry in the GTP log.

STEP 4 | (Optional) Configure logging for PFCP traffic.

1. Select when you want the firewall to generate a log entry.
 - **Log at PFCP association start**
 - **Log at PFCP association end**
 - **Log at PFCP session start**
 - **Log at PFCP session end**
2. On the **Other Log Settings** tab, select the type of **PFCP Allowed Messages** you want to include in the logs.
 - **Session Establishment**—These PFCP messages set up the session, including establishing the GTP-U tunnel.
 - **Session Modification**—These PFCP messages are sent if the session ID or PDR ID changes (for example, as a result of moving from a 4G to a 5G network). It includes

messages such as PFCP Session Modification Request and PFCP Session Modification Response.

- **Session Deletion**—These PFCP messages terminate the PFCP session, including releasing associated resources.

Mobile Network Protection Profile

Name: 5G Multi-Edge Security

Description:

GTP Inspection | Filtering Options | GTP Tunnel Limit | Overbilling Protection | **Other Log Settings**

GTPv1-C Allowed Messages

- Tunnel Management
- Path Management
- Others

GTPv2-C Allowed Messages

- Tunnel Management
- Path Management
- Others

GTP-U Allowed Messages

- Tunnel Management
- Path Management
- G-PDU

G-PDU per New GTP-U Tunnel: 1

5G Allowed Messages

- N11

PFCP Allowed Messages

- Session Establishment
- Session Modification
- Session Deletion

Log User Location

Packet Capture

OK Cancel

STEP 5 | Click **OK**.

STEP 6 | Create a **Security policy rule** that applies your Mobile Network Protection profile to PFCP traffic by selecting the **Mobile Network Protection** profile you created.

STEP 7 | Create another Security policy rule based on **equipment ID**, **subscriber ID**, or **network slice**.

STEP 8 | **Commit** your changes.

