paloalto TECHDOCS

Panorama Administrator's Guide

Version 10.2



docs.paloaltonetworks.com

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

© 2021-2024 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

November 27, 2024

Table of Contents

Panorama Overview 11

- About Panorama 12
- Panorama Models 14
- Centralized Firewall Configuration and Update Management 17
  - Context Switch—Firewall or Panorama 17
  - Total Configuration Size for Panorama 18
  - Templates and Template Stacks 18
  - Device Groups 20
- Centralized Logging and Reporting 26
  - Managed Collectors and Collector Groups 26
  - Local and Distributed Log Collection 27
  - Caveats for a Collector Group with Multiple Log Collectors 28
  - Log Forwarding Options 30
  - Centralized Reporting 31
- Data Redistribution Using Panorama 33
- Role-Based Access Control 34
  - Administrative Roles 34
  - Authentication Profiles and Sequences 36
  - Access Domains 36
  - Administrative Authentication 37
- Panorama Commit, Validation, and Preview Operations 39
- Plan Your Panorama Deployment 41
- Deploy Panorama: Task Overview 43

Set Up Panorama 45

- Determine Panorama Log Storage Requirements 46
- Manage Large-Scale Firewall Deployments 48
  - Determine the Optimal Large-Scale Firewall Deployment Solution 48
  - Increased Device Management Capacity for M-Series and Panorama Virtual Appliance 48
- Set Up the Panorama Virtual Appliance 52
  - Setup Prerequisites for the Panorama Virtual Appliance 52
  - Install the Panorama Virtual Appliance 57
  - Perform Initial Configuration of the Panorama Virtual Appliance 119
  - Set Up The Panorama Virtual Appliance as a Log Collector 122
  - Set Up the Panorama Virtual Appliance with Local Log Collector 131
  - Set up a Panorama Virtual Appliance in Panorama Mode 136
  - Set up a Panorama Virtual Appliance in Management Only Mode 137
  - Expand Log Storage Capacity on the Panorama Virtual Appliance 138

 ©2025 Palo Alto Networks, Inc.

Increase CPUs and Memory on the Panorama Virtual Appliance....166
Increase the System Disk on the Panorama Virtual Appliance....172
Complete the Panorama Virtual Appliance Setup....178
Convert Your Panorama Virtual Appliance....178
Set Up the M-Series Appliance....190
M-Series Appliance Interfaces....190
Perform Initial Configuration of the M-Series Appliance....192
Perform Initial Configuration of an Air Gapped M-Series Appliance....197
M-Series Setup Overview....201
Set Up the M-Series Appliance as a Log Collector....203
Increase Storage on the M-Series Appliance....214
Configure Panorama to Use Multiple Interfaces....220
Register Panorama and Install Licenses....228
Register Panorama....228
Activate a Panorama Support License....230
Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected....231
Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected....232
Activate/Retrieve a Firewall Management License on the M-Series Appliance....234
Install the Panorama Device Certificate....237
Install the Device Certificate for a Dedicated Log Collector....240
Transition to a Different Panorama Model....243
Migrate from a Panorama Virtual Appliance to an M-Series Appliance....243
Migrate a Panorama Virtual Appliance to a Different Hypervisor....247
Migrate from an M-Series Appliance to a Panorama Virtual Appliance....253
Migrate from an M-500 Appliance to an M-700 Appliance....259
Migrate from an M-600 Appliance to an M-700 Appliance....266
Migrate from an M-100 Appliance to an M-500 Appliance....271
Migrate from an M-100 or M-500 Appliance to an M-200 or M-600 Appliance....274
Access and Navigate Panorama Management Interfaces....279
Log in to the Panorama Web Interface....279
Navigate the Panorama Web Interface....280
Log in to the Panorama CLI....280
Set Up Administrative Access to Panorama....282
Configure an Admin Role Profile....282
Configure an Admin Role Profile for Selective Push to Managed Firewalls....283
Configure an Access Domain....284

 ©2025 Palo Alto Networks, Inc.

Configure Administrative Accounts and Authentication 285
Configure Tracking of Administrator Activity 298
Set Up Authentication Using Custom Certificates 301
How Are SSL/TLS Connections Mutually Authenticated? 301
Configure Authentication Using Custom Certificates on Panorama 302
Configure Authentication Using Custom Certificates on Managed Devices 304
Add New Client Devices 306
Change Certificates 306

Manage Firewalls 309

Add a Firewall as a Managed Device 310
Install the Device Certificate for Managed Firewalls 319
Install the Device Certificate for a Managed Firewall 319
Install the Device Certificate for All Managed Firewalls Without a Device Certificate 323
Set Up Zero Touch Provisioning 328
ZTP Overview 328
Install the ZTP Plugin 330
Configure the ZTP Installer Administrator Account 336
Add ZTP Firewalls to Panorama 337
Configure ZTP for Firewalls in HA 344
Use the CLI for ZTP Tasks 347
Uninstall the ZTP Plugin 350
Manage Device Groups 351
Add a Device Group 351
Create a Device Group Hierarchy 352
Create Objects for Use in Shared or Device Group Policy 354
Revert to Inherited Object Values 356
Manage Unused Shared Objects 356
Manage Precedence of Inherited Objects 357
Move or Clone a Policy Rule or Object to a Different Device Group 358
Push a Policy Rule to a Subset of Firewalls 359
Device Group Push to a Multi-VSYS Firewall 362
Manage the Rule Hierarchy 364
Manage Templates and Template Stacks 366
Template Capabilities and Exceptions 366
Add a Template 366
Configure a Template Stack 369
Configure a Template or Template Stack Variable 373
Import and Overwrite Existing Template Stack Variables 376

 ©2025 Palo Alto Networks, Inc.

Override a Template or Template Stack Value....378
Disable/Remove Template Settings....380
Manage the Master Key from Panorama....382
Schedule a Configuration Push to Managed Firewalls....388
Redistribute Data to Managed Firewalls....391
Transition a Firewall to Panorama Management....394
Plan the Transition to Panorama Management....394
Migrate a Firewall to Panorama Management and Reuse Existing Configuration....395
Migrate a Firewall to Panorama Management and Push a New Configuration....400
Migrate a Firewall HA Pair to Panorama Management and Reuse Existing Configuration....402
Migrate a Firewall HA Pair to Panorama Management and Push a New Configuration....407
Load a Partial Firewall Configuration into Panorama....410
Localize a Panorama Pushed Configuration on a Managed Firewall....412
Device Monitoring on Panorama....415
Monitor Device Health....415
Monitor Policy Rule Usage....417
Use Case: Configure Firewalls Using Panorama....423
Device Groups in this Use Case....423
Templates in this Use Case....424
Set Up Your Centralized Configuration and Policies....425

Manage Log Collection....433

Configure a Managed Collector....434
Monitor Managed Collector Health Status....442
Configure Authentication for a Dedicated Log Collector....443
Configure an Administrative Account for a Dedicated Log Collector....443
Configure RADIUS Authentication for a Dedicated Log Collector....445
Configure TACACS+ Authentication for a Dedicated Log Collector....449
Configure LDAP Authentication for a Dedicated Log Collector....452
Manage Collector Groups....457
Configure a Collector Group....457
Configure Authentication with Custom Certificates Between Log Collectors....460
Move a Log Collector to a Different Collector Group....463
Remove a Firewall from a Collector Group....464
Configure Log Forwarding to Panorama....465
Configure Syslog Forwarding to External Destinations....470
Forward Logs to Strata Logging Service....474

 ©2025 Palo Alto Networks, Inc.

Verify Log Forwarding to Panorama....475
Modify Log Forwarding and Buffering Defaults....477
Configure Log Forwarding from Panorama to External Destinations....479
Log Collection Deployments....482
Deploy Panorama with Dedicated Log Collectors....482
Deploy Panorama M-Series Appliances with Local Log Collectors....491
Deploy Panorama Virtual Appliances with Local Log Collectors....497
Deploy Panorama Virtual Appliances in Legacy Mode with Local Log Collection....502

Manage WildFire Appliances....505

Add Standalone WildFire Appliances to Manage with Panorama....506
Configure Basic WildFire Appliance Settings on Panorama....511
Configure Authentication for a WildFire Appliance....511
Set Up Authentication Using Custom Certificates on WildFire Appliances and Clusters....525
Configure a Custom Certificate for a Panorama Managed WildFire Appliance....525
Configure Authentication with a Single Custom Certificate for a WildFire Cluster....527
Apply Custom Certificates on a WildFire Appliance Configured through Panorama....529
Remove a WildFire Appliance from Panorama Management....532
Manage WildFire Clusters....533
Configure a Cluster Centrally on Panorama....533
View WildFire Cluster Status Using Panorama....557

Manage Licenses and Updates....559

Manage Licenses on Firewalls Using Panorama....560

Monitor Network Activity....563

Use Panorama for Visibility....564
Monitor the Network with the ACC and AppScope....564
Analyze Log Data....566
Generate, Schedule, and Email Reports....567
Configure Key Limits for Scheduled Reports....570
Ingest Traps ESM Logs on Panorama....573
Use Case: Monitor Applications Using Panorama....575
Use Case: Respond to an Incident Using Panorama....578
Incident Notification....578
Review the Widgets in the ACC....579
Review Threat Logs....579

 ©2025 Palo Alto Networks, Inc.

Review WildFire Logs ... 580
Review Data Filtering Logs ... 580
Update Security Rules ... 581

Panorama High Availability ... 583

Panorama HA Prerequisites ... 584
Priority and Failover on Panorama in HA ... 586
Failover Triggers ... 588
HA Heartbeat Polling and Hello Messages ... 588
HA Path Monitoring ... 588
Logging Considerations in Panorama HA ... 590
Logging Failover on a Panorama Virtual Appliance in Legacy Mode ... 590
Logging Failover on an M-Series Appliance or Panorama Virtual Appliance in Panorama Mode ... 591
Synchronization Between Panorama HA Peers ... 592
Manage a Panorama HA Pair ... 593
Set Up HA on Panorama ... 593
Set Up Authentication Using Custom Certificates Between HA Peers ... 596
Test Panorama HA Failover ... 598
Switch Priority after Panorama Failover to Resume NFS Logging ... 598
Restore the Primary Panorama to the Active State ... 599

Administer Panorama ... 601

Preview, Validate, or Commit Configuration Changes ... 602
Commit Selective Configuration Changes for Managed Devices ... 606
Push Selective Configuration Changes to Managed Devices ... 608
Enable Automated Commit Recovery ... 611
Manage Panorama and Firewall Configuration Backups ... 613
Schedule Export of Configuration Files ... 613
Save and Export Panorama and Firewall Configurations ... 615
Revert Panorama Configuration Changes ... 617
Configure the Maximum Number of Configuration Backups on Panorama ... 620
Load a Configuration Backup on a Managed Firewall ... 621
Compare Changes in Panorama Configurations ... 622
Manage Locks for Restricting Configuration Changes ... 623
Add Custom Logos to Panorama ... 625
Use the Panorama Task Manager ... 626
Manage Storage Quotas and Expiration Periods for Logs and Reports ... 627
Log and Report Storage ... 627
Log and Report Expiration Periods ... 628
Configure Storage Quotas and Expiration Periods for Logs and Reports ... 628

 ©2025 Palo Alto Networks, Inc.

Configure the Run Time for Panorama Reports 631
Monitor Panorama 632
Panorama System and Configuration Logs 632
Monitor Panorama and Log Collector Statistics Using SNMP 633
Reboot or Shut Down Panorama 636
Configure Panorama Password Profiles and Complexity 637
Panorama Plugins 639
About Panorama Plugins 640
Install Panorama Plugins 642
VM-Series Plugin and Panorama Plugins 644
Install the VM-Series Plugin on Panorama 644
Endpoint Monitoring for Cisco TrustSec 646
Panorama Plugin for Cisco TrustSec 646
Install the Panorama Plugin for Cisco TrustSec 647
Configure the Panorama Plugin for Cisco TrustSec 648
Troubleshoot the Panorama Plugin for Cisco TrustSec 655
Troubleshooting 659
Troubleshoot Panorama System Issues 660
Generate Diagnostic Files for Panorama 660
Diagnose Panorama Suspended State 660
Monitor the File System Integrity Check 660
Manage Panorama Storage for Software and Content Updates 661
Recover from Split Brain in Panorama HA Deployments 662
Troubleshoot Log Storage and Connection Issues 663
Verify Panorama Port Usage 663
Resolve Zero Log Storage for a Collector Group 666
Replace a Failed Disk on an M-Series Appliance 666
Replace the Virtual Disk on an ESXi Server 666
Replace the Virtual Disk on vCloud Air 667
Migrate Logs to a New M-Series Appliance in Log Collector Mode 668
Migrate Logs to a New M-Series Appliance in Panorama Mode 674
Migrate Logs to a New M-Series Appliance Model in Panorama Mode in High Availability 682
Migrate Logs to the Same M-Series Appliance Model in Panorama Mode in High Availability 692
Migrate Log Collectors after Failure/RMA of Non-HA Panorama 701
Regenerate Metadata for M-Series Appliance RAID Pairs 705
View Log Query Jobs 706
Replace an RMA Firewall 708
Partial Device State Generation for Firewalls 708

 ©2025 Palo Alto Networks, Inc.

Before Starting RMA Firewall Replacement....708
Restore the Firewall Configuration after Replacement....710
Troubleshoot Commit Failures....715
Triage Commit Issues on Panorama....716
Troubleshoot Template or Device Group Push Failures....718
Troubleshoot Panorama Push Failure Due to Pending Local Firewall Changes....720
Troubleshoot Registration or Serial Number Errors....721
Troubleshoot Reporting Errors....722
Troubleshoot Device Management License Errors....723
Troubleshoot Automatically Reverted Firewall Configurations....724
View Task Success or Failure Status....726
Test Policy Match and Connectivity for Managed Devices....727
Troubleshoot Policy Rule Traffic Match....727
Troubleshoot Connectivity to Network Resources....728
Generate a Stats Dump File for a Managed Firewall....730
Recover Managed Device Connectivity to Panorama....732
Restore an Expired Device Certificate....735

 ©2025 Palo Alto Networks, Inc.

Panorama Overview

The Panorama™ management server provides centralized monitoring and management of multiple Palo Alto Networks next-generation firewalls and of WildFire appliances and appliance clusters. It provides a single location from which you can oversee all applications, users, and content traversing your network, and then use this knowledge to create application enablement policies that protect and control the network. Using Panorama for centralized policy and firewall management increases operational efficiency in managing and maintaining a distributed network of firewalls. Using Panorama for centralized WildFire appliance and WildFire appliance cluster management increases the number of firewalls a single network supports, provides high availability for fault tolerance, and increases management efficiency.

- About Panorama
- Panorama Models
- Centralized Firewall Configuration and Update Management
- Centralized Logging and Reporting
- Data Redistribution Using Panorama
- Role-Based Access Control
- Panorama Commit, Validation, and Preview Operations
- Plan Your Panorama Deployment
- Deploy Panorama: Task Overview

About Panorama

Panorama enables you to effectively configure, manage, and monitor your Palo Alto Networks firewalls with central oversight. The three main areas in which Panorama adds value are:

- **Centralized configuration and deployment**—To simplify central management and rapid deployment of the firewalls and WildFire appliances on your network, use Panorama to pre-stage the firewalls and WildFire appliances for deployment. You can then assemble the firewalls into groups, and create templates to apply a base network and device configuration and use device groups to administer globally shared and local policy rules. See Centralized Firewall Configuration and Update Management.
- **Aggregated logging with central oversight for analysis and reporting**—Collect information on activity across all the managed firewalls on the network and centrally analyze, investigate and report on the data. This comprehensive view of network traffic, user activity, and the associated risks empowers you to respond to potential threats using the rich set of policies to securely enable applications on your network. See Centralized Logging and Reporting.
- **Distributed administration**—Enables you to delegate or restrict access to global and local firewall configurations and policies. See Role-Based Access Control for delegating appropriate levels of access for distributed administration.

Six Panorama Models are available: the Panorama virtual appliance, M-600 appliance, M-500 appliance, and M-200 appliance are supported in PAN-OS 10.0. The M-300 appliance and M-700 appliance are supported in PAN-OS 10.2.Panorama Centralized Management illustrates how you can deploy Panorama in a high availability (HA) configuration to manage firewalls.

 ©2025 Palo Alto Networks, Inc.

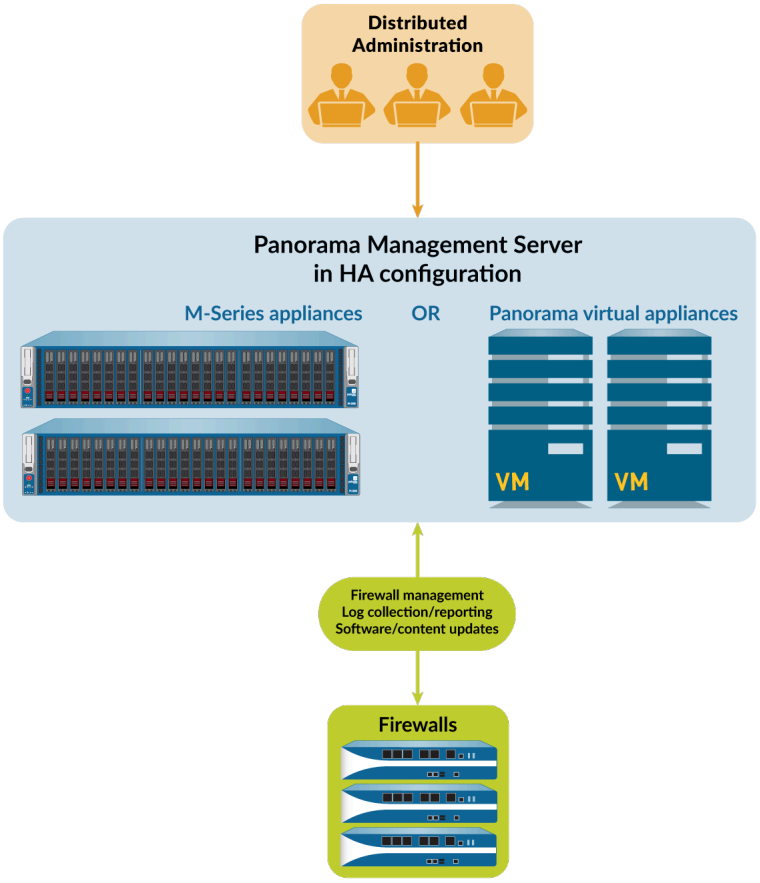


Figure 1: Panorama Centralized Management

 ©2025 Palo Alto Networks, Inc.

Panorama Models

Panorama is available as one of the following virtual or physical appliances, each of which supports licenses for managing up to 25, 100, or 1,000 firewalls. Additionally, M-600 and M-700 appliances support licenses for managing up to 5,000 firewalls and similarly resourced Panorama virtual appliances support licenses for managing up to 2,500 firewalls:

- **Panorama virtual appliance**—This model provides simple installation and facilitates server consolidation for sites that need a virtual management appliance. You can install Panorama on Alibaba Cloud, Amazon Web Services (AWS), AWS GovCloud, Microsoft Azure, Google Cloud Platform (GCP), KVM, Hyper-V, Oracle Cloud Infrastructure (OCI), a VMware ESXi server, or on VMware vCloud Air. The virtual appliance can collect firewall logs locally at rates of up to 20,000 logs per second and can manage Dedicated Log Collectors for higher logging rates. The virtual appliance can function as a dedicated management server, a Panorama management server with local log collection capabilities, or as a Dedicated Log Collector. For the supported interfaces, log storage capacity, and maximum log collection rates, see the Setup Prerequisites for the Panorama Virtual Appliance. You can deploy the virtual appliance in the following modes:
  - **Panorama mode**—In this mode, the Panorama virtual appliance supports a local Log Collector with 1 to 12 virtual logging disks (see Deploy Panorama Virtual Appliances with Local Log Collectors). Each logging disk has 2TB of storage capacity for a total maximum of 24TB on a single virtual appliance and 48TB on a high availability (HA) pair. Only Panorama mode enables you to add multiple virtual logging disks without losing logs on existing disks. Panorama mode also provides the benefit of faster report generation. In Panorama mode, the virtual appliance does not support NFS storage.
  - **Legacy mode** (ESXi and vCloud Air only)—In this mode, the Panorama virtual appliance receives and stores firewall logs without using a local Log Collector (see Deploy Panorama Virtual Appliances in Legacy Mode with Local Log Collection). By default, the virtual appliance in Legacy mode has one disk partition for all data. Approximately 11GB of the partition is allocated to log storage. If you need more local log storage, you can add one virtual disk of up to 8TB on ESXi 5.5 and later versions or on vCloud Air. Earlier ESXi versions support one virtual disk of up to 2TB. If you need more than 8TB, you can mount the virtual appliance in Legacy mode to an NFS datastore but only on the ESXi server, not in vCloud Air. This mode is only available if your Panorama virtual appliance is in Legacy mode on upgrade to PAN-OS 10.0. On upgrade to PAN-OS 9.0 and later releases, Legacy mode is no longer available if you change to any other mode. If you change your Panorama virtual appliance from Legacy mode to one of the available modes, you will no longer be able to change back into Legacy mode.

    As a best practice, deploy the virtual appliance in Panorama mode to optimize log storage and report generation.

    While supported, Legacy mode is not recommended for production environments but may still be used for lab or demo environments.
  - **Management Only mode**—In this mode, the Panorama virtual appliance is a dedicated management appliance for your managed devices and Dedicated Log Collectors. Additionally, an appropriately resourced Panorama virtual appliance can manage up to 2,500 firewalls in this mode. The Panorama virtual appliance has no log collection capabilities

 ©2025 Palo Alto Networks, Inc.

except for config and system logs and requires a Dedicated Log Collector to these store logs. By default, the virtual appliance in Management Only mode has only one disk partition for all data so all logs forwarded to a Panorama virtual appliance in Management Only mode are dropped. Therefore, to store the log data from your managed appliances, you must configure log forwarding in order to store the log data from your managed devices. For more information, see Increased Device Management Capacity Requirements.

- **Log Collector mode**—The Panorama virtual appliance functions as a Dedicated Log Collector. If multiple firewalls forward large volumes of log data, a Panorama virtual appliance in Log Collector mode provides increased scale and performance. In this mode, the appliance does not have a web interface for administrative access; it has only a command line interface (CLI). However, you can manage the appliance using the web interface of the Panorama management server. CLI access to a Panorama virtual appliance in Log Collector mode is necessary only for initial setup and debugging. For configuration details, see Deploy Panorama with Dedicated Log Collectors.

- **M-Series appliance**—The M-200, M-300, M-500, M-600, and M-700 appliances are dedicated hardware appliances intended for large-scale deployments. In environments with high logging rates (over 10,000 logs per second) and log retention requirements, these appliances enable scaling of your log collection infrastructure. For the supported interfaces, log storage capacity, and maximum log collection rates, see M-Series Appliance Interfaces. All M-Series models share the following attributes:
  - RAID drives to store firewall logs and RAID 1 mirroring to protect against disk failures
  - SSD to store the logs that Panorama and Log Collectors generate
  - MGT, Eth1, Eth2, and Eth3 interfaces that support 1Gbps throughput
  - Redundant, hot-swappable power supplies
  - front-to-back airflow

  The M-500 and M-600 appliances have the following additional attributes, which make them more suitable for data centers:

  - Eth4 and Eth5 interfaces that support 10Gbps throughput

  Additionally, the following attribute makes the M-600 and M-700 appliances more suitable for large-scale firewall deployments:

  - The M-600 and M-700 appliances in Management Only mode can manage up to 5,000 firewalls.

  You can deploy the M-Series appliances in the following modes:

  - **Panorama mode**—The appliance functions as a Panorama management server to manage firewalls and Dedicated Log Collectors. The appliance also supports a local Log Collector to aggregate firewall logs. Panorama mode is the default mode. For configuration details, see Deploy Panorama M-Series Appliances with Local Log Collectors.
  - **Management Only mode**—The Panorama appliance is a dedicated management appliance for your managed devices and Dedicated Log Collectors. The Panorama appliance has no log collection capabilities except for config and system logs and your deployment requires a Dedicated Log Collector to store these logs. By default, the Panorama appliance in Management Only mode has only one disk partition for all data so all logs forwarded to a Panorama virtual appliance in Management Only mode are dropped. Therefore, to store the

 ©2025 Palo Alto Networks, Inc.

log data from your managed appliances, you must configure log forwarding in order to store the log data from your managed devices.

- **Log Collector mode**—The appliance functions as a Dedicated Log Collector. If multiple firewalls forward large volumes of log data, an M-Series appliance in Log Collector mode provides increased scale and performance. In this mode, the appliance does not have a web interface for administrative access; it has only a command line interface (CLI). However, you can manage the appliance using the web interface of the Panorama management server. CLI access to an M-Series appliance in Log Collector mode is necessary only for initial setup and debugging. For configuration details, see Deploy Panorama with Dedicated Log Collectors.

For more details and specifications for the M-Series appliances, see the M-Series Appliance Hardware Reference Guides.

 ©2025 Palo Alto Networks, Inc.

Centralized Firewall Configuration and Update Management

Panorama™ uses *device groups* and *templates* to group firewalls into logical sets that require similar configuration. You use device groups and templates to centrally manage all configuration elements, policies, and objects on the managed firewalls. Panorama also enables you to centrally manage licenses, software (PAN-OS® software, SSL-VPN client software, GlobalProtect™ agent/app software), and content updates (Applications, Threats, WildFire®, and Antivirus). All device group, template, and template stack configuration objects are required to have a unique name.

In the event an unforeseen restart of your managed firewall or Panorama occurs, all uncommitted configuration changes in your device groups and templates are preserved locally until you successfully commit the changes. A restart can be the restart of the firewall or Panorama or of a PAN-OS management process related to configuration management. 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.

- Context Switch—Firewall or Panorama
- Total Configuration Size for Panorama
- Templates and Template Stacks
- Device Groups

Context Switch—Firewall or Panorama

The Panorama™ web interface enables you to toggle between a Panorama-centric view and a firewall-centric view using the **Context** drop-down at the top-left of every tab. Set the **Context** to **Panorama** to manage firewalls centrally or switch context to the web interface of a specific firewall to configure it locally. The similarity of the Panorama and firewall web interfaces enables you to seamlessly move between them to monitor and manage firewalls.

The **Context** drop-down lists only the firewalls that are connected to Panorama. For a Device Group and Template administrator, the drop-down lists only the connected firewalls that are within the Access Domains assigned to that administrator. To search a long list, use the Filters within the drop-down.

For firewalls in a high availability (HA) configuration, the icons have colored backgrounds to indicate the HA state (as follows). Knowing the HA state is useful when selecting a firewall context. For example, you generally make firewall-specific configuration changes on an active firewall.

- **Green**—Active.
- **Yellow**—Passive or the firewall is initiating (the initiating state lasts for up to 60 seconds after boot up).
- **Red**—The firewall is non-functional (error state), suspended (an administrator disabled the firewall), or tentative (for a link or path monitoring event in an active/active HA configuration).

When you configure an admin role profile for a Device Group and Template admin, you must assign a **Device Admin Role** that is pushed to your managed firewalls to context switch between the Panorama and firewall web interface.

 ©2025 Palo Alto Networks, Inc.

During the context switch, Panorama validates if the admin has access to a specific vsys or for all vsys. If the admin has access to all vsys, then Panorama uses the device admin role context switch. If the admin has access to one or some of the vsys, then Panorama uses the vsys admin role to context switch.

Total Configuration Size for Panorama

The total configuration file size of Panorama™ M-Series and virtual appliances is an important piece of the performance metric when determining which M-Series appliance or the minimum amount of virtual resources you need to allocate on your Panorama virtual appliance to ensure that you meet your Security requirements. Exceeding the supported total configuration file size of the Panorama management server results in reduced performance when performing configuration changes, commits, and pushes to managed firewalls.

The Panorama management server in Panorama mode supports a total configuration file size of 80MB for all template, device group, and Panorama-specific configurations. The maximum configuration file size supported by Panorama in Management Only mode depends on the Panorama model or resources you allocate to the Panorama virtual appliance. Refer to the table below for the recommended maximum configuration file size based on the Panorama M-Series appliance model or on the resources you allocate to the Panorama virtual appliance.

Panorama Model	Virtual Resources Required	Maximum Configuration File Size in Management Only Mode	Maximum Configuration File Size in Panorama Mode
M-200	N/A	120 MB	80 MB
M-300		150 MB	
M-500		120 MB	
M-600		150 MB	
M-700		180 MB	
Panorama Virtual Appliance Refer to the Setup Prerequisites for the Panorama Virtual Appliance for additional setup information.	• 16 vCPU • 128GB memory	120 MB	
	• 56 vCPU • 256GB memory	150 MB	

Templates and Template Stacks

You use templates and template stacks to configure the settings that enable firewalls to operate on the network. Templates are the basic building blocks you use to configure the

 ©2025 Palo Alto Networks, Inc.

Network and **Device** tabs on Panorama™. You can use templates to define interface and zone configurations, to manage the server profiles for logging and syslog access, or to define VPN configurations. Template stacks give you the ability to layer multiple templates and create a combined configuration. Template stacks simplify management because they allow you to define a common base configuration for all devices attached to the template stack and they give you the ability to layer templates to create a combined configuration. This enables you to define templates with location- or function-specific settings and then stack the templates in descending order of priority so that firewalls inherit the settings based on the order of the templates in the stack.

Both templates and template stacks support variables. Variables allow you to create placeholder objects with their value specified in the template or template stack based on your configuration needs. Create a template or template stack variable to replace IP addresses, Group IDs, and interfaces in your configurations. Template variables are inherited by the template stack and you can override them to create a template stack variable. However, templates do not inherit variables defined in the template stack. When a variable is defined in the template or template stack and pushed to the firewall, the value defined for the variable is displayed on the firewall.

Use templates to accommodate firewalls that have unique settings. Alternatively, you can push a broader, common base configuration and then override certain pushed settings with firewall-specific values on individual firewalls. When you override a setting on the firewall, the firewall saves that setting to its local configuration and Panorama no longer manages the setting. To restore template values after you override them, use Panorama to force the template or template stack configuration onto the firewall. For example, after you define a common NTP server in a template and override the NTP server configuration on a firewall to accommodate a local time zone, you can later revert to the NTP server defined in the template.

When defining a template stack, consider assigning firewalls that are the same hardware model and require access to similar network resources, such as gateways and syslog servers. This enables you to avoid the redundancy of adding every setting to every template stack. The following figure illustrates an example configuration in which you assign data center firewalls in the Asia-Pacific (APAC) region to a stack 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 can then re-use the global and APAC-specific templates by adding them to another stack that includes a template with branch-specific settings. Templates in a stack have a configurable priority order that ensures Panorama pushes only one value for any duplicate setting. Panorama evaluates the templates listed in a stack configuration from top to bottom with higher templates having priority. The following figure illustrates a data center stack in which the data center template has a higher priority than the global template: Panorama pushes the idle timeout value from the data center template and ignores the value from the global template.

 ©2025 Palo Alto Networks, Inc.

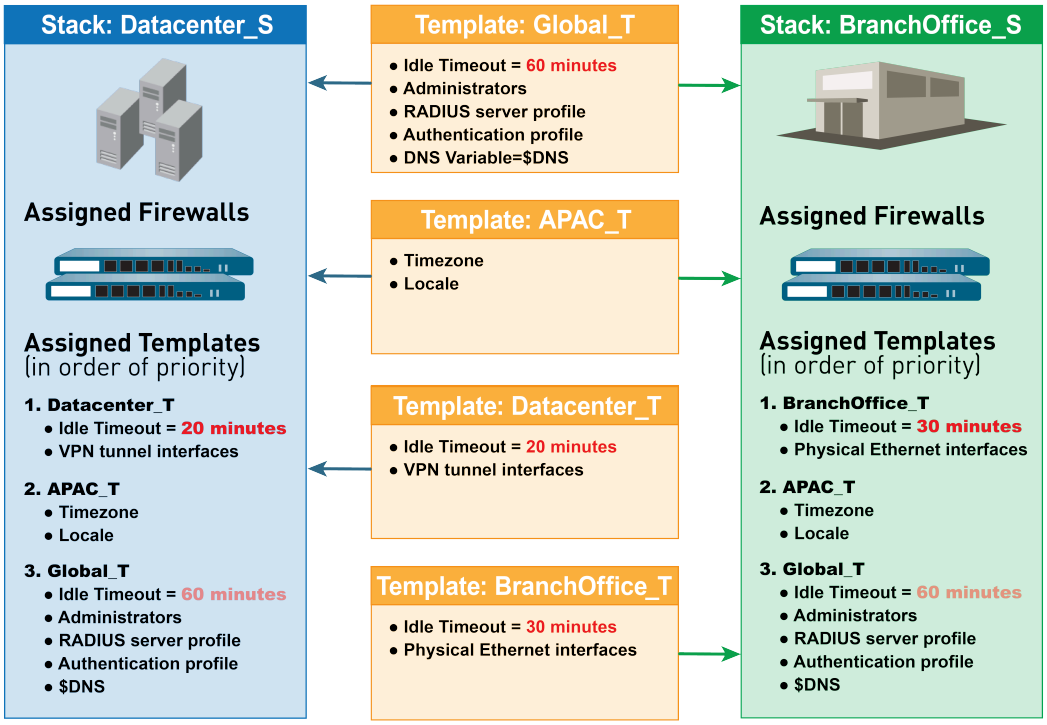


Figure 2: Template Stacks

You cannot use templates or template stacks to set firewall modes: virtual private network (VPN) mode, multiple virtual systems (multi-vsys) mode, or operational modes (normal or FIPS-CC mode). For details, see Template Capabilities and Exceptions. However, you can assign firewalls that have non-matching modes to the same template or stack. In such cases, 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 vsys in a template to firewalls that don't support virtual systems or that don't have any virtual systems configured.

For the relevant procedures, see Manage Templates and Template Stacks.

Device Groups

To use Panorama effectively, you have to group the firewalls in your network into logical units called *device groups*. A device group enables grouping based on network segmentation, geographic location, organizational function, or any other common aspect of firewalls that require similar policy configurations. Using device groups, you can configure policy rules and the objects they reference. You can organize device group hierarchically, with shared rules and objects at the top, and device group-specific rules and objects at subsequent levels. This enables you to create a hierarchy of rules that enforce how firewalls handle traffic. For example, you can define a set of shared rules as a corporate acceptable use policy. Then, to allow only regional offices to access peer-to-peer traffic such as BitTorrent, you can define a device group rule that Panorama pushes only to the regional offices (or define a shared security rule and target it to the regional offices). For the relevant procedures, see Manage Device Groups. The following topics describe device group concepts and components in more detail:

- Device Group Hierarchy
- Device Group Policies

 ©2025 Palo Alto Networks, Inc.

- Device Group Objects

Device Group Hierarchy

You can Create a Device Group Hierarchy to nest device groups in a tree hierarchy of up to four levels, with lower-level groups inheriting the settings (policy rules and objects) of higher-level groups. At the bottom level, a device group can have parent, grandparent, and great-grandparent device groups (*ancestors*). At the top level, a device group can have child, grandchild, and great-grandchild device groups (*descendants*). All device groups inheriting settings from the *Shared* location—a container at the top of the hierarchy for configurations that are common to all device groups.

Creating a device group hierarchy enables you to organize firewalls based on common policy requirements 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. Without a hierarchy, you would have to configure both function- and location-specific settings for every device group in a single level under Shared.

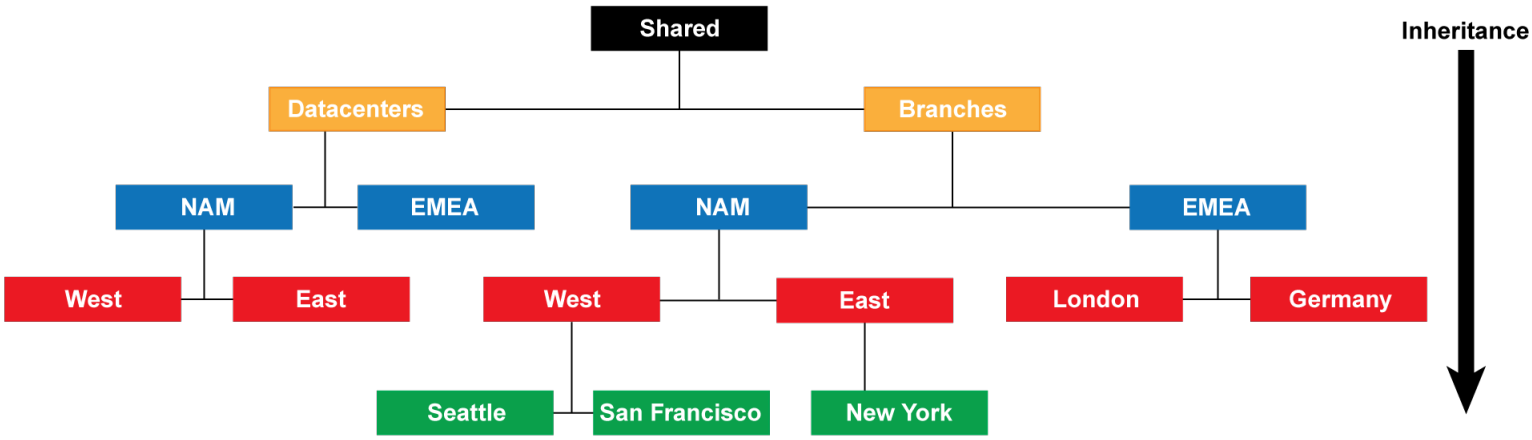
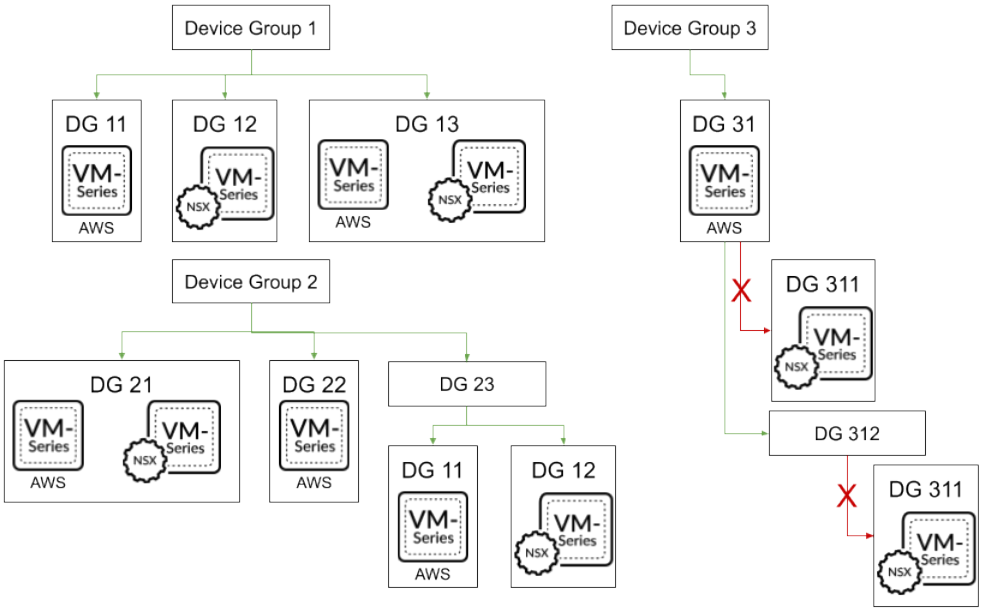


Figure 3: Device Group Hierarchy

For details on the order in which firewalls evaluate policy rules in a device group hierarchy, see Device Group Policies. For details on overriding the values of objects that device groups inherit from ancestor device groups, see Device Group Objects.

In a multiple Panorama plugin deployment to perform, a device group containing firewalls deployed in a particular hypervisor cannot be the child or parent of a device group containing firewalls deployed in a different hypervisor. For example, if Panorama receives IP address updates from VMware NSX-V and AWS, you cannot create a device group of NSX-V VM-Series firewalls that is a child of an AWS VM-Series firewall device group.

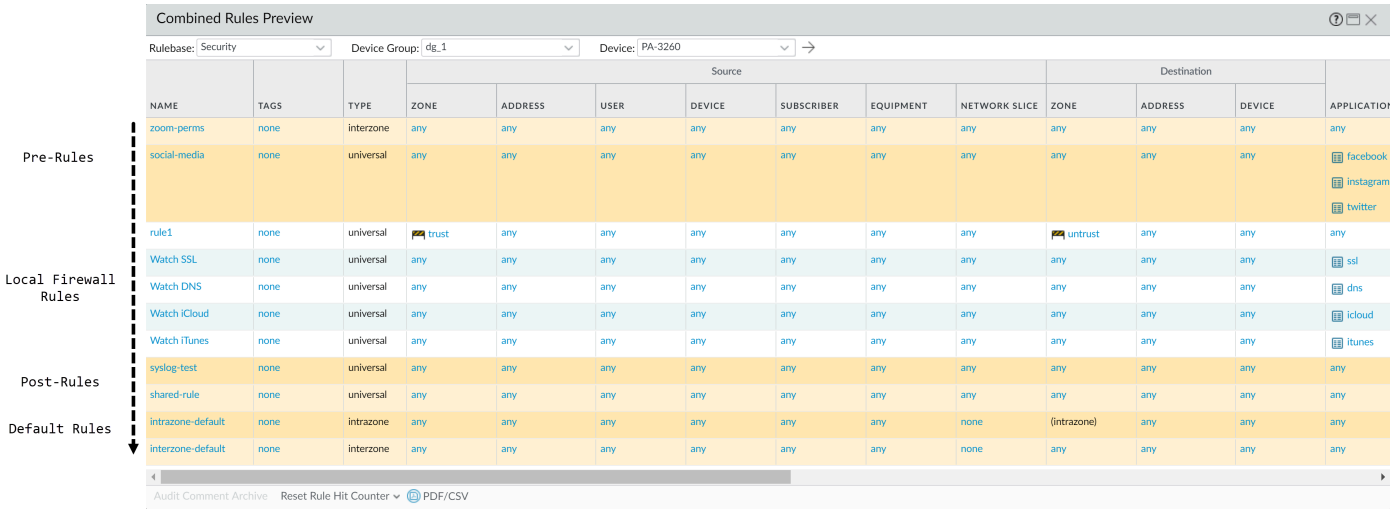
 ©2025 Palo Alto Networks, Inc.



Device Group Policies

Device groups provide a way to implement a layered approach for managing policies across a network of managed firewalls. A firewall evaluates policy rules by layer (shared, device group, and local) and by type (pre-rules, post-rules, and default rules) in the following order from top to bottom. When the firewall receives traffic, it performs the action defined in the first evaluated rule that matches the traffic and disregards all subsequent rules. To change the evaluation order for rules within a particular layer, type, and rulebase (for example, shared Security pre-rules), see Manage the Rule Hierarchy.

Whether you view rules on a firewall or in Panorama, the web interface displays them in evaluation order. All the shared, device group, and default rules that the firewall inherits from Panorama are shaded orange. Local firewall rules display between the pre-rules and post-rules.



 ©2025 Palo Alto Networks, Inc.

Evaluation Order	Rule Scope and Description	Administration Device
Shared pre-rules Device group pre-rules	Panorama pushes shared pre-rules to all the firewalls in all device groups. Panorama pushes device group-specific pre-rules to all the firewalls in a particular device group and its descendant device groups. If a firewall inherits rules from device groups at multiple levels in the device group hierarchy, it evaluates pre-rules in the order of highest to lowest level. This means the firewall first evaluates shared rules and last evaluates the rules of device groups with no descendants. You can use pre-rules to enforce the acceptable use policy of an organization. For example, a pre-rule might block access to specific URL categories or allow Domain Name System (DNS) traffic for all users.	These rules are visible on firewalls but you can only manage them in Panorama.
Local firewall rules	Local rules are specific to a single firewall or virtual system (vsys).	A local firewall administrator, or a Panorama administrator who switches to a local firewall context, can edit local firewall rules.
Device group post-rules Shared post-rules	Panorama pushes shared post-rules to all the firewalls in all device groups. Panorama pushes device group-specific post-rules to all the firewalls in a particular device group and its descendant device groups. If a firewall inherits rules from device groups at multiple levels in the device group hierarchy, it evaluates post-rules in the order of lowest to highest level. This means the firewall first evaluates the rules of device groups with no descendants and last evaluates shared rules. Post-rules typically include rules to deny access to traffic based on the App-ID™ signatures, User-ID™	These rules are visible on firewalls but you can only manage them in Panorama.

 ©2025 Palo Alto Networks, Inc.

Evaluation Order	Rule Scope and Description	Administration Device
	information (users or user groups), or service.	
intrazone-default interzone-default	The default rules apply only to the Security rulebase, and are predefined on Panorama (at the Shared level) and the firewall (in each vsys). These rules specify how PAN-OS handles traffic that doesn't match any other rule. The intrazone-default rule allows all traffic within a zone. The interzone-default rule denies all traffic between zones. If you override default rules, their order of precedence runs from the lowest context to the highest: overridden settings at the firewall level take precedence over settings at the device group level, which take precedence over settings at the Shared level.	Default rules are initially read-only, either because they are part of the predefined configuration or because Panorama pushed them to firewalls. However, you can override the rule settings for tags, action, logging, and security profiles. The context determines the level at which you can override the rules: • Panorama—At the Shared or device group level, you can override default rules that are part of the predefined configuration. • Firewall—You can override default rules that are part of the predefined configuration on the firewall or vsys, or that Panorama pushed from the Shared location or a device group.

Device Group Objects

Objects are configuration elements that policy rules reference, for example: IP addresses, URL categories, security profiles, users, services, and applications. Rules of any type (pre-rules, post-rules, default rules, and rules locally defined on a firewall) and any rulebase (Security, NAT, QoS, Policy Based Forwarding, Decryption, Application Override, Captive Portal, and DoS Protection) can reference objects. You can reuse an object in any number of rules that have the same scope as that object in the Device Group Hierarchy. For example, if you add an object to the Shared location, all rules in the hierarchy can reference that *shared object* because all device groups inherit objects from Shared. If you add an object to a particular device group, only the rules in that device group and its descendant device groups can reference that *device group object*. If object values in a device group must differ from those inherited from an ancestor device group, you can Override inherited object values (see Step Override inherited object values.). You can also Revert to Inherited Object Values at any time. When you Create Objects for Use in Shared or Device Group Policy once and use them many times, you reduce administrative overhead and ensure consistency across firewall policies.

You can configure how Panorama handles objects system-wide:

- **Pushing unused objects**—By default, Panorama pushes all objects to firewalls regardless of whether any shared or device group policy rules reference the objects. Optionally, you can

 ©2025 Palo Alto Networks, Inc.

configure Panorama to push only referenced objects. For details, see Manage Unused Shared Objects.

- **Precedence of ancestor and descendant objects**—By default, when device groups at multiple levels in the hierarchy have an object with the same name but different values (because of overrides, as an example), policy rules in a descendant device group use the object values in that descendant instead of object values inherited from ancestor device groups or Shared. Optionally, you can reverse this order of precedence to push values from Shared or the highest ancestor containing the object to all descendant device groups. For details, see Manage Precedence of Inherited Objects.

 ©2025 Palo Alto Networks, Inc.

Centralized Logging and Reporting

Panorama aggregates logs from all managed firewalls and provides visibility across all the traffic on the network. It also provides an audit trail for all policy modifications and configuration changes made to the managed firewalls. In addition to aggregating logs, Panorama can forward them as SNMP traps, email notifications, syslog messages, and HTTP payloads to an external server.

For centralized logging and reporting, you also have the option to use the cloud-based Strata Logging Service that is architected to work seamlessly with Panorama. The Strata Logging Service allows your managed firewalls to forward logs to the Strata Logging Service infrastructure instead of to Panorama or to the managed Log Collectors, so you can augment your existing distributed log collection setup or to scale your current logging infrastructure without having to invest time and effort yourself.

The Application Command Center (ACC) on Panorama provides a single pane for unified reporting across all the firewalls. It enables you to centrally Monitor Network Activity, to analyze, investigate, and report on traffic and security incidents. On Panorama, you can view logs and generate reports from logs forwarded to the Strata Logging Service, Panorama or to the managed Log Collectors, if configured, or you can query the managed firewalls directly. For example, you can generate reports about traffic, threat, and/or user activity in the managed network based on logs stored on Panorama (and the managed collectors) or by accessing the logs stored locally on the managed firewalls, or in the Strata Logging Service.

If you don't Configure Log Forwarding to Panorama or the Strata Logging Service, you can schedule reports to run on each managed firewall and forward the results to Panorama for a combined view of user activity and network traffic. Although reports don't provide a granular drill-down on specific information and activities, they still provide a unified monitoring approach.

- Managed Collectors and Collector Groups
- Local and Distributed Log Collection
- Caveats for a Collector Group with Multiple Log Collectors
- Log Forwarding Options
- Centralized Reporting

Managed Collectors and Collector Groups

Panorama uses Log Collectors to aggregate logs from managed firewalls. When generating reports, Panorama queries the Log Collectors for log information, providing you visibility into all the network activity that your firewalls monitor. Because you use Panorama to configure and manage Log Collectors, they are also known as *managed collectors*. Panorama can manage two types of Log Collectors:

- **Local Log Collector**—This type of Log Collector runs locally on the Panorama management server. Only an M-700, M-600, M-500, M-300, or M-100 appliance, or Panorama virtual appliance in Panorama mode supports a local Log Collector.

  If you forward logs to a Panorama virtual appliance in Legacy mode, it stores the logs locally without a Log Collector.

 ©2025 Palo Alto Networks, Inc.

- **Dedicated Log Collector**—This is an M-700, M-600, M-500, M-300, M-200, or M-100 appliance or Panorama virtual appliance in Log Collector mode. You can use an M-Series appliance in Panorama mode or a Panorama virtual appliance in Panorama or Legacy (ESXi and vCloud Air) mode to manage Dedicated Log Collectors. To use the Panorama web interface for managing Dedicated Log Collectors, you must add them as managed collectors. Otherwise, administrative access to a Dedicated Log Collector is only available through its CLI using the predefined administrative user (*admin*) account. Dedicated Log Collectors don't support additional administrative user accounts.

You can use either or both types of Log Collectors to achieve the best logging solution for your environment (see Local and Distributed Log Collection).

A Collector Group is 1 to 16 managed collectors that operate as a single logical log collection unit. If the Collector Group contains Dedicated Log Collectors, Panorama uniformly distributes the logs across all the disks in each Log Collector and across all Log Collectors in the group. This distribution optimizes the available storage space. To enable a Log Collector to receive logs, you must add it to a Collector Group. You can enable log redundancy by assigning multiple Log Collectors to a Collector Group (see Caveats for a Collector Group with Multiple Log Collectors). The Collector Group configuration specifies which managed firewalls can send logs to the Log Collectors in the group.

To configure Log Collectors and Collector Groups, see Manage Log Collection.

Local and Distributed Log Collection

Before you Configure Log Forwarding to Panorama, you must decide whether to use local Log Collectors, Dedicated Log Collectors, or both.

A local Log Collector is easy to deploy because it requires no additional hardware or virtual machine instance. In a high availability (HA) configuration, you can send logs to the local Log Collector on both Panorama peers; the passive Panorama doesn't wait for failover to start collecting logs.

For local log collection, you can also forward logs to a Panorama virtual appliance in Legacy mode, which stores the logs without using a Log Collector as a logical container.

Dedicated Log Collectors are M-700, M-600, M-500, M-300, M-200, or Panorama virtual appliance in Log Collector mode. Because they perform only log collection, not firewall management, Dedicated Log Collectors allow for a more robust environment than local Log Collectors. Dedicated Log Collectors provide the following benefits:

- Enable the Panorama management server to use more resources for management functions instead of logging.
- Provide high-volume log storage on a dedicated hardware appliance.
- Enable higher logging rates.
- Provide horizontal scalability and redundancy with RAID 1 storage.
- Optimize bandwidth resources in networks where more bandwidth is available for firewalls to send logs to nearby Log Collectors than to a remote Panorama management server.
- Enable you to meet regional regulatory requirements (for example, regulations might not allow logs to leave a particular region).

 ©2025 Palo Alto Networks, Inc.

Distributed Log Collection illustrates a topology in which the Panorama peers in an HA configuration manage the deployment and configuration of firewalls and Dedicated Log Collectors.

You can deploy the Panorama management server in an HA configuration but not the Dedicated Log Collectors.

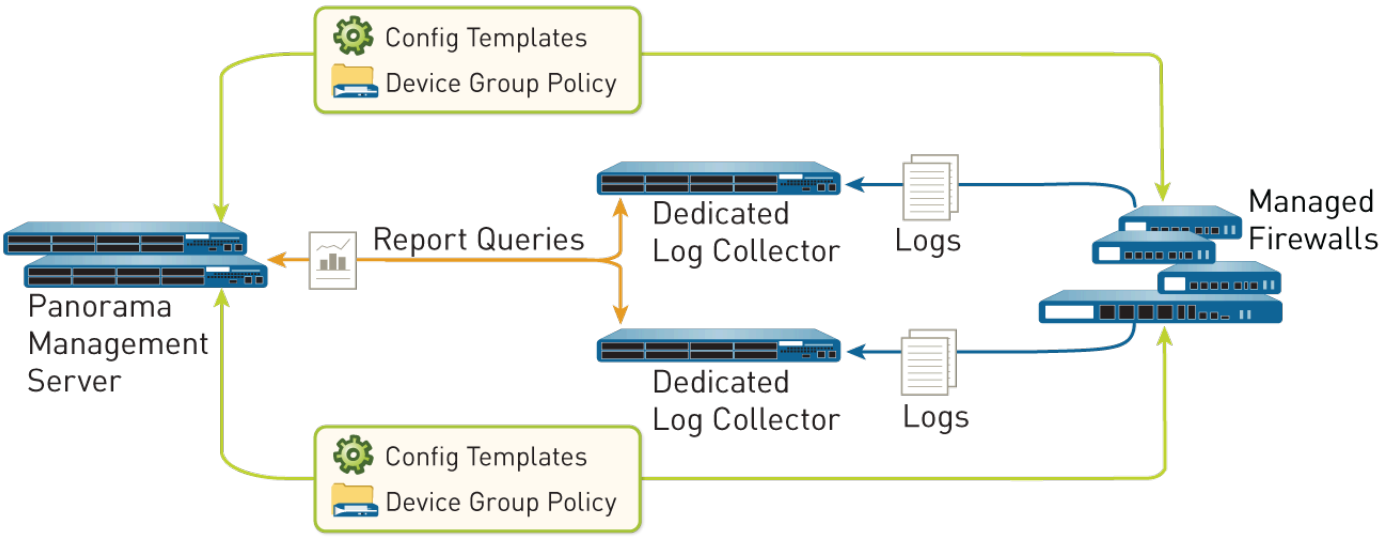


Figure 4: Distributed Log Collection

Caveats for a Collector Group with Multiple Log Collectors

You can Configure a Collector Group with multiple Log Collectors (up to 16) to ensure log redundancy, increase the log retention period, and accommodate logging rates that exceed the capacity of a single Log Collector (see Panorama Models for capacity information). In any single Collector Group, all the Log Collectors must run on the same Panorama model: all M-700 appliances, all M-600 appliances, all M-500 appliances, all M-300 appliances, all M-200 appliances, or all Panorama virtual appliances. For example, if a single managed firewall generates 48TB of logs, the Collector Group that receives those logs will require at least three Log Collectors that are M-300 appliances or one Log Collector that is an M-700 appliance or similarly resourced Panorama virtual appliance.

A Collector Group with multiple Log Collectors uses the available storage space as one logical unit and uniformly distributes the logs across all its Log Collectors. The log distribution is based on the disk capacity of the Log Collectors (see Panorama Models) and a hash algorithm that dynamically decides which Log Collector owns the logs and writes to disk. Although Panorama uses a preference list to prioritize the list of Log Collectors to which a managed firewall can forward logs, Panorama does not necessarily write the logs to the first Log Collector specified in the preference list. For example, consider the following preference list:

Managed Firewall	Log Forwarding Preference List Defined in a Collector Group
FW1	L1,L2,L3

 ©2025 Palo Alto Networks, Inc.

Managed Firewall	Log Forwarding Preference List Defined in a Collector Group
FW2	L4,L5,L6

Using this list, FW1 will forward logs to L1 so long as that primary Log Collector is available. However, based on the hash algorithm, Panorama might choose L2 as the owner that writes the logs to its disks. If L2 becomes inaccessible or has a chassis failure, FW1 will not know because it can still connect to L1.

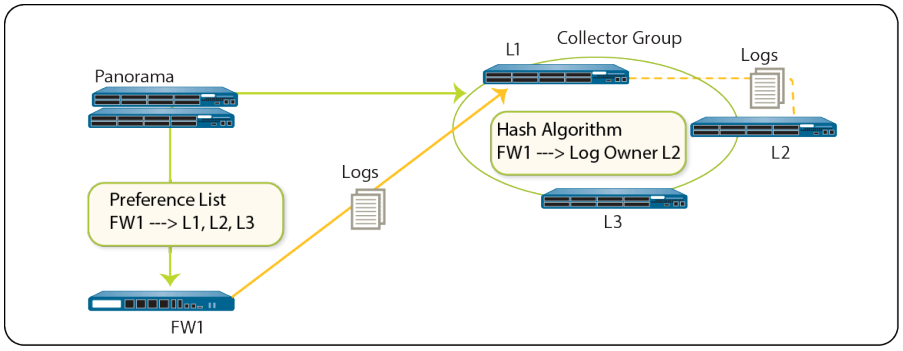


Figure 5: Example - Typical Log Collector Group Setup

In the case where a Collector Group has only one Log Collector and the Log Collector fails, the firewall stores the logs to its HDD/SSD (the available storage space varies by firewall model). As soon as connectivity is restored to the Log Collector, the firewall resumes forwarding logs where it left off before the failure occurred.

In the case of a Collector Group with multiple Log Collectors, the firewall does not buffer logs to its local storage if only one Log Collector is down. In the example scenario where L2 is down, FW1 continues sending logs to L1, and L1 stores the log data that would be sent to L2. Once L2 is back up, L1 no longer stores log data intended for L2 and distribution resumes as expected. If one of the Log Collectors in a Collector Group goes down, the logs that would be written to the down Log Collector are redistributed to the next Log Collector in the preference list.

Palo Alto Networks recommends adding at least three Log Collectors to a Collector Group to avoid split brain and log ingestion issues should one Log Collector go down. See the changes to default Collector Group behavior for more information.

(PAN-OS 10.2.9 and later 10.2 releases) Two Log Collectors in a Collector Group are supported and the Collector Group remains operational even if one Log Collector goes down.

 ©2025 Palo Alto Networks, Inc.

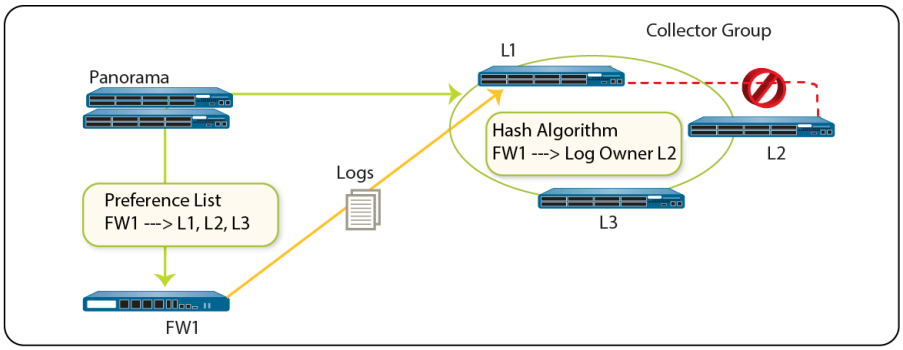


Figure 6: Example - When a Log Collector Fails

Palo Alto Networks recommends the following mitigations if using multiple Log Collectors in a Collector Group:

- Enable log redundancy when you Configure a Collector Group. This ensures that no logs are lost if any one Log Collector in the Collector Group becomes unavailable. Each log will have two copies and each copy will reside on a different Log Collector. Log redundancy is available only if each Log Collector has the same number of logging disks.

  Because enabling redundancy creates more logs, this configuration requires more storage capacity. When a Collector Group runs out of space, it deletes older logs.

  Enabling redundancy doubles the log processing traffic in a Collector Group, which reduces its maximum logging rate by half, as each Log Collector must distribute a copy of each log it receives.

- Obtain an On-Site-Spare (OSS) to enable prompt replacement if a Log Collector failure occurs.
- In addition to forwarding logs to Panorama, configure forwarding to an external service as backup storage. The external service can be a syslog server, email server, SNMP trap server, or HTTP server.

Log Forwarding Options

By default, each firewall stores its log files locally. To use Panorama for centralized log monitoring and report generation, you must Configure Log Forwarding to Panorama. Logs are forwarded over the management interface by default unless you configure a dedicated service route to forward logs. Panorama supports forwarding logs to either a Log Collector, the Strata Logging Service, or both in parallel. You can also use external services for archiving, notification, or analysis by forwarding logs to the services directly from the firewalls or from Panorama. External services include the syslog servers, email servers, SNMP trap servers, or HTTP-based services. In addition to forwarding firewall logs, you can forward the logs that the Panorama management server and Log Collectors generate. The Panorama management server, Log Collector, or firewall that forwards the logs converts them to a format that is appropriate for the destination (syslog message, email notification, SNMP trap, or HTTP payload).

Palo Alto Networks firewalls and Panorama support the following log forwarding options. Before choosing an option, consider the logging capacities of your Panorama Models and Determine Panorama Log Storage Requirements.

- Forward logs from firewalls to Panorama and from Panorama to external services—This configuration is best for deployments in which the connections between firewalls and external

 ©2025 Palo Alto Networks, Inc.

services have insufficient bandwidth to sustain the logging rate, which is often the case when the connections are remote. This configuration improves firewall performance by offloading some processing to Panorama.

You can configure each Collector Group to forward logs to different destinations.

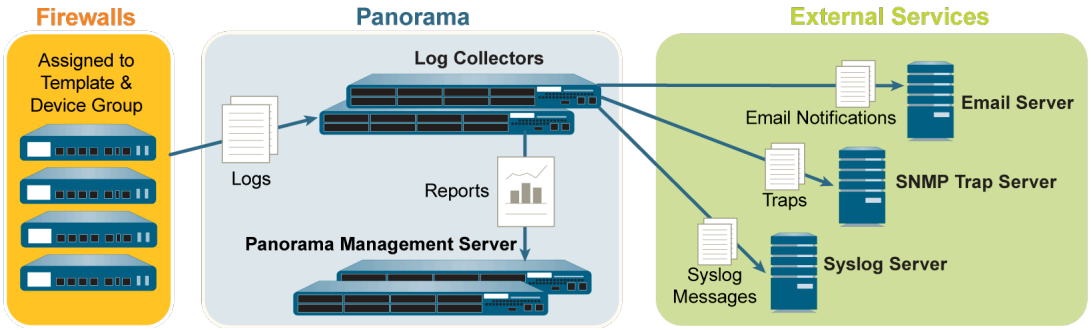


Figure 7: Log Forwarding to Panorama and then to External Services

- Forward logs from firewalls to Panorama and to external services in parallel—In this configuration, both Panorama and the external services are endpoints of separate log forwarding flows; the firewalls don't rely on Panorama to forward logs to external services. This configuration is best for deployments in which the connections between firewalls and external services have sufficient bandwidth to sustain the logging rate, which is often the case when the connections are local.

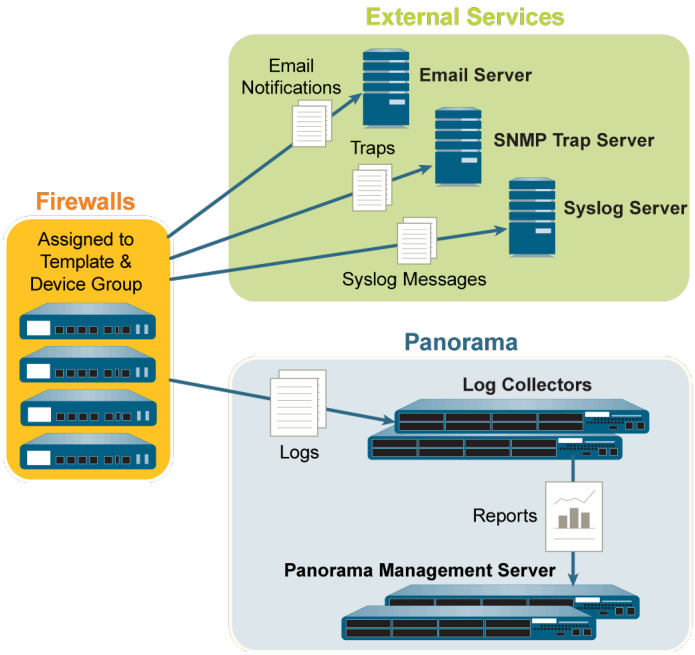


Figure 8: Log Forwarding to External Services and Panorama in Parallel

Centralized Reporting

Panorama aggregates logs from all managed firewalls and enables reporting on the aggregated data for a global view of application use, user activity, and traffic patterns across the entire

 ©2025 Palo Alto Networks, Inc.

network. As soon as the firewalls are added to Panorama, the ACC can display all traffic traversing your network. With logging enabled, clicking into a log entry in the ACC provides direct access to granular details about the application.

For generating reports, Panorama uses two sources: the local Panorama database and the remote firewalls that it manages. The Panorama database refers to the local storage on Panorama that is allocated for storing both summarized logs and some detailed logs. If you have a distributed Log Collection deployment, the Panorama database includes the local storage on Panorama and all the managed Log Collectors. Panorama summarizes the information—traffic, application, threat — collected from all managed firewalls at 15-minute intervals. Using the local Panorama database allows for faster response times, however, if you prefer to not forward logs to Panorama, Panorama can directly access the remote firewall and run reports on data that is stored locally on the managed firewalls.

Panorama offers more than 40 predefined reports that can be used as is, or they can be customized by combining elements of other reports to generate custom reports and report groups that can be saved. Reports can be generated on demand, on a recurring schedule, and can be scheduled for email delivery. These reports provide information on the user and the context so that you correlate events and identify patterns, trends, and potential areas of interest. With the integrated approach to logging and reporting, the ACC enables correlation of entries from multiple logs relating to the same event.

For more information, see Monitor Network Activity.

 ©2025 Palo Alto Networks, Inc.

Data Redistribution Using Panorama

With data redistribution, you only have to configure each source once, then you can redistribute multiple data types to as many clients as needed. This helps you to scale your network so that you can easily add or remove source and clients as your network needs change.

Data redistribution also provides granularity by redistributing only the types of information to only the firewalls or Panorama management systems that you specify. You can use subnets, ranges, and regions to further reduce network traffic and maximize device capacity.

One of the key benefits of the Palo Alto Networks firewall is that it can enforce policies and generate reports based on usernames and tags instead of IP addresses. The challenge for large-scale networks is ensuring every firewall that enforces policies and generates reports has the mappings and tags that apply for all of your policy rules. Additionally, every firewall that enforces Authentication Policy requires a complete, identical set of authentication timestamps for your user base. Whenever users authenticate to access services and applications, individual firewalls record the associated timestamps but don't automatically share them with other firewalls to ensure consistency. Data redistribution solves these challenges for large-scale networks by enabling you to redistribute the necessary data. However, instead of setting up extra connections to redistribute the data between firewalls, you can leverage your Panorama infrastructure to Redistribute Data to Managed Firewalls. The infrastructure has existing connections that enable you to redistribute data in layers, from firewalls to Panorama. Panorama can then redistribute the information to the firewalls that enforce policies and generate reports.

Each firewall or Panorama management server can receive data from up to 100 redistribution points. The redistribution points can be other firewalls or Panorama management servers. However, you can also use Windows-based User-ID agents to perform the mapping and redistribute the information to firewalls. Only the firewalls record authentication timestamps when user traffic matches Authentication policy rules.

 ©2025 Palo Alto Networks, Inc.

Role-Based Access Control

Role-based access control (RBAC) enables you to define the privileges and responsibilities of administrative users (administrators). Every administrator must have a user account that specifies a role and authentication method. Administrative Roles define access to specific configuration settings, logs, and reports within Panorama and firewall contexts. For Device Group and Template administrators, you can map roles to Access Domains, which define access to specific device groups, templates, and firewalls (through context switching). By combining each access domain with a role, you can enforce the separation of information among the functional or regional areas of your organization. For example, you can limit an administrator to monitoring activities for data center firewalls but allow that administrator to set policies for test lab firewalls. By default, every Panorama appliance (virtual appliance or M-Series appliance) has a predefined administrative account (admin) that provides full read-write access (superuser access) to all functional areas and to all device groups, templates, and firewalls. For each administrator, you can define an authentication profile that determines how Panorama verifies user access credentials.



Instead of using the default account for all administrators, it is a best practice to create a separate administrative account for each person who needs access to the administrative or reporting functions on Panorama. This provides better protection against unauthorized configuration changes and enables Panorama to log and identify the actions of each administrator.

- Administrative Roles
- Authentication Profiles and Sequences
- Access Domains
- Administrative Authentication

Administrative Roles

You configure administrator accounts based on the security requirements of your organization, any existing authentication services that your network uses, and the required administrative roles. A *role* defines the type of system access that is available to an administrator. You can define and restrict access as broadly or granularly as required, depending on the security requirements of your organization. For example, you might decide that a data center administrator can have access to all device and networking configurations, but a security administrator can control only security policy definitions, while other key individuals can have limited CLI or XML API access. The role types are:

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

Dynamic Role	Privileges
Superuser	Full read-write access to Panorama

 ©2025 Palo Alto Networks, Inc.

Dynamic Role	Privileges
Superuser (read-only)	Read-only access to Panorama
Panorama administrator	Full access to Panorama except for the following actions: • Create, modify, or delete Panorama or firewall administrators and roles. • Export, validate, revert, save, load, or import a configuration in the **Device** > **Setup** > **Operations** page. • Configure **Scheduled Config Export** functionality in the **Panorama** tab. • **Generate Tech Support File**, **Generate Stats Dump File**, and **Download Core Files** (**Panorama** > **Support**)

- **Admin Role Profiles**—To provide more granular access control over the functional areas of the web interface, CLI, and XML API, you can create custom roles. When new features are added to the product, you must update the roles with corresponding access privileges: Panorama does not automatically add new features to custom role definitions. You select one of the following profile types when you Configure an Admin Role Profile.

Admin Role Profile	Description
Panorama	For these roles, you can assign read-write access, read-only access, or no access to all the Panorama features that are available to the superuser dynamic role except the management of Panorama administrators and Panorama roles. For the latter two features, you can assign read-only access or no access, but you cannot assign read-write access. An example use of a Panorama role would be for security administrators who require access to security policy definitions, logs, and reports on Panorama. Custom Panorama admin roles have the following limitations: • No access to **Reboot Panorama** (**Panorama** > **Setup** > **Operations**) • No access to **Generate Tech Support File**, **Generate Stats Dump File**, and **Download Core Files** (**Panorama** > **Support**)
Device Group and Template	For these roles, you can assign read-write access, read-only access, or no access to specific functional areas within device groups, templates, and firewall contexts. By combining these roles with Access Domains, you can enforce the separation of information among the functional or regional areas of your organization. Device Group and Template roles have the following limitations: • No access to the CLI or XML API

 ©2025 Palo Alto Networks, Inc.

Admin Role Profile	Description

	• No access to configuration or system logs
	• No access to VM information sources
	• No access to **Reboot Panorama** (**Panorama** > **Setup** > **Operations**)
	• No access to **Generate Tech Support File**, **Generate Stats Dump File**, and **Download Core Files** (**Panorama** > **Support**)
	• In the **Panorama** tab, access is limited to:
	  • Device deployment features (read-write, read-only, or no access)
	  • The device groups specified in the administrator account (read-write, read-only, or no access)
	  • The templates and managed firewalls specified in the administrator account (read-only or no access)
	An example use of this role would be for administrators in your operations staff who require access to the device and network configuration areas of the web interface for specific device groups and/or templates.

Authentication Profiles and Sequences

An authentication profile defines the authentication service that validates the login credentials of administrators when they access Panorama. The service can be local authentication or an external authentication service. Some services (SAML, TACACS+, and RADIUS) provide the option to manage both authentication and authorization for administrative accounts on the external server instead of on Panorama. In addition to the authentication service, the authentication profile defines options such as Kerberos single sign-on (SSO) and SAML single logout (SSO).

Some networks have multiple databases (such as TACACS+ and LDAP) for different users and user groups. To authenticate administrators in such cases, configure an authentication sequence—a ranked order of authentication profiles that Panorama matches an administrator against during login. Panorama checks against each profile in sequence until one successfully authenticates the administrator. An administrator is denied access only if authentication fails for all the profiles in the sequence.

Access Domains

Access domains control administrative access to specific Device Groups and templates, and also control the ability to switch context to the web interface of managed firewalls. Access domains apply only to administrators with Device Group and Template roles. Mapping Administrative Roles to access domains enables very granular control over the information that administrators access on Panorama. For example, consider a scenario where you configure an access domain that includes all the device groups for firewalls in your data centers and you assign that access domain to an administrator who is allowed to monitor data center traffic but who is not allowed to configure the firewalls. In this case, you would map the access domain to a role that enables all monitoring privileges but disables access to device group settings. Additionally, Device Group and Template admins can perform administrative tasks for managed firewalls in their access

 ©2025 Palo Alto Networks, Inc.

domain such as viewing the configuration and system logs, perform configuration audits, review pending tasks, and directly access firewall operations such as reboot, generating a tech support file, executing a stats dump, and exporting a core file.

You configure access domains in the local Panorama configuration and then assign them to administrative accounts and roles. You can perform the assignment locally or use an external SAML, TACACS+, or RADIUS server. Using an external server enables you to quickly reassign access domains through your directory service instead of reconfiguring settings on Panorama. To use an external server, you must define a server profile that enables Panorama to access the server. You must also define Vendor-Specific Attributes (VSAs) on the RADIUS or TACACS+ server, or SAML attributes on the SAML IdP server.

For example, if you use a RADIUS server, you would define a VSA number and value for each administrator. The value defined has to match the access domain configured on Panorama. When an administrator tries to log in to Panorama, Panorama queries the RADIUS server for the administrator access domain and attribute number. Based on the response from the RADIUS server, the administrator is authorized for access and is restricted to the firewalls, virtual systems, device groups, and templates that are assigned to the access domain.

For the relevant procedures, see:

- Configure an Access Domain.
- Configure RADIUS Authentication for Panorama Administrators.
- Configure TACACS+ Authentication for Panorama Administrators.
- Configure SAML Authentication for Panorama Administrators.

Administrative Authentication

You can configure the following types of authentication and authorization (Administrative Roles and Access Domains) for Panorama administrators:

Authentication Method	Authorization Method	Description
Local	Local	The administrative account credentials and authentication mechanisms are local to Panorama. You use Panorama to assign administrative roles and access domains to the accounts. To further secure the accounts, you can create a password profile that defines a validity period for passwords and set Panorama-wide password complexity settings. For details, see Configure Local or External Authentication for Panorama Administrators.
SSH Keys	Local	The administrative accounts are local to Panorama, but authentication to the CLI is based on SSH keys. You use Panorama to assign administrative roles and access domains to the accounts. For details, see Configure an Administrator with SSH Key-Based Authentication for the CLI.
Certificates	Local	The administrative accounts are local to Panorama, but authentication to the web interface is based on client certificates.

 ©2025 Palo Alto Networks, Inc.

Authentication Method	Authorization Method	Description
		You use Panorama to assign administrative roles and access domains to the accounts. For details, see Configure a Panorama Administrator with Certificate-Based Authentication for the Web Interface.
External service	Local	The administrative accounts you define locally on Panorama serve as references to the accounts defined on an external Multi-Factor Authentication, SAML, Kerberos, TACACS+, RADIUS, or LDAP server. The external server performs authentication. You use Panorama to assign administrative roles and access domains to the accounts. For details, see Configure Local or External Authentication for Panorama Administrators.
External	External service	The administrative accounts are defined only on an external SAML, TACACS+, or RADIUS server. The server performs both authentication and authorization. For authorization, you define Vendor-Specific Attributes (VSAs) on the TACACS+ or RADIUS server, or SAML attributes on the SAML server. Panorama maps the attributes to administrator roles and access domains that you define on Panorama. For details, see: • Configure SAML Authentication for Panorama Administrators • Configure TACACS+ Authentication for Panorama Administrators • Configure RADIUS Authentication for Panorama Administrators

©2025 Palo Alto Networks, Inc.

Panorama Commit, Validation, and Preview Operations

When you are ready to activate changes that you made to the candidate configuration on Panorama or to push changes to the devices that Panorama manages (firewalls, Log Collectors, and WildFire appliances and appliance clusters), you can Preview, Validate, or Commit Configuration Changes. For example, if you add a Log Collector to the Panorama configuration, firewalls cannot send logs to that Log Collector until you commit the change to Panorama and then push the change to the Collector Group that contains the Log Collector.

You can filter changes by administrator or *location* and then commit, push, validate, or preview only those changes. The location can be specific device groups, templates, Collector Groups, Log Collectors, shared settings, or the Panorama management server.

When you commit changes, they become part of the running configuration. Changes that you haven't committed are part of the candidate configuration. Panorama queues commit requests so that you can initiate a new commit while a previous commit is in progress. Panorama performs the commits in the order they are initiated but prioritizes auto-commits that are initiated by Panorama (such as FQDN refreshes). However, if the queue already has the maximum number of administrator-initiated commits (10), you must wait for Panorama to finish processing a pending commit before initiating a new one. You can Use the Panorama Task Manager ( ) to cancel pending commits or to see details about commits that are pending, in progress, completed, or failed. To check which changes a commit will activate, you can run a commit preview.

When you initiate a commit, Panorama checks the validity of the changes before activating them. The validation output displays conditions that block the commit (errors) or that are important to know (warnings). For example, validation could indicate an invalid route destination that you need to fix for the commit to succeed. The validation process enables you to find and fix errors before you commit (it makes no changes to the running configuration). This is useful if you have a fixed commit window and want to be sure the commit will succeed without errors.

When you preview your configuration commit, any configuration object added between existing any other existing object is displayed as a modified configuration object rather than an added configuration object. For example, `Address1` and `Address2` are existing Address objects. A Panorama admin later creates `Address3` and adds the Address object between `Address1` and `Address2`. When the Panorama admin goes to preview the configuration changes, `Address3` is displayed as a modified configuration object.

Automated commit recovery is enabled by default, allowing the managed firewalls to locally test the configuration pushed from Panorama to verify that the new changes do not break the connection between Panorama and the managed firewall. If the committed configuration breaks the connection between Panorama and a managed firewall then the firewall automatically fails the commit and the configuration is reverted to the previous running configuration and the Shared Policy or Template Status (**Panorama** > **Managed Devices** > **Summary**) gets out of sync depending on which configuration objects were pushed. Additionally, the managed firewalls test their connection to Panorama every 60 minutes and if a managed firewall detects that it can no longer successfully connect to Panorama then it reverts its configuration to the previous running configuration.

 ©2025 Palo Alto Networks, Inc.

For details on candidate and running configurations, see Manage Panorama and Firewall Configuration Backups.

To prevent multiple administrators from making configuration changes during concurrent sessions, see Manage Locks for Restricting Configuration Changes.

When pushing configurations to managed firewalls, Panorama pushes the running configuration. Because of this, Panorama does not let you push changes to managed firewalls until you first commit the changes to Panorama.

 ©2025 Palo Alto Networks, Inc.

Plan Your Panorama Deployment

- Determine the management approach. Do you plan to use Panorama to centrally configure and manage the policies, to centrally administer software, content and license updates, and/or centralize logging and reporting across the managed firewalls in the network?

  If you already deployed and configured Palo Alto Networks firewalls on your network, determine whether to transition the firewalls to centralized management. This process requires a migration of all configuration and policies from your firewalls to Panorama. For details, see Transition a Firewall to Panorama Management.
- Verify the Panorama and firewall software versions. Panorama can manage firewalls running PAN-OS versions that match the Panorama version or are earlier than the Panorama version. See Panorama Management Compatibility for more information.
- (Multi-vsys firewalls) If you already deployed and configured multi-vsys Palo Alto Networks firewalls on your network, Palo Alto Networks recommends you transition and manage all vsys configurations of the multi-vsys firewall from Panorama. This is required to avoid commit issues on the multi-vsys firewall and allows you to take advantage of the optimized shared object pushes from Panorama.
- (Multi-vsys firewalls) Delete or rename any locally configured firewall **Shared** object that has an identical name to an object in the Panorama **Shared** configuration. Otherwise, configuration pushes from Panorama fail after the upgrade and display the error `<object-name> is already in use`.
- Determine your authentication method between Panorama and its managed devices and high availability peer. By default, Panorama uses predefined certificates to authenticate the SSL connections used for management and inter-device communication. However, you can configure custom certificate-based authentication to enhance the security of the SSL connections between Panorama, firewalls, and log collectors. By using custom certificates, you can establish a unique chain of trust to ensure mutual authentication between Panorama and the devices it manages. You can import the certificates from your enterprise public key infrastructure (PKI) or generate it on Panorama.
- Plan to use Panorama in a high availability configuration; set it up as an active/passive high availability pair. See Panorama High Availability.
- Plan how to accommodate network segmentation and security requirements in a large-scale deployment. By default, Panorama running on an M-Series appliance uses the management (MGT) interface for administrative access to Panorama and for managing devices (firewalls, Log Collectors, and WildFire appliances and appliance clusters), collecting logs, communicating with Collector Groups, and deploying software and content updates to devices. However, to improve security and enable network segmentation, you can reserve the MGT interface for administrative access and use dedicated M-Series Appliance Interfaces (Eth1, Eth2, Eth3, Eth4, and Eth5) for the other services.
- For meaningful reports on network activity, plan a logging solution:
  - Verify the resource allocation for your Panorama virtual appliance deployed in Log Collector mode on AWS or Azure. The Panorama virtual appliance does not retain Log Collector mode if resized. This results in log data loss.
  - Estimate the log storage capacity your network needs to meet security and compliance requirements. Consider such factors as the logging capacities of your Panorama Models,

 ©2025 Palo Alto Networks, Inc.

network topology, number of firewalls sending logs, type of log traffic (for example, URL Filtering and Threat logs versus Traffic logs), the rate at which firewalls generate logs, and the number of days for which you want to store logs on Panorama. For details, see Determine Panorama Log Storage Requirements.

- Do you need to forward logs to external services (such as a syslog server) in addition to Panorama? See Log Forwarding Options.
- Do you want to own or manage your own log storage on premises, or do you want to leverage the Strata Logging Service provided by Palo Alto Networks?
- If you need a long-term storage solution, do you have a Security Information and Event Management (SIEM) solution, such as Splunk or ArcSight, to which you can forward logs?
- Do you need redundancy in logging?

  If you configure a Collector Group with multiple Log Collectors, you can enable redundancy to ensure that no logs are lost if any one Log Collector becomes unavailable (see Caveats for a Collector Group with Multiple Log Collectors).

  If you deploy Panorama virtual appliances in Legacy mode in an HA configuration, the managed firewalls can send logs to both HA peers so that a copy of each log resides on each peer. This redundancy option is enabled by default (see Modify Log Forwarding and Buffering Defaults).
- Will you log to a Network File System (NFS)? If the Panorama virtual appliance is in Legacy mode and does not manage Dedicated Log Collectors, NFS storage is the only option for increasing log storage capacity beyond 8TB. NFS storage is available only if Panorama runs on an ESXi server. If you use NFS storage, keep in mind that the firewalls can send logs only to the primary peer in the HA pair; only the primary peer is mounted to the NFS and can write to it.

- ❑ Determine which role-based access privileges administrators require to access managed firewalls and Panorama. See Set Up Administrative Access to Panorama.
- ❑ Plan the required Device Groups. Consider whether to group firewalls based on function, security policy, geographic location, or network segmentation. An example of a function-based device group is one that contains all the firewalls that a Research and Development team uses. Consider whether to create smaller device groups based on commonality, larger device groups to scale more easily, or a Device Group Hierarchy to simplify complex layers of administration.
- ❑ Plan a layering strategy for administering policies. Consider how firewalls inherit and evaluate policy rules within the Device Group Hierarchy, and how to best implement shared rules, device-group rules, and firewall-specific rules to meet your network needs. For visibility and centralized policy management, consider using Panorama for administering rules even if you need firewall-specific exceptions for shared or device group rules. If necessary, you can Push a Policy Rule to a Subset of Firewalls within a device group.
- ❑ Plan the organization of your firewalls based on how they inherit network configuration settings from Templates and Template Stacks. For example, consider assigning firewalls to templates based on hardware models, geographic proximity, and similar network needs for time zones, a DNS server, and interface settings.

 ©2025 Palo Alto Networks, Inc.

Deploy Panorama: Task Overview

The following task list summarizes the steps to get started with Panorama. For an example of how to use Panorama for central management, see Use Case: Configure Firewalls Using Panorama.

STEP 1 | (M-Series appliance only) Rack mount the appliance.

STEP 2 | Perform initial configuration to enable network access to Panorama. See Set Up the Panorama Virtual Appliance or Set Up the M-Series Appliance.

STEP 3 | Register Panorama and Install Licenses.

STEP 4 | Install Content and Software Updates for Panorama.

STEP 5 | (Recommended) Set up Panorama in a high availability configuration. See Panorama High Availability.

STEP 6 | Add a Firewall as a Managed Device.

STEP 7 | Add a Device Group or Create a Device Group Hierarchy, Add a Template, and (if applicable) Configure a Template Stack.

STEP 8 | (Optional) Configure log forwarding to Panorama and/or to external services. See Manage Log Collection.

STEP 9 | Monitor Network Activity using the visibility and reporting tools on Panorama.

 ©2025 Palo Alto Networks, Inc.

©2025 Palo Alto Networks, Inc.

Set Up Panorama

For centralized reporting and cohesive policy management across all the firewalls on your network, you can deploy the Panorama™ management server as a virtual appliance or as a hardware appliance (the M-200, M-300, M-500, M-600, or M-700 appliance).

The following topics describe how to set up Panorama on your network:

- Determine Panorama Log Storage Requirements
- Manage Large-Scale Firewall Deployments
- Set Up the Panorama Virtual Appliance
- Set Up the M-Series Appliance
- Register Panorama and Install Licenses
- Install the Panorama Device Certificate
- Install the Device Certificate for a Dedicated Log Collector
- Transition to a Different Panorama Model
- Access and Navigate Panorama Management Interfaces
- Set Up Administrative Access to Panorama
- Set Up Authentication Using Custom Certificates

Determine Panorama Log Storage Requirements

When you Plan Your Panorama Deployment, estimate how much log storage capacity Panorama requires to determine which Panorama Models to deploy, whether to expand the storage on those appliances beyond their default capacities, whether to deploy Dedicated Log Collectors, and whether to Configure Log Forwarding from Panorama to External Destinations. When log storage reaches the maximum capacity, Panorama automatically deletes older logs to create space for new ones.

Perform the following steps to determine the approximate log storage that Panorama requires. For details and use cases, refer to Panorama Sizing and Design Guide.

STEP 1 | Determine the log retention requirements of your organization.

Factors that affect log retention requirements include:

- IT policy of your organization
- Log redundancy—If you enable log redundancy when you Configure a Collector Group, each log will have two copies, which doubles your required log storage capacity.
- Regulatory requirements, such as those specified by the Payment Card Industry Data Security Standard (PCI DSS), Sarbanes-Oxley Act, and Health Insurance Portability and Accountability Act (HIPAA).



If your organization requires the removal of logs after a certain period, you can set the expiration period for each log type. You can also set a storage quota for each log type as a percentage of the total space if you need to prioritize log retention by type. For details, see Manage Storage Quotas and Expiration Periods for Logs and Reports.

STEP 2 | Determine the average daily logging rates.

Do this multiple times each day at peak and non-peak times to estimate the average. The more often you sample the rates, the more accurate your estimate.

1. Display the current log generation rate in logs per second:
   - If Panorama is not yet collecting logs, access the CLI of each firewall, run the following command, and calculate the total rates for all the firewalls. This command displays the number of logs received in the last second.

     ```
     > debug log-receiver statistics
     ```

   - If Panorama is already collecting logs, run the following command at the CLI of each appliance that receives logs (Panorama management server or Dedicated Log

 ©2025 Palo Alto Networks, Inc.

Collector) and calculate the total rates. This command gives the average logging rate for the last five minutes.

```
> debug log-collector log-collection-stats show incoming-logs
```

You can also use an SNMP manager to determine the logging rates of Log Collectors (see the panLogCollector MIB, OID 1.3.6.1.4.1.25461.1.1.6) and firewalls (see the panDeviceLogging, OID 1.3.6.1.4.1.25461.2.1.2.7).

2. Calculate the average of the sampled rates.
3. Calculate the daily logging rate by multiplying the average logs-per-second by 86,400.

STEP 3 | Estimate the required storage capacity.

This formula provides only an estimate; the exact amount of required storage will differ from the formula result.

Use the formula:

[(<logs_per_second> x 86400) x <days_of retention>] x <average_log_size> ÷ (1024 x 1024 x 1024)

The average log size varies considerably by log type. However, you can use 489 bytes as an approximate average log size.

For example, if Panorama must store logs for 30 days with a log rate of 1,500 LPS, then the required log storage capacity is: **[(1500 x 86400) x 30] x 489 ÷ (1024 x 1024 x 1024) = 1770GB**.

The results above are calculations for the detailed logs only with the default quota settings reserve 60% of the available storage for detail logs. This means that the calculated number represents 60% of the storage used by the Log Collector. To calculate the total storage required, divide this number by .60: **1770 ÷ .6 = 2951GB**.

One third, roughly 33%, of the avilable disk space is allocated to logd formatted logs to support upgrades, downgrades, and to support fixing database corruption. To calculate the total storage, divide the storage required by .66: **2951 ÷ .66 = 4471 total storage**.

STEP 4 | Next steps...

If you determine that Panorama requires more log storage capacity:

- Expand Log Storage Capacity on the Panorama Virtual Appliance.
- Increase Storage on the M-Series Appliance.

 ©2025 Palo Alto Networks, Inc.

Manage Large-Scale Firewall Deployments

Panorama™ provides multiple options to manage a large-scale firewall deployment. For consolidation of all management functions, Panorama supports management of up to 5,000 firewalls using an M-600, M-700 appliance, or Panorama virtual appliance on ESXi in Management Only mode or up to 2,500 firewalls with a Panorama virtual appliance in Management Only mode. To simplify the deployment and operational management of a large-scale firewall deployment greater than 5,000 firewalls, the Panorama Interconnect plugin allows you to manage multiple Panorama management server Nodes from a single Panorama Controller.

- Determine the Optimal Large-Scale Firewall Deployment Solution
- Increased Device Management Capacity for M-Series and Panorama Virtual Appliance

Determine the Optimal Large-Scale Firewall Deployment Solution

To ease the operational burden of managing the configuration of your large-scale firewall deployment, Palo Alto Networks provides different firewall management options to best suit your deployment scenario.

If your large-scale firewall deployment is composed of one or very few Panorama management servers, you can deploy an M-600, M-700 appliance, or Panorama virtual appliance on ESXi to manage up to 5,000 firewalls, or Panorama virtual appliance to manage up to 2,500 firewalls, to leverage all Panorama capabilities from a single Panorama management server. The Increased Device Management Capacity for M-Series and Panorama Virtual Appliance is ideal for vertically scaled deployments where you manage a large number of firewalls from a single Panorama management server rather than deploying multiple Panorama management servers to manage fewer firewalls.

If your large-scale firewall deployment is composed of multiple Panorama management servers with similar configurations, the Panorama Interconnect plugin allows you to manage multiple Panorama Nodes from a single Panorama Controller. This plugin simplifies the deployment and operational management of large scale firewall deployments because you can centrally manage policy and configuration from a Panorama Controller. From the Panorama Controller, the device group and template stack configuration is synchronized to the Panorama Nodes and pushed to managed devices. The Panorama Interconnect plugin is ideal for horizontally-scaled firewall deployments with multiple distributed Panorama management servers.

Increased Device Management Capacity for M-Series and Panorama Virtual Appliance

You can manage up to 5,000 firewalls using a single M-600, M-700 appliance, or Panorama™ virtual appliance installed on VMware ESXi, or up to 2,5000 firewalls with all other supported Panorama virtual appliances in order to reduce the management footprint of your large-scale firewall deployment.

- Increased Device Management Capacity Requirements
- Install Panorama for Increased Device Management Capacity

 ©2025 Palo Alto Networks, Inc.

Increased Device Management Capacity Requirements

You can manage up to 5,000 firewalls using a single M-600, M-700 appliance, or Panorama™ virtual appliance installed on VMware ESXi, or up to 2,500 firewalls with all other supported Panorama virtual appliances. For managing such large deployments from a single Panorama management server alleviates the operational complexity of configuration management and reduces the security and compliance risk of managing multiple Panorama management servers. For details about the maximum number of template stacks that Panorama supports, see Template Stacks.

For log collection, a single Panorama management server is ideal because it provides a centralized location to view and analyze log data from managed devices rather than requiring you to access each individual Panorama management server. To provide redundancy in the event of system or network failure, Palo Alto Networks recommends deploying two Panorama management servers in a high availability (HA) configuration. For Panorama system and config logs, an additional disk with a minimum 92GB capacity is required. This additional disk is automatically detected by the Panorama virtual appliance when Panorama is rebooted and mounted as a partition for system and config log storage.

For generating pre-defined reports, you must enable Panorama to use Panorama data for pre-defined reports. This generates pre-defined reports using log data already collected by Panorama or the Dedicated Log Collector, which reduces the resource utilization when generating reports. Enabling this setting is required, otherwise Panorama performance may be impacted, and Panorama may become unresponsive.

The Panorama management server must meet the following minimum requirements to support increased device management capacity:

Requirement	5,000 Firewalls	2,500 Firewalls
Model	M-600 M-700 (10.2.5 and later) VMware ESXi	All supported Panorama hypervisors. For more information, see Panorama Models.
Panorama Mode	Management Only	Management Only
System Disk	Used to store the operating system files, system logs, software updates, and content updates. • **M-Series Appliances**—240GB SSD • (10.2.5 and later)**ESXi**—224GB You must manually increase the system disk to 224GB.	• 81GB—Used to store the operating system files and system logs. • Additional disk with a minimum 92GB capacity used for storing Panorama system and config logs.
CPUs	56	32

 ©2025 Palo Alto Networks, Inc.

Requirement	5,000 Firewalls	2,500 Firewalls
Memory	256GB	256GB
Log Collection	Local log collection is not supported. See Deploy Panorama with Dedicated Log Collectors to set up log collection.	
Logging and Reporting	Enable the **Use Panorama Data for Pre-Defined Reports** setting (**Panorama** > **Setup** > **Management** > **Logging and Reporting Settings** > **Log Export and Reporting**)	

Install Panorama for Increased Device Management Capacity

Activate the device management license to manage up 5,000 firewalls from a single Panorama™ management server.

STEP 1 | Contact your Palo Alto Networks sales representative to obtain the Panorama device management license that enables you to manage up to 5,000 firewalls.

- If you are deploying an M-600 appliance, obtain the `PAN-M-600-P-1K` device management license.
- If you are deploying an M-700 appliance, obtain the `PAN-M-700-P-1K` device management license.
- If you are deploying a Panorama virtual appliance, obtain the `PAN-PRA-1000` device management license.

STEP 2 | Set up the Panorama management server.

- (M-600 and M-700 appliances only) Set Up the M-Series Appliance.

or

- Set Up the Panorama Virtual Appliance.

STEP 3 | (ESXi on PAN-OS 10.2.5 or later) Increase the System Disk for Panorama on an ESXi Server to 224GB.

A 224GB system disk is required for a Panorama virtual appliance installed on VMware ESXi to manage up to 5,000 firewalls. Review the Increased Device Management Capacity Requirements for more information.

STEP 4 | Change the Panorama management server to Management Only mode if Panorama is not already in this mode.

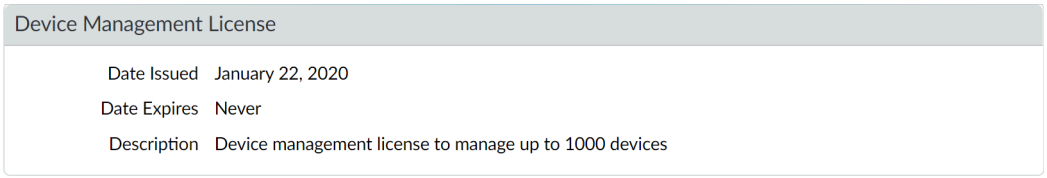
- Begin at Step 5 to Set Up an M-Series Appliance in Management Only Mode.
- Set up a Panorama Virtual Appliance in Management Only Mode.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Register your Panorama management server and install licenses.

1. Register Panorama.
2. Activate a Panorama Support License.
3. Activate the device management license on the Panorama management server.
   - Activate/Retrieve a Firewall Management License on the M-Series Appliance.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected.

STEP 6 | Select **Panorama** > **Licenses** and verify that the device management license is successfully activated.



If you are activating a new device management license on a Panorama, you can manage up to 5,000 firewalls with an M-600, M-700 appliance, or Panorama virtual appliance on ESXi, or up to 2,500 firewalls with a Panorama virtual appliance, but the Description still displays `Device management license to manage up to 1000 devices or more`.

 ©2025 Palo Alto Networks, Inc.

Set Up the Panorama Virtual Appliance

The Panorama virtual appliance enables you to use your existing VMware virtual infrastructure to centrally manage and monitor Palo Alto Networks firewalls and Dedicated Log Collectors. You can install the virtual appliance on an ESXi server, Alibaba Cloud, Amazon Web Services (AWS), AWS GovCloud, Microsoft Azure, Google Cloud Platform (GCP), KVM, Hyper-V, or in vCloud Air. In addition to or instead of deploying Dedicated Log Collectors, you can forward firewall logs directly to the Panorama virtual appliance. For greater log storage capacity and faster reporting, you have the option to switch the virtual appliance from Legacy mode to Panorama mode and configure a local Log Collector. For more details about the Panorama virtual appliance and its modes, see Panorama Models.

These topics assume you are familiar with the public and private hypervisor products required to create the virtual appliance, and don't cover any related concepts or terminology.

- Setup Prerequisites for the Panorama Virtual Appliance
- Install the Panorama Virtual Appliance
- Perform Initial Configuration of the Panorama Virtual Appliance
- Set Up The Panorama Virtual Appliance as a Log Collector
- Set Up the Panorama Virtual Appliance with Local Log Collector
- Set up a Panorama Virtual Appliance in Panorama Mode
- Set up a Panorama Virtual Appliance in Management Only Mode
- Expand Log Storage Capacity on the Panorama Virtual Appliance
- Increase CPUs and Memory on the Panorama Virtual Appliance
- Increase the System Disk on the Panorama Virtual Appliance
- Complete the Panorama Virtual Appliance Setup
- Convert Your Panorama Virtual Appliance

Setup Prerequisites for the Panorama Virtual Appliance

Complete the following tasks before you Install the Panorama Virtual Appliance:

- Use your browser to access the Palo Alto Networks Customer Support web site and Register Panorama You will need the Panorama serial number that you received in the order fulfillment email. After registering Panorama, you can access the Panorama software downloads page.
- Review the supported Panorama hypervisors to verify the hypervisor meets the minimum version requirements to deploy Panorama.
- If you will install Panorama on a VMware ESXi server, verify that the server meets the minimum requirements as listed in the System Requirements for the Panorama Virtual Appliance. These requirements apply to Panorama 5.1 and later releases. The requirements

 ©2025 Palo Alto Networks, Inc.

vary based on whether you will run the virtual appliance in Panorama mode or Management Only mode. For details on the modes, see Panorama Models.

If you install Panorama on VMware vCloud Air, you set the system settings during installation.

Review the minimum resource requirements for deploying the Panorama virtual appliance on Alibaba Cloud, Amazon Web Services (AWS), AWS GovCloud, Microsoft Azure, Google Cloud Platform (GCP), Hyper-V, KVM, Oracle Cloud Infrastructure (OCI), and VMware ESXi to ensure that the virtual machine meets the minimum required resources for the desired mode (Panorama, Management Only, or Log Collector). The minimum resource requirements for the Panorama virtual appliance are designed to help you achieve the maximum number of logs per second (LPS) for log collection in Panorama and Log Collector mode. If you add or remove virtual logging disks that results in a configuration that does not meet or exceed the number of virtual logging disks recommended (below), your LPS will be reduced.

If the minimum resource requirements are not met for Panorama mode when you Install the Panorama Virtual Appliance, Panorama defaults to Management Only mode for all supported public (Alibaba Cloud, AWS, AWS GovCloud, Azure, GCP, and OCI) and private (Hyper-V, KVM, and VMware ESXi) hypervisors. If the minimum resource requirements are not met for Management Only mode, Panorama defaults to Maintenance mode for all supported public hypervisors, Hyper-V, and KVM. If the minimum resource requirements for Management Only mode are not met when you Install Panorama on VMware, Panorama defaults to Legacy mode.

While still supported, Legacy mode is not recommended for production environments. Additionally, you can no longer switch Panorama to Legacy mode. For more information on supported modes, see Panorama Models.

Table 1: System Requirements for the Panorama Virtual Appliance

Requirement	Panorama Virtual Appliance in Management Only Mode	Panorama Virtual Appliance in Panorama Mode	Panorama Virtual Appliance in Log Collector Mode
Virtual hardware version	• **VMware ESXi and vCloud Air**—64-bit kernel-based VMware ESXi 6.0, 6.5, 6.7, 7.0, or 8.0. The supported version of the virtual hardware family type (also known as the VMware virtual hardware version) on the ESXi server is vmx-10. *The Panorama virtual appliance for ESXi **does not support** the following:* • *Creation of quiesced snapshots.* *Disable **Quiesce guest file system** in the vSphere client or set the **`quiesce`** flag to 0 or false in the vSphere CLI before creating a snapshot of your virtual Panorama appliance.* • *VMware vMotion to migrate a Panorama virtual appliance from one ESXi server to another.*		

 ©2025 Palo Alto Networks, Inc.

Requirement	Panorama Virtual Appliance in Management Only Mode	Panorama Virtual Appliance in Panorama Mode	Panorama Virtual Appliance in Log Collector Mode
	• **Hyper-V**—Windows Server 2016 with Hyper-V role or Hyper-V 2016, Windows Server 2019 with Hyper-V role or Hyper-V 2019, and Windows Server 2022 with Hyper-V role or Hyper-V 2022. • **KVM**—Ubuntu version 16.04 or CentOS7 In Panorama mode, the virtual appliance running on any ESXi version supports up to 12 virtual logging disks with 2TB of log storage each, for a total maximum capacity of 24TB. (VMware ESXi and vCloud Air only) In Legacy mode, the virtual appliance supports one virtual logging disk. ESXi 5.5 and later versions supports one disk of up to 8TB. Earlier ESXi versions support one disk of up to 2TB.		
(ESXi and vCloud Air only) Client computer	To install the Panorama virtual appliance and manage its resources, you must install a VMware vSphere Client or VMware Infrastructure Client that is compatible with your ESXi server.		
System disk	• **Default**—81GB • (ESXi and GCP only) **Upgraded**—224GB An upgraded system disk is required for SD-WAN. (Panorama and Log Collector mode) An upgraded system disk is required if you added more than 8 logging disks. For log storage, Panorama uses virtual logging disks instead of the system disk or an NFS datastore. Panorama must be initially installed with the **Default** system disk size, with the option to increase the system disk size after initial installation.		

 ©2025 Palo Alto Networks, Inc.

Requirement	Panorama Virtual Appliance in Management Only Mode	Panorama Virtual Appliance in Panorama Mode	Panorama Virtual Appliance in Log Collector Mode

CPUs, memory, and logging disks	• Manage up to 500 managed devices
	  • 16 CPUs
	  • 32GB memory
	  • Local log storage not supported
	• Manage up to 1,000 managed devices
	  • 32 CPUs
	  • 128GB memory
	  • Local log storage not supported
	• To manage more than 1,000 firewalls, see Increased Device Management Capacity Requirements.
		The minimum resources below are required to achieve the specified logging rate.
		• Up to 10,000 logs/sec (LPS):
		  • 16 CPUs
		  • Based on your logging needs:
		  • 2x2TB logging disks, 32GB memory
		  • 4x2TB logging disks, 64GB memory or more
		  • Manage up to 500 managed devices
		• Up to 20,000 log/sec (LPS)
		  • 32 CPUs
		  • 128GB memory
		  • 8x2TB logging disks
		  • Manage up to 1,000 managed devices
			The minimum resources below are required to achieve the specified logging rate.
			• Up to 15,000 log/sec (LPS)
			  • 16 CPUs
			  • Based on your logging needs:
			  • 2x2TB logging disks, 32GB memory
			  • 4x2TB logging disks, 64GB memory or more
			• Up to 25,000 logs/sec (LPS)
			  • 32 CPUs
			  • 128GB memory
			  • 8x2TB logging disks

		The first logging disk on the Panorama virtual appliance must be 2TB in order to add additional logging disks. If the first logging disk is smaller than 2TB, you are unable to add additional logging disks.	
Minimum CPUs and memory	• 16 CPUs • 32GB memory	The minimum resources below do not take LPS into consideration and are only required for the Panorama virtual appliance to function based on the number of logging disks added. Palo Alto Networks recommends you refer to the recommended resources above. For larger Panorama deployments, be aware that you may be under-provisioning your Panorama. This may lead to impacted performance and may cause Panorama to become unresponsive depending on the number of firewalls managed, the configuration size,	

 ©2025 Palo Alto Networks, Inc.

Requirement	Panorama Virtual Appliance in Management Only Mode	Panorama Virtual Appliance in Panorama Mode	Panorama Virtual Appliance in Log Collector Mode
		the number of administrators logged in to Panorama, and the volume of logs ingested. • **2TB to 8TB**—16 CPUs, 32GB memory • **10TB to 24TB**— 16 CPUs, 64GB memory	
Log storage capacity	Panorama in Management Only mode requires log forwarding to a Dedicated Log Collector.	2TB to 24TB	2TB to 24TB

Supported Interfaces

Interfaces can be used for device management, log collection, Collector Group communication, licensing and software updates. The Panorama virtual appliance supports up to six interfaces (MGT and Eth1 - Eth5).

Table 2: Supported interfaces for public hypervisors

Function	Alibaba Cloud	Amazon Web Services (AWS) and AWS GovCloud		Microsoft Azure	Google Cloud Platform (GCP)	OCI
Device Management	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported
Device Log Collection	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported
Collector Group Communication	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported
Licensing and Software Updates	MGT interface only	MGT interface only	MGT interface only	MGT interface only	MGT interface only	MGT interface only

 ©2025 Palo Alto Networks, Inc.

Table 3: Supported Interfaces for Private Hypervisors

Function	KVM	Hyper-V	VMware (ESXi, vCloud Air)
Device Management	Any interface supported	Any interface supported	Any interface supported
Device Log Collection	Any interface supported	Any interface supported	Any interface supported
Collector Group Communication	Any interface supported	Any interface supported	Any interface supported
Licensing and Software Updates	Any interface supported	Any interface supported	Any interface supported

Install the Panorama Virtual Appliance

Before installation, decide whether to run the virtual appliance in Panorama mode, Management Only mode, Log Collector mode, or Legacy mode (VMware only). Each mode has different resource requirements, as described in Setup Prerequisites for the Panorama Virtual Appliance. You must complete the prerequisites before starting the installation.

As a best practice, install the virtual appliance in Panorama mode to optimize log storage and report generation. For details on Panorama and Legacy mode, see Panorama Models.

- Install Panorama on VMware
- Set Up Panorama on Alibaba Cloud
- Install Panorama on AWS
- Install Panorama on AWS GovCloud
- Install Panorama on Azure
- Install Panorama on Google Cloud Platform
- Install Panorama on KVM
- Install Panorama on Hyper-V
- Set Up Panorama on Oracle Cloud Infrastructure (OCI)

Install Panorama on VMware

You can install the Panorama virtual appliance on the ESXi and vCloud Air VMware platforms.

- Install Panorama on an ESXi Server
- Install Panorama on vCloud Air
- Support for VMware Tools on the Panorama Virtual Appliance

 ©2025 Palo Alto Networks, Inc.

Install Panorama on an ESXi Server

Use these instructions to install a new Panorama virtual appliance on a VMware ESXi server. For upgrades to an existing Panorama virtual appliance, skip to Install Content and Software Updates for Panorama.

STEP 1 | Download the Panorama 10.2 base image Open Virtual Appliance (OVA) file.

1. Log in to the Palo Alto Networks Support Portal.
2. Select **Updates** > **Software Updates** and filter by **Panorama Base Images** to download the OVA file (`Panorama-ESX-10.2.0.ova`).

STEP 2 | Install Panorama.

1. Launch the VMware vSphere Client and connect to the VMware server.
2. Select **File** > **Deploy OVF Template**.
3. **Browse** to select the Panorama OVA file and click **Next**.
4. Confirm that the product name and description match the downloaded version, and click **Next**.
5. Enter a descriptive name for the Panorama virtual appliance, and click **Next**.
6. Select a datastore location (system disk) on which to install the Panorama image. See the Setup Prerequisites for the Panorama Virtual Appliance for the supported system disk sizes. After selecting the datastore, click **Next**.

   *The Panorama virtual appliance must be initially installed with the **Default** system disk size. Installing the Panorama virtual appliance with a system disk larger than the **Default** system disk size is not supported and may result in limited utilization. You have the option to increase the system disk size after initial installation*
7. Select **Thick Provision Lazy Zeroed** as the disk format, and click **Next**.
8. Specify which networks in the inventory to use for the Panorama virtual appliance, and click **Next**.
9. Confirm the selected options, click **Finish** to start the installation process, and click **Close** when it finishes. Do not power on the Panorama virtual appliance yet.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Configure resources on the Panorama virtual appliance.

1. Right-click the Panorama virtual appliance and **Edit Settings**.
2. In the **Hardware** settings, allocate the CPUs and memory as necessary.

   *The virtual appliance boots up in Panorama mode if you allocate sufficient **CPUs** and **Memory** and add a virtual logging disk (later in this procedure). Otherwise, the appliance boots up in Management Only mode. For details on the modes, see Panorama Models.*
3. Set the **SCSI Controller** to **LSI Logic Parallel**.
4. (Optional) Add a virtual logging disk.

   This step is required in the following scenarios:

   - *In Panorama mode to store logs on a dedicated logging disk.*
   - *Manage your SD-WAN deployment in Management Only mode.*

   1. **Add** a disk, select **Hard Disk** as the hardware type, and click **Next**.
   2. **Create a new virtual disk** and click **Next**.
   3. Set the **Disk Size** to exactly 2TB.

      In Panorama mode, you can later add additional logging disks (for a total of 12) with 2TB of storage each. Expanding the size of a logging disk that is already added to Panorama is not supported.
   4. Select your preferred **Disk Provisioning** disk format.

      Consider your business needs when selecting the disk provisioning format. For more information regarding the disk provisioning performance considerations, refer to the VMware Thick vs Thin Disks and All Flash Arrays document, or additional VMware documentation.

      *When adding multiple logging disks, it is a best practice to select the same **Disk Provisioning** format for all disks to avoid any unexpected performance issues that may arise.*
   5. Select **Specify a datastore or datastore structure** as the location, **Browse** to a datastore that has sufficient storage, click **OK**, and click **Next**.
   6. Select a SCSI **Virtual Device Node** (you can use the default selection) and click **Next**.

      Panorama will fail to boot if you select a format other than SCSI.
   7. Verify that the settings are correct and click **Finish**.
5. Click **OK** to save your changes.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Power on the Panorama virtual appliance.

1. In the vSphere Client, right-click the Panorama virtual appliance and select **Power** > **Power On**. Wait for Panorama to boot up before continuing.
2. Log in to the Panorama virtual appliance CLI from the ESXi console:
   1. Right-click the Panorama virtual appliance and select **Open Console**.
   2. Enter your username and password to log in (default is **admin** for both).

STEP 5 | Configure a new administrative password for the Panorama virtual appliance.

You must configure a unique administrative password before you can access the web interface or CLI of the Panorama virtual appliance. The new password must be a minimum of eight characters and include a minimum of one lowercase character, one uppercase character, and one number or special character.

When you first log in to the Panorama CLI, you are prompted to enter the **Old Password** and the **New Password** for the `admin` user before you can continue.

STEP 6 | Verify the Panorama is running the correct system mode.

```
admin> show system info
```

In the output, the `system-mode` indicates either `panorama` or `management-only` mode.

STEP 7 | Register the Panorama virtual appliance and activate the device management license and support licenses.

1. (VM Flex Licensing Only) Provisioning the Panorama Virtual Appliance Serial Number.

   When leveraging VM Flex licensing, this step is required to generate the Panorama virtual appliance serial number needed to register the Panorama virtual appliance with the Palo Alto Networks Customer Support Portal (CSP).
2. Register Panorama.

   You must register the Panorama virtual appliance using the serial number provided by Palo Alto Networks in the order fulfillment email.

   This step is not required when leveraging VM Flex licensing as the serial number is automatically registered with the CSP when generated.
3. Activate the firewall management license.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected.
4. Activate a Panorama Support License.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Increase the System Disk for Panorama on an ESXi Server if you intend to use the Panorama virtual appliance for the following:

- Manage your SD-WAN deployment in Panorama mode.
- Requires additional storage space for dynamic updates when managing large-scale firewall deployments.

STEP 9 | Complete configuring the Panorama virtual appliance for your deployment needs.

- For Panorama in Log Collector Mode.
  1. Add a Virtual Disk to Panorama on an ESXi Server as needed.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Log Collector mode.
  2. Begin at Step 6 to switch to Log Collector mode.

     *Enter the Public IP address of the Dedicated Log Collector when you add the Log Collector as a managed collector to the Panorama management server. You cannot specify the **IP Address**, **Netmask**, or **Gateway**.*
- For Panorama in Panorama mode.
  1. Add a Virtual Disk to Panorama on an ESXi Server.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Panorama mode.
  2. Set up a Panorama Virtual Appliance in Panorama Mode.
  3. Configure a Managed Collector.
- For Panorama in Management Only mode.
  1. Set up a Panorama Virtual Appliance in Management Only Mode.
  2. Configure a Managed Collector to add a Dedicated Log Collector to the Panorama virtual appliance.

     Management Only mode does not support local log collection, and requires a Dedicated Log Collector to store managed device logs.
- For SD-WAN deployments.
  1. Increase the System Disk for Panorama on an ESXi Server

     To leverage SD-WAN on Panorama deployed on ESXi, you must increase the system disk to 224GB.

     You cannot migrate back to a 81GB system disk after successfully increasing the system disk to 224GB.
  2. Set up a Panorama Virtual Appliance in Management Only Mode.
  3. Add a Virtual Disk to Panorama on an ESXi Server.

     To leverage SD-WAN, you must add a single 2TB logging disk to Panorama in Management Only mode.

 ©2025 Palo Alto Networks, Inc.

Install Panorama on vCloud Air

Use these instructions to install a new Panorama virtual appliance in VMware vCloud Air. If you are upgrading a Panorama virtual appliance deployed in vCloud Air, skip to Install Content and Software Updates for Panorama.

STEP 1 | Download the Panorama 10.2 base image Open Virtual Appliance (OVA) file.

1. Go to the Palo Alto Networks software downloads site. (If you can't log in, go to the Palo Alto Networks Customer Support web site for assistance.)
2. In the Download column in the Panorama Base Images section, download the Panorama 10.2 release OVA file (`Panorama-ESX-10.0.0.ova`).

STEP 2 | Import the Panorama image to the vCloud Air catalog.

For details on these steps, refer to the OVF Tool User's Guide.

1. Install the OVF Tool on your client system.
2. Access the client system CLI.
3. Navigate to the OVF Tool directory (for example, C:\Program Files\VMware\VMware OVF Tool).
4. Convert the OVA file to an OVF package:

```
ovftool.exe <OVA-file-pathname> <OVF-file-pathname>
```

5. Use a browser to access the vCloud Air web console, select your **Virtual Private Cloud OnDemand** location, and record the browser URL. You will use the URL information to complete the next step. The URL format is: **https://<*virtual-cloud-location*>.vchs.vmware.com/compute/cloud/org/<*vCloud-account-number*>/#/catalogVAppTemplateList?catalog=<*catalog-ID*>**.
6. Import the OVF package, using the information from the vCloud Air URL to complete the <virtual#cloud#location>, <vCloud#account#number>, and <catalog#ID> variables. The other variables are your vCloud Air username and domain <user>@<domain>, a virtual data center <datacenter>, and a vCloud Air template <template>.

```
ovftool.exe -st="OVF" "<OVF-file-pathname>"
 "vcloud://<user>@<domain>:password@<virtual-cloud-
location>.vchs.vmware.com?vdc=<datacenter>&org=<vCloud-
account-number>&vappTemplate=<template>.ovf&catalog=default-
catalog"
```

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Install Panorama.

1. Access the vCloud Air web console and select your **Virtual Private Cloud OnDemand** region.
2. Create a Panorama virtual machine. For the steps, refer to Add a Virtual Machine from a Template in the vCloud Air Documentation Center. Configure the **CPU**, **Memory** and **Storage** as follows:
   - Set the **CPU** and **Memory** based on whether the virtual appliance mode: see Setup Prerequisites for the Panorama Virtual Appliance.
   - Set the **Storage** to configure the Panorama virtual appliance system disk. See Setup Prerequisites for the Panorama Virtual Appliance for the supported disk sizes based on the Panorama virtual appliance mode. For better logging and reporting performance, select the **SSD-Accelerated** option.

     To increase the log storage capacity, you must Add a Virtual Disk to Panorama on vCloud Air. In Panorama mode, the virtual appliance does not use the system disk for log storage; you must add a virtual logging disk.

STEP 4 | Create vCloud Air NAT rules on the gateway to allow inbound and outbound traffic for the Panorama virtual appliance.

Refer to Add a NAT Rule in the vCloud Air Documentation Center for the detailed instructions:

1. Add a NAT rule that allows Panorama to receive traffic from the firewalls and allows administrators to access Panorama.
2. Add a NAT rule that allows Panorama to retrieve updates from the Palo Alto Networks update server and to access the firewalls.

STEP 5 | Create a vCloud Air firewall rule to allow inbound traffic on the Panorama virtual appliance.

Outbound traffic is allowed by default.

Refer to Add a Firewall Rule in the vCloud Air Documentation Center for the detailed instructions.

STEP 6 | Power on the Panorama virtual appliance if it isn't already on.

In the vCloud Air web console, select the **Virtual Machines** tab, select the Panorama virtual machine, and click **Power On**.

You are now ready to Perform Initial Configuration of the Panorama Virtual Appliance.

Support for VMware Tools on the Panorama Virtual Appliance

VMware Tools is bundled with the software image (ovf) for the Panorama virtual appliance. The support for VMware Tools allows you to use the vSphere environment—vCloud Director and vCenter server—for the following:

- View the IP address assigned to the Panorama management interface.
- View resource utilization metrics on hard disk, memory, and CPU. You can use these metrics to enable alarms or actions on the vCenter server or vCloud Director.
- Graceful shutdown and restart of Panorama using the power off function on the vCenter server or vCloud Director.

 ©2025 Palo Alto Networks, Inc.

- Enables a heartbeat mechanism between the vCenter server and Panorama for verifying that Panorama is functioning, or if the firewall/Panorama is rebooting. If the firewall goes into maintenance mode, heartbeats are disabled so that the vCenter server does not shut down the firewall. Disabling heartbeats allows the firewall to stay operational in maintenance mode when it cannot not send heartbeats to the vCenter server.

Set Up Panorama on Alibaba Cloud

Set up a Panorama™ virtual appliance on Alibaba Cloud to centrally managed the configuration of physical and VM-Series firewalls.

- Upload the Panorama Virtual Appliance Image to Alibaba Cloud
- Install Panorama on Alibaba Cloud

Upload the Panorama Virtual Appliance Image to Alibaba Cloud

Complete the following procedure to upload a Panorama™ management server qcow2 file for KVM and create a custom image that you need to launch the Panorama virtual appliance. Uploading and creating the image is required only once. You can use the same image for all subsequent deployments of the Panorama virtual appliance.

STEP 1 | Download the Panorama qcow2 file for KVM from the Palo Alto Networks Customer Support Portal (CSP).

1. Log in to the Palo Alto Networks CSP.
2. Select **Updates** > **Software Updates** and select **Panorama Base Images** from the software updates filter drop-down.
3. Download the latest version of the `Panorama-KVM` qcow2 file.

STEP 2 | Log in to the Alibaba Cloud Console.

STEP 3 | Create an Object Storage Service (OSS) bucket for the Panorama virtual appliance image.

1. From the Alibaba Cloud menu, select **Object Storage Service** > **Buckets** and **Create Bucket**.
2. Enter a descriptive **Bucket Name**.
3. Select the bucket **Region**.

   This region must be in the same region you plan on deploying your Panorama virtual appliance and in the same region as the firewalls you plan to manage with Panorama.
4. Configure the remaining OSS bucket settings as needed.
5. Click **OK**.

   You are automatically taken to the OSS bucket Overview page after successful creation.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Upload the qcow2 file to the OSS bucket.

1. In the OSS bucket Overview, select **Files** and **Upload** the qcow2 file you downloaded in the previous step.
2. For **Upload To** target, select **Current**.
3. For the **File ACL**, select **Inherited from Bucket**.
4. Click **Select Files** and select the qcow2 file.

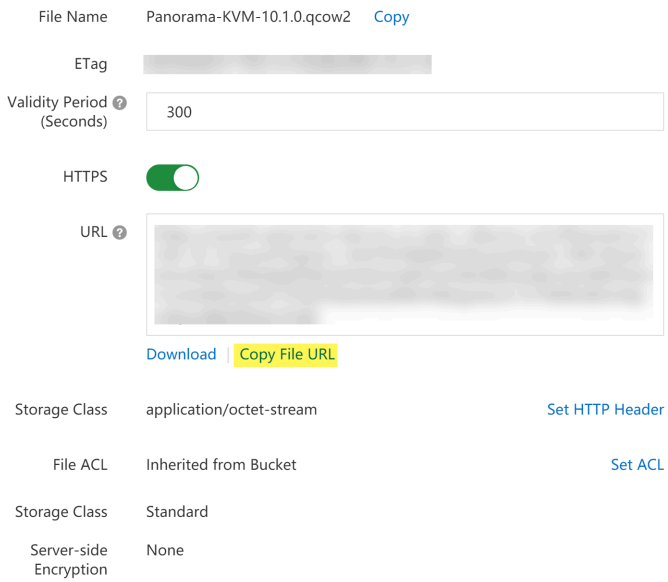
   Alternatively, you can drag and drop the qcow2 file into the **Files to Upload** section.
5. **Upload** the qcow2 file.

   A Task List window appears displaying the upload Status. Continue to the next step after the qcow2 file upload **Status** displays **Uploaded**.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Make the qcow2 file a bootable image.

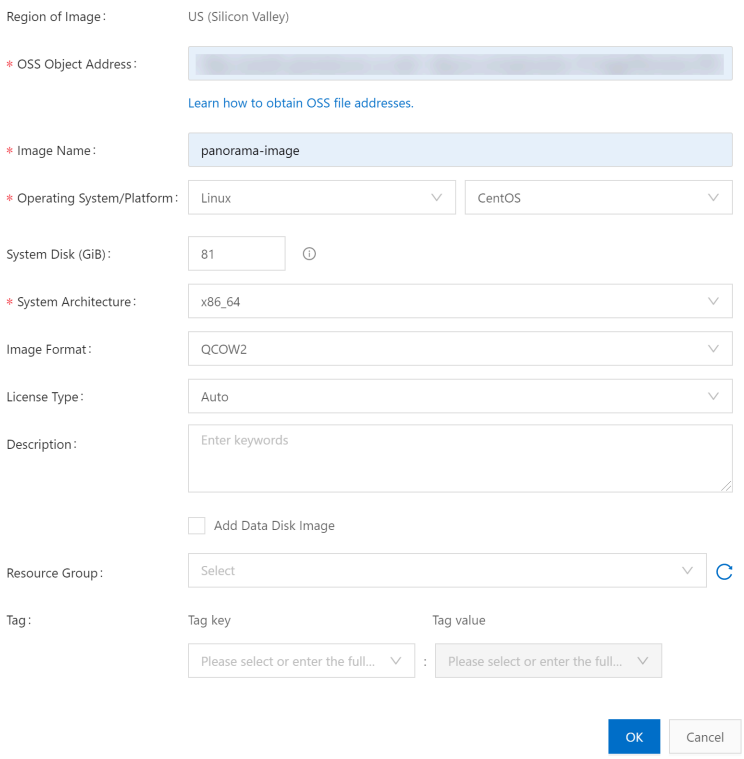
1. In the OSS bucket Overview, select **Files** and click the qcow2 file you uploaded to view the file Details.
2. Click **Copy File URL** and exit the file Details.



3. From the Alibaba Cloud menu, select **Elastic Compute Service** > **Instances & Images** > **Images** and **Import Image**.
4. Paste the **OSS Object Address** for the qcow2 file.

   This the file URL you copied in the previous step.
5. Enter an **Image Name**.
6. For the **Operating System/Platform**, select **Linux CentOS**.
7. For the **System Disk (GiB)**, enter **81**.
8. For the **System Architecture**, select **x86_64**.
9. For the **Image Format**, select **QCOW2**.
10. Click **OK**.

 ©2025 Palo Alto Networks, Inc.



Install Panorama on Alibaba Cloud

Use the Elastic Compute Service (ECS) to create a Panorama™ virtual appliance instance on Alibaba Cloud. An ECS instance supports a single NIC by default and automatically attached an Elastic Network Interface (ENI) to it. You must manually upload a Panorama virtual appliance qcow2 image downloaded from the Palo Alto Networks Customer Supported Portal (CSP) to Alibaba Cloud to successfully install the Panorama virtual appliance on Alibaba Cloud.

A Panorama virtual appliance deployed on Alibaba Cloud is Bring Your Own License (BYOL), supports all deployment modes (Panorama, Log Collector, and Management Only), and shares the same processes and functionality as the M-Series hardware appliances. For more information on Panorama modes, see Panorama Models.

Review the Setup Prerequisites for the Panorama Virtual Appliance to determine the correct Elastic Computer Service (ECS) instance type for your needs. The virtual resources requirement for the Panorama virtual appliance is based on the total number of firewalls managed by the Panorama virtual appliance and the required Logs Per Second (LPS) for forwarding logs from your managed firewalls to your Log Collector.

Palo Alto Networks supports the following instance types.

- `ecs.g5.xlarge`, `ecs.g5.2xlarge`, `ecs.g5.4xlarge`
- `ecs.sn2ne.xlarge`, `ecs.sn2ne.2xlarge`, `ecs.sn2ne.4xlarge`

Under-provisioning the Panorama virtual appliance will impact management performance. This includes the Panorama virtual appliance becoming slow or unresponsive depending on how under-provisioning the Panorama virtual appliance is.

STEP 1 | Log in to the Alibaba Cloud Console.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Upload the Panorama Virtual Appliance Image to Alibaba Cloud.

STEP 3 | Set up the virtual private cloud (VPC) for your network needs.

Whether you launch the Panorama virtual appliance in an existing VPC or you create a new VPC, the Panorama virtual appliance must be able to receive traffic from other instances in the VPC and perform inbound and outbound communication between the VPC and the internet as needed.

Refer to the Alibaba Cloud VPC documentation for more information.

1. Create a VPC and Configure Networks or use an existing VPC.
2. Verify that the network and security components are appropriately defined.
   - Create an internet gateway to enable internet access to the subnet of your Panorama virtual appliance. Internet access is required to install software and content updates, activate licenses, and leverage Palo Alto Networks cloud services. Otherwise, you must manually install updates and activate licenses.
   - Create subnets. Subnets are segments of the IP address range assigned to the VPC in which you can launch Alibaba Cloud instances. It is recommended that the Panorama virtual appliance belong to the management subnet so that you can configure it to access the internet if needed.
   - Add routes to the route table for a private subnet to ensure traffic can be routed across subnets in the VPC and from the internet if applicable.

     Ensure you create routes between subnets to allow communication between:

     - Panorama, managed firewalls, and Log Collectors.
     - (Optional) Panorama and the internet.
   - Ensure that the following ingress security rules are allowed for the VPC to manage VPC traffic. The ingress traffic source for each rule is unique to your deployment topology.

     See Ports Used for Panorama for more information.

     - Allow SSH (port **22**) traffic to enable access to the Panorama CLI.
     - Allow HTTPS (port **443** and **27280**) traffic to enable access to the Panorama web interface.
     - Allow traffic on port **3978** to enable communication between Panorama, manage firewalls, and managed Log Collectors. This port is also used by Log Collectors to forward logs to Panorama.
     - Allow traffic on port **28443** to enable managed firewalls to get software and content updates from Panorama.

STEP 4 | Select **Elastic Compute Service** > **Instances & Images** > **Instances** and click **Create Instance** in the upper right corner.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Create the Panorama virtual appliance instance.

1. Select **Custom Launch**.
2. Configure the Panorama virtual appliance instance.
   - **Billing Method**—Select the desired subscription method for the instance.
   - **Region**—Select a region of your choice. The region you select must provide on of the supported instance types.
   - **Instance Type**—Select one of the supported instance types. You can select Type-based Selection to search for the instance type.
   - **Image**— Select **Custom Image** and select the Panorama virtual appliance image you uploaded.
   - **Storage**—Choose a disk type and enter **81**GiB as the system disk capacity.
   - (Optional) **Add Disk**—Add additional logging disks.

     If you intend to use the Panorama virtual appliance in Panorama mode or as a Dedicated Log Collector, add the virtual logging disks during the initial deployment. By default, the Panorama virtual appliance is in Panorama mode for the initial deployment when you meet the Panorama mode resource requirements and have added at least one virtual logging disk. Otherwise, the Panorama virtual appliance defaults to Management Only mode. Change the Panorama virtual appliance to Management Only mode if you just want to manage devices and Dedicated Log Collectors, and to not collect logs locally.

     The Panorama virtual appliance on Alibaba Cloud only supports 2TB logging disks, and in total supports up to 24TB of log storage. You are unable to add a logging disk smaller than 2TB, or a logging disk with a size not divisible by the 2TB logging disk requirement. The Panorama virtual appliance partitions logging disks larger than 2TB into 2TB partitions.
   - (Optional) **Snapshot**—Specify how often a snapshot is automatically taken of the Panorama virtual appliance instance to prevent risks and accidental data deletion.
   - **Duration**—Specify the duration for the Panorama virtual appliance instance.

STEP 6 | Configure the Panorama virtual appliance network settings.

1. Select **Next: Networking**.
2. Configure the network settings for the Panorama virtual appliance instance.
   - **Network Type**—Select the VPC and management VSwitch you created.
   - **Public IP Address**—If you do not have a public IP address, enable (check) **Assign Public IPv4 Address** and a public IPv4 address is automatically assigned to the Panorama virtual appliance instance.

     If you must use a specific IP address, or an address in a specific range, you can request a custom IP address. Refer to the Elastic IP Address User Guide.
   - **Security Group**—Select the management security group you created and enable **Port 443 (HTTPS)**, **Port 22**, and **Port 3389**.
   - **Elastic Network Interface**—No configuration needed. The Management interface is already attached to eth0.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Configure the Panorama virtual appliance instance system settings.

1. Select **Next: System Configurations**.
2. Configure system settings for the Panorama virtual appliance instance.
   - **Logon Credentials**—Select **Key Pair** and select the key pair. If a key pair has not already been created, select **Create Key Pair** to create an new key pair on Alibaba Cloud or import an existing key pair.

     Password authentication is not supported.
   - **Instance Name**—Enter a descriptive name for the Panorama virtual appliance. This the name displayed for the instance throughout the Alibaba Cloud Console.
   - **Host**—Enter a hostname for the Panorama virtual appliance instance.

STEP 8 | (Optional) Select **Next: Grouping** to configuring grouping for all Alibaba Cloud resources associated with the Panorama virtual appliance instance.

STEP 9 | Select **Preview** to view the configuration before ordering.

STEP 10 | View and check the **ECS Terms of Service** and **Product Terms of Service**.

STEP 11 | **Create Instance** to create the Panorama virtual appliance instance.

When prompted, click **Console** to view the instance creation status.

STEP 12 | Allocate Elastic IP (EIP) addresses.

The EIP is a public IP address used to connect to the Panorama virtual appliance.

This step is required only if you want to enable internet access for the Panorama virtual appliance.

1. Select **Elastic Compute Service** > **Network & Security** > **VPC** > **Elastic IP Addresses** > **Elastic IP Addresses**.

   Select **Create EIP** if you do not have any existing EIPs.
2. In the **Actions** column, select **Bind Resource** to bind an EIP to any interface exposed to the Internet.

 ©2025 Palo Alto Networks, Inc.

STEP 13 | Log in to the Panorama CLI using the SSH to configure the Panorama virtual appliance network settings.

You must configure the `admin` password, system IP address , netmask, and default gateway. Additionally, you must add the Alibaba Cloud DNS servers to successfully connect to the Palo Alto Networks update server.

*You can also access the Panorama CLI from the Alibaba console. To access the Panorama CLI from the Alibaba console, select **Elastic Compute Service** > **Instances & Images** > **Instances** and select the Panorama virtual appliance instance. In the Instance Details, select **Connect**.*

You are prompted to create a VCN password for the Panorama virtual appliance instance on first connection from the Alibaba VCN. Be sure to save this password as it cannot be recovered and is required to connect using the VCN or update the password in the future.

STEP 14 | Configure a new administrative password for the Panorama virtual appliance.

You must configure a unique administrative password before you can access the web interface or CLI of the Panorama virtual appliance. To access the CLI, the private key used to launch the Panorama virtual appliance is required.

The new password must be a minimum of eight characters and include a minimum of one lowercase character, one uppercase character, and one number or special character.

Configure a new password using the following commands and follow the on screen prompts:

```
admin> configure
admin# set mgt-config users admin password
```

 ©2025 Palo Alto Networks, Inc.

STEP 15 | Configure the initial network settings for the Panorama virtual appliance.

```
admin> configure
```

```
admin#  set deviceconfig system type static
```

```
admin#  set deviceconfig system ip-address <instance-private-IP
 address> netmask <netmask> default-gateway <default-gateway-IP>
```

The default gateway on Alibaba Cloud ends in ***.253****. For example, if the private IP address for your Panorama virtual appliance instance is 192.168.100.20, the default gateway is 192.168.100.253.*

```
admin#  set deviceconfig system dns-setting servers primary
 100.100.2.136
```

```
admin#  set deviceconfig system dns-setting servers secondary
 100.100.2.138
```

```
admin# commit
```

STEP 16 | Register the Panorama virtual appliance and activate the device management license and support licenses.

1. (VM Flex Licensing Only) Provisioning the Panorama Virtual Appliance Serial Number.

   When leveraging VM Flex licensing, this step is required to generate the Panorama virtual appliance serial number needed to register the Panorama virtual appliance with the Palo Alto Networks Customer Support Portal (CSP).
2. Register Panorama.

   You must register the Panorama virtual appliance using the serial number provided by Palo Alto Networks in the order fulfillment email.

   This step is not required when leveraging VM Flex licensing as the serial number is automatically registered with the CSP when generated.
3. Activate the firewall management license.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected.
4. Activate a Panorama Support License.

 ©2025 Palo Alto Networks, Inc.

STEP 17 | Complete configuring the Panorama virtual appliance for your deployment needs.

- (Management Only mode) Set up a Panorama Virtual Appliance in Management Only Mode.
- (Log Collector mode) Begin at Step 6 to Switch from Panorama mode to Log Collector mode.

  *Enter the Public IP address of the Dedicated Log Collector when you Add the Log Collector as a managed collector to the Panorama management server. You cannot specify the **IP Address**, **Netmask**, or **Gateway**.*
- (Panorama and Management Only mode) Configure a Managed Collector to add a Dedicated Log Collector to the Panorama virtual appliance. Management Only mode does not support local log collection, and requires a Dedicated Log Collector to store managed device logs.

STEP 18 | Complete configuring the Panorama virtual appliance for your deployment needs.

- For Panorama in Log Collector Mode.
  1. Add a Virtual Disk to Panorama on Alibaba Cloud as needed.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Log Collector mode.
  2. Begin at Step 6 to switch to Log Collector mode.

     *Enter the Public IP address of the Dedicated Log Collector when you add the Log Collector as a managed collector to the Panorama management server. You cannot specify the **IP Address**, **Netmask**, or **Gateway**.*
- For Panorama in Panorama mode.
  1. Add a Virtual Disk to Panorama on Alibaba Cloud as needed.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Panorama mode.
  2. Set up a Panorama Virtual Appliance in Panorama Mode.
  3. Configure a Managed Collector.
- For Panorama in Management Only mode.
  1. Set up a Panorama Virtual Appliance in Management Only Mode.
  2. Configure a Managed Collector to add a Dedicated Log Collector to the Panorama virtual appliance.

     Management Only mode does not support local log collection, and requires a Dedicated Log Collector to store managed device logs.

Install Panorama on AWS

You can now deploy Panorama™ and a Dedicated Log Collector on Amazon Web Services (AWS). Panorama deployed on AWS is Bring Your Own License (BYOL), supports all deployment modes (Panorama, Log Collector, and Management Only), and shares the same processes and functionality as the M-Series hardware appliances. For more information on Panorama modes, see Panorama Models.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Log in to AWS Web Service console and select the EC2 Dashboard.

- Amazon Web Service Console
- AWS GovCloud Web Service Console

STEP 2 | Set up the virtual private cloud (VPC) for your network needs.

Whether you launch the Panorama virtual appliance in an existing VPC or you create a new VPC, the Panorama virtual appliance must be able to receive traffic from other instances in the VPC and perform inbound and outbound communication between the VPC and the internet as needed.

Refer to the AWS VPC documentation for instructions on creating a VPC and setting it up for access.

1. Create a new VPC or use an existing VPC. Refer to the AWS Getting Started documentation
2. Verify that the network and security components are appropriately defined.
   - Create an internet gateway to enable internet access to the subnet of your Panorama virtual appliance. Internet access is required to install software and content updates, activate licenses, and leverage Palo Alto Networks cloud services. Otherwise, you must manually install updates and activate licenses.
   - Create subnets. Subnets are segments of the IP address range assigned to the VPC in which you can launch AWS instances. It is recommended that the Panorama virtual appliance belong to the management subnet so that you can configure it to access the internet if needed.
   - Add routes to the route table for a private subnet to ensure traffic can be routed across subnets in the VPC and from the internet if applicable.

     Ensure you create routes between subnets to allow communication between:

     - Panorama, managed firewalls, and Log Collectors.
     - (Optional) Panorama and the internet.
   - Ensure that the following inbound security rules are allowed for the VPC to manage VPC traffic. The ingress traffic source for each rule is unique to your deployment topology.

     See Ports Used for Panorama for more information.

     - Allow SSH (port **22**) traffic to enable access to the Panorama CLI.
     - Allow HTTPS (port **443**) traffic to enable access to the Panorama web interface.
     - Allow traffic on port **3978** to enable communication between Panorama, manage firewalls, and managed Log Collectors. This port is also used by Log Collectors to forward logs to Panorama.
     - Allow traffic on port **28443** to enable managed firewalls to get software and content updates from Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Deploy Panorama on Amazon Web Services.

1. Select **Services** > **EC2** > **Instances** and **Launch Instance**.
2. Select **AWS Marketplace**, search for **`Palo Alto Networks Panorama`**, and **Select** the Panorama AMI and **Continue**.
3. Choose the **EC2 instance type** for allocating the resources required for the Panorama virtual appliance, and click **Next: Configure Instance Details**. Review the Setup Prerequisites for the Panorama Virtual Appliance for resource requirements.

   If you plan to use the Panorama virtual appliance as a Dedicated Log Collector, ensure that you configure the appliance with the required resources during initial deployment. The Panorama virtual appliance does not remain in Log Collector mode if you resize the virtual machine after you deploy it, and this results in a loss of log data.
4. Configure the instance details.
   1. Select **Next: Configure Instance Details**.
   2. For the **Network**, select the VPC.
   3. Select the **Subnet**.
   4. To **Auto-assign Public IP** select **Enable.**

      This IP must be accessible by the firewalls you plan to manage using Panorama. This allows you to obtain a publicly accessible IP address for the management interface of the Panorama virtual appliance. You can later attach an Elastic IP address to the management interface. Unlike the public IP address that is disassociated from the virtual appliance when the instance is terminated, the Elastic IP address provides persistence and you can the IP address to a new (or replacement) instance of the Panorama virtual appliance without the need to reconfigure the IP address whenever the Panorama virtual appliance instance is powered off.
   5. Configure any additional instance details as needed.
5. (Optional) Configure the Panorama virtual appliance storage.
   1. Select **Next: Add Storage**.
   2. **Add New Volume** to add additional log storage.

      (SD-WAN only) If you plan on managing your SD-WAN deployment in Management Only mode, you must add a 2TB logging disk.

      If you intend to use the Panorama virtual appliance in Panorama mode or as a Dedicated Log Collector, add the virtual logging disks during the initial deployment. By default, the Panorama virtual appliance is in Panorama mode for the initial deployment when you meet the Panorama mode resource requirements and have added at least one virtual logging disk. Otherwise, the Panorama virtual appliance defaults to Management Only mode. Change the Panorama virtual appliance to Management Only mode if you just want to manage devices and Dedicated Log Collectors, and to not collect logs locally.

      The Panorama virtual appliance on AWS only supports 2TB logging disks, and in total supports up to 24TB of log storage. You are unable to add a logging disk smaller than 2TB, or a logging disk with a size not divisible by the 2TB logging disk requirement.

 ©2025 Palo Alto Networks, Inc.

The Panorama virtual appliance partitions logging disks larger than 2TB into 2TB partitions.

6. (Optional) Select **Next: Add Tags** and add one or more tags as metadata to help you identify and group the Panorama virtual appliance. For example, add a **Name** tag with a **Value** that helps you identify which firewalls the Panorama virtual appliance manages.
7. Configure the instance security group.
   1. Select **Next: Configure Security Group**.
   2. **Select an existing security group** to assign a security group for the Panorama virtual appliance instance.
   3. Select the security group you previously created.

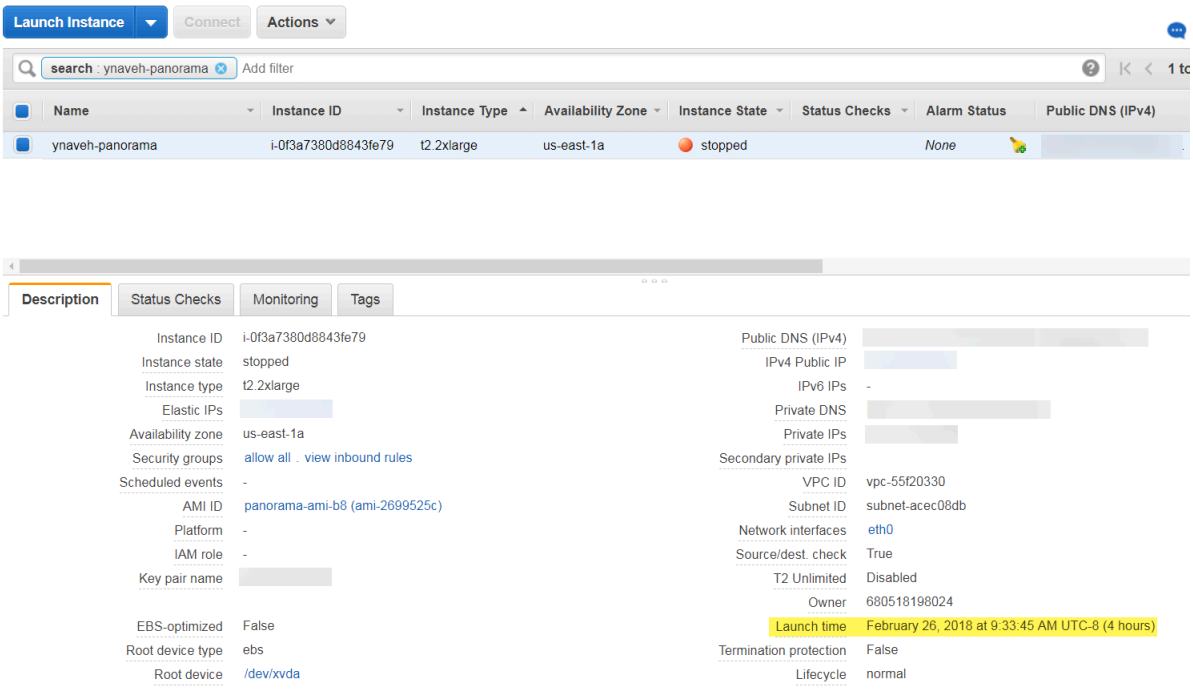
      You can select the `default` security group to allow all inbound and outbound traffic types.
8. **Review and Launch** the Panorama virtual appliance instance to verify that your selections are accurate before you **Launch.**
9. Select an existing key pair or create a new one and acknowledge the disclaimer.

   If you created a new key from AWS, download and save the key to a safe location. The file extension is `.pem`. You must load the public key into PuTTYgen *and save it in `.ppk` format. You cannot regenerate this key if lost.*

It takes about 30 minutes to finish deploying the Panorama virtual appliance after you launch it on AWS. Deploying the Panorama virtual appliance may take longer depending

 ©2025 Palo Alto Networks, Inc.

on the number and size of the disks attached to the instance. View the Launch Time by selecting the Panorama virtual appliance instance (**Instances**).



If you plan to use the Panorama virtual appliance as a Dedicated Log Collector, ensure that you provision the appliance with the required resources. The Panorama virtual appliance does not remain in Log Collector mode if you resize the virtual machine after you deploy it and this results in a loss of log data.

STEP 4 | Shut down the Panorama virtual appliance.

1. On the EC2 Dashboard, select **Instances**.
2. Select the Panorama virtual appliance and click **Instance State** > **Stop Instance**.

STEP 5 | Create or assign an Elastic IP (EIP) address to the management interface.

1. Select **Services** > **EC2** > **Elastic IPs** and **Allocate Elastic IP address**.
2. Select a **Network Border Group** to specify the logical group of zones from where the public IP4v address is advertised.
3. For the Public IPv4 address pool, select **Amazon's pool of IPv4 addresses**.
4. **Allocate** the EIP.
5. Click the IPv4 address in the Allocated IPv4 address column and **Associate Elastic IP address**.
6. Select the Panorama virtual appliance **Instance.**
7. Select the Panorama virtual appliance **Private IP address** to with which to associate the EIP.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Power on the Panorama virtual appliance.

1. On the EC2 Dashboard, select **Instance**.
2. From the list, select the Panorama virtual appliance and click **Actions** > **Instance State** > **Start**.

STEP 7 | Configure a new administrative password for the Panorama virtual appliance.

You must configure a unique administrative password before you can access the web interface of the Panorama virtual appliance. To access the CLI, the private key used to launch the Panorama virtual appliance is required.

The new password must be a minimum of eight characters and include a minimum of one lowercase character, one uppercase character, and one number or special character.

- If you have an SSH service installed on your computer:

1. Enter the following command to log into the Panorama virtual appliance:

```
ssh -i <private_key.ppk> admin@<public-ip_address>
```

2. Configure a new password using the following commands and follow the on screen prompts:

```
admin> configure
```

 ©2025 Palo Alto Networks, Inc.

```
admin# set mgt-config users admin password
```

3. If you need to activate a BYOL, set the DNS server IP address so that the Panorama virtual appliance can access the Palo Alto Networks licensing server. Enter the following command to set the DNS server IP address:

```
admin# set deviceconfig system dns-setting servers
 primary <ip_address>
```

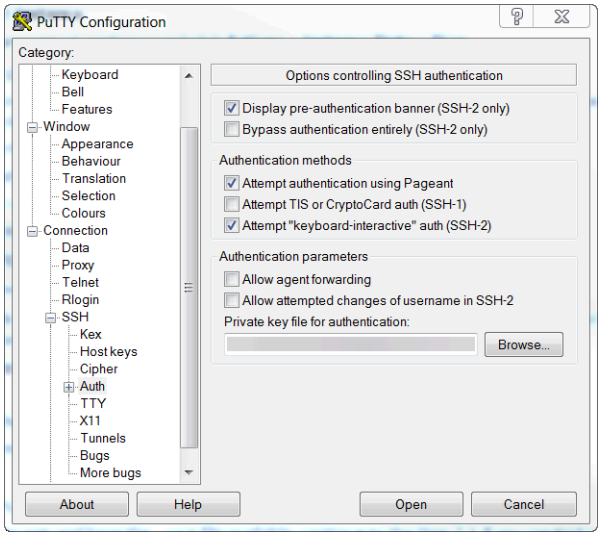
4. Commit your changes with the command:

```
admin# commit
```

5. Terminate the SSH session.

- If you are using PuTTY to SSH into the Panorama virtual appliance:

1. If you are using an existing key pair and have the `.ppk` file available, continue to the Step 7.3. If you created a new key pair or have only the `.pem` file of the existing key pair, open PuTTYgen and **Load** the `.pem` file.
2. **Save the private key** to a local accessible destination.
3. Open PuTTY and select **SSH** > **Auth** and then **Browse** to the `.ppk` file you saved in the previous step.



4. Select **Sessions** and enter the public IP address of the Panorama virtual appliance. Click **Open** and click **Yes** when the security prompt appears.
5. Log in as admin when prompted.
6. Configure a new password using the following commands and follow the onscreen prompts:

```
admin> configure
```

 ©2025 Palo Alto Networks, Inc.

```
admin# set mgt-config users admin password
```

7. Set the DNS server IP address so that the Panorama virtual appliance can access the Palo Alto Networks licensing server. Enter the following command to set the DNS server IP address:

```
admin# set deviceconfig system dns-setting servers
 primary <ip_address>
```

8. Commit your changes with the command:

```
admin# commit
```

9. Terminate the SSH session.

STEP 8 | Register the Panorama virtual appliance and activate the device management license and support licenses.

1. (VM Flex Licensing Only) Provisioning the Panorama Virtual Appliance Serial Number.

   When leveraging VM Flex licensing, this step is required to generate the Panorama virtual appliance serial number needed to register the Panorama virtual appliance with the Palo Alto Networks Customer Support Portal (CSP).

2. Register Panorama.

   You must register the Panorama virtual appliance using the serial number provided by Palo Alto Networks in the order fulfillment email.

   This step is not required when leveraging VM Flex licensing as the serial number is automatically registered with the CSP when generated.

3. Activate the firewall management license.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected.
4. Activate a Panorama Support License.

 ©2025 Palo Alto Networks, Inc.

STEP 9 | Complete configuring the Panorama virtual appliance for your deployment needs.

- For Panorama in Log Collector Mode.
  1. Add a Virtual Disk to Panorama on AWS as needed.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Log Collector mode.
  2. Begin at Step 6 to switch to Log Collector mode.

     *Enter the Public IP address of the Dedicated Log Collector when you add the Log Collector as a managed collector to the Panorama management server. You cannot specify the **IP Address**, **Netmask**, or **Gateway**.*
- For Panorama in Panorama mode.
  1. Add a Virtual Disk to Panorama on AWS.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Panorama mode.
  2. Set up a Panorama Virtual Appliance in Panorama Mode.
  3. Configure a Managed Collector.
- For Panorama in Management Only mode.
  1. Set up a Panorama Virtual Appliance in Management Only Mode.
  2. Configure a Managed Collector to add a Dedicated Log Collector to the Panorama virtual appliance.

     Management Only mode does not support local log collection, and requires a Dedicated Log Collector to store managed device logs.

Install Panorama on AWS GovCloud

You can now deploy Panorama™ and a Dedicated Log Collector on Amazon Web Services (AWS) GovCloud. AWS GovCloud is an isolated AWS region that meets the regulatory and compliance requirements of the US government agencies and customers. Panorama deployed on AWS GovCloud is Bring Your Own License (BYOL), supports all deployment modes (Panorama, Log Collector, and Management Only). For more information on Panorama modes, see Panorama Models.

To secure your workloads that contain all categories of Controlled Unclassified Information (CUI) data and government-oriented, publicly available data in the AWS GovCloud (US) region, the Panorama virtual appliance provides the same security features offered in the standard AWS public cloud on AWS GovCloud. The Panorama virtual appliance on AWS GovCloud and the standard AWS public cloud support the same features and capabilities.

Review the Setup Prerequisites for the Panorama Virtual Appliance to review the supported EC2 instance types. Once you are ready, refer to Install Panorama on AWS to install the Panorama virtual appliance on AWS GovCloud.

See the following procedures to add additional logging storage to your Panorama virtual appliance, or to increase the allocated CPU cores and memory:

- Add a Virtual Disk to Panorama on AWS

 ©2025 Palo Alto Networks, Inc.

- Increase CPUs and Memory for Panorama on AWS

Install Panorama on Azure

You can now deploy Panorama™ and a Dedicated Log Collector on Microsoft Azure. Panorama deployed on Azure is Bring Your Own License (BYOL), supports all deployment modes (Panorama, Log Collector, and Management Only), and shares the same processes and functionality as the M-Series hardware appliances. For more information on Panorama modes, see Panorama Models.

STEP 1 | Log into to the Microsoft Azure portal.

STEP 2 | Set up the virtual network for your network needs.

Whether you launch the Panorama virtual appliance in an existing virtual network or you create a new virtual network, the Panorama virtual appliance must be able to receive traffic from other instances in the virtual network and perform inbound and outbound communication between the virtual network and the internet as needed.

Refer to the Microsft Azure Virtual Network documentation for more information.

1. Create a Virtual Network or use an existing virtual network.
2. Verify that the network and security components are appropriately defined.
   - Create a NAT gateway if you want to enable only outbound internet access for the subnet to which the Panorama virtual appliance belongs.
   - Create subnets. Subnets are segments of the IP address range assigned to the VNet in which you can launch Microsoft Azure instances. It is recommended that the Panorama virtual appliance belong to the management subnet so that you can configure it to access the internet if needed.
   - Add routes to the route table for a private subnet to ensure traffic can be routed across subnets in the VNet and from the internet if applicable.

     Ensure you create routes between subnets to allow communication between:

     - Panorama, managed firewalls, and Log Collectors.
     - (Optional) Panorama and the internet.
   - Ensure that the following ingress security rules are allowed for the VNet to manage VNet traffic. The ingress traffic source for each rule is unique to your deployment topology.

     See Ports Used for Panorama for more information.

     - Allow SSH (port **22**) traffic to enable access to the Panorama CLI.
     - Allow HTTPS (port **443**) traffic to enable access to the Panorama web interface.
     - Allow traffic on port **3978** to enable communication between Panorama, manage firewalls, and managed Log Collectors. This port is also used by Log Collectors to forward logs to Panorama.
     - Allow traffic on port **28443** to enable managed firewalls to get software and content updates from Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Deploy the Panorama virtual appliance.

1. In the Azure Dashboard, select **Virtual machines** and **Add** a new virtual machine.
2. Search for Palo Alto Networks and select the latest Panorama virtual appliance image.
3. **Create** the Panorama virtual appliance.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the Panorama virtual appliance.

1. Select your Azure **Subscription**.
2. Select the Azure **Resource Group** to contain all your Azure instance resources.
3. Enter a **Virtual machine name** for the Panorama virtual appliance.
4. Select the **Region** for the Panorama virtual appliance to be deployed in.
5. (Optional) Select the **Availability options**. See How to use availability sets for more information.
6. Select the **Image** used to deploy the Panorama management server. **Browse all public and private images** to deploy the Panorama management server from the Panorama image on the Azure marketplace.
7. Configure the Panorama virtual appliance size. Review the Setup Prerequisites for the Panorama Virtual Appliance for sizing requirements.

   If you plan to use the Panorama virtual appliance as a Dedicated Log Collector, ensure that you configure the appliance with the required resources during initial deployment. The Panorama virtual appliance does not remain in Log Collector mode if you resize the virtual machine after you deploy it, and this results in a loss of log data.
8. Configure the unique Panorama virtual appliance administrator credentials.

   You must configure a unique administrative password before you can access the web interface and CLI of the Panorama virtual appliance.

   1. Enter a **Username** for the Panorama virtual appliance administrator. To ensure that your username is secure, admin is not a valid entry.
   2. Enter a **Password** or copy and paste an **SSH public key** for securing administrative access to the Panorama virtual appliance.

      You must enable SSH key authentication if you plan to use this instance of the Panorama virtual appliance in FIPS-CC operational mode. Although you can deploy the Panorama virtual appliance using a username and password, you will be unable to authenticate using the username and password after changing the operational mode to FIPS-CC. After resetting to FIPS-CC mode, you must use the SSH key to log in and can then configure a username and password that you can use for subsequently logging in to the Panorama web interface. For details on creating the SSH key, refer to the Azure documentation.
9. Configure the Panorama virtual appliance instance **Networking**

   1. Select an existing **Virtual network** or create a new virtual network.
   2. Configure the **Subnet**. The subnet is dependent on the virtual network you selected or created in the previous step. If you selected an existing virtual network, you can choose one of the subnets for the selected virtual network.
   3. Select an existing **Public IP address** or create a new one. This creates the management interface used to access your Panorama virtual appliance.
   4. Select an existing **NIC network security group** or create a new security group. Network security groups control traffic to the virtual machine. Make sure that HTTPS and SSH are allowed for the Inbound rules.

 ©2025 Palo Alto Networks, Inc.

10. Configure the instance **Management** settings.
    1. Select whether to enable **Auto-shutdown**. Auto-shutdown allows you to configure a daily time to automatically shut down the virtual machine that you disable auto-shutdown to avoid the possibility that a new public IP address gets assigned to the virtual machine, that logs are dropped, that logs are not or that you are unable to manage your firewalls while the Panorama virtual appliance is shut down.
    2. Select whether to enable boot **Monitoring**. Select the Diagnostic storage account if enabled. Monitoring automatically sends boot-up diagnostic logs to your Diagnostics storage account. For more information, see Overview of Monitoring in Microsoft Azure.
    3. Configure any other settings as needed.
11. Review the summary, accept the terms of use and privacy policy, and **Create** the Panorama virtual appliance.

STEP 5 | Verify that you the Panorama virtual appliance has been successfully deployed.

1. Select **Dashboard** > **Resource Groups** and select the resource group containing the Panorama virtual appliance.
2. Under Settings, select **Deployments** for the virtual machine deployment status.

It takes about 30 minutes to deploy the Panorama virtual appliance. Launching the Panorama virtual appliance may take longer depending on the resources configured for the virtual machine. Microsoft Azure does not permit the ICMP protocol to test whether it deployed successfully.

If you plan to use the Panorama virtual appliance as a Dedicated Log Collector, ensure that you correctly configured the appliance the required resources. The Panorama virtual appliance does not remain in Log Collector mode if you resize the virtual machine after you deploy it and this results in a loss of log data.

STEP 6 | Configure a static Public IP address.

1. On the Azure portal, select **Virtual machines** and select the Panorama virtual machine.
2. Select **Overview** and click the **Public IP address**.
3. Under Assignment, select **Static** and **Save** the new IP address configuration.

STEP 7 | Log in to the web interface of the Panorama virtual appliance.

1. On the Azure portal, in **All Resources**, select the Panorama virtual appliance and view the public IP address located in the Overview section.
2. Use a secure (https) connection from your web browser to log in to the Panorama virtual appliance using the public IP address.
3. Enter the username and password of the Panorama virtual appliance. You are prompted with a certificate warning. Accept the certificate warning and continue to the web page.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Register the Panorama virtual appliance and activate the device management license and support licenses.

1. (VM Flex Licensing Only) Provisioning the Panorama Virtual Appliance Serial Number.

   When leveraging VM Flex licensing, this step is required to generate the Panorama virtual appliance serial number needed to register the Panorama virtual appliance with the Palo Alto Networks Customer Support Portal (CSP).
2. Register Panorama.

   You must register the Panorama virtual appliance using the serial number provided by Palo Alto Networks in the order fulfillment email.

   This step is not required when leveraging VM Flex licensing as the serial number is automatically registered with the CSP when generated.
3. Activate the firewall management license.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected.
4. Activate a Panorama Support License.

STEP 9 | Complete configuring the Panorama virtual appliance for your deployment needs.

- For Panorama in Log Collector Mode.
  1. Add a Virtual Disk to Panorama on Azure as needed.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Log Collector mode.
  2. Begin at Step 6 to switch to Log Collector mode.

     *Enter the Public IP address of the Dedicated Log Collector when you add the Log Collector as a managed collector to the Panorama management server. You cannot specify the **IP Address**, **Netmask**, or **Gateway**.*
- For Panorama in Panorama mode.
  1. Add a Virtual Disk to Panorama on Azure.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Panorama mode.
  2. Set up a Panorama Virtual Appliance in Panorama Mode.
  3. Configure a Managed Collector.
- For Panorama in Management Only mode.
  1. Set up a Panorama Virtual Appliance in Management Only Mode.
  2. Configure a Managed Collector to add a Dedicated Log Collector to the Panorama virtual appliance.

     Management Only mode does not support local log collection, and requires a Dedicated Log Collector to store managed device logs.

 ©2025 Palo Alto Networks, Inc.

Install Panorama on Google Cloud Platform

You can now deploy Panorama™ and a Dedicated Log Collector on Google Cloud Platform (GCP). Panorama deployed on GCP is Bring Your Own License (BYOL), supports all deployment modes (Panorama, Log Collector, and Management Only), and shares the same processes and functionality as the M-Series hardware appliances. For more information on Panorama modes, see Panorama Models.

To deploy the Panorama virtual appliance on GCP, you need to build a custom image. To begin this process, you must download the Panorama `tar.gz` from the Palo Alto Networks Customer Support portal and upload it to a GCP storage bucket. You can then create the custom image and use the image to deploy the Panorama virtual appliance on GCP.

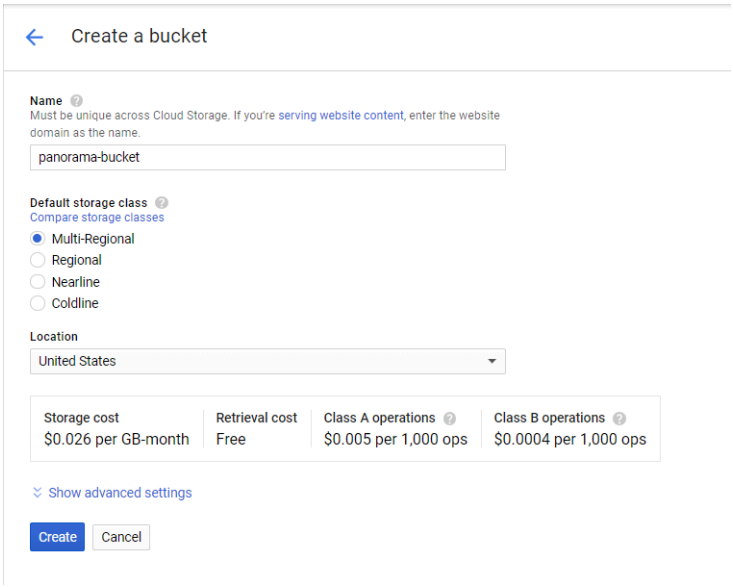
STEP 1 | Download the Panorama virtual appliance image.

1. Log in to the Palo Alto Networks Support Portal.
2. Select **Updates** > **Software Updates** and filter by **Panorama Base Images**.
3. Download the latest version of the Panorama on GCP `tar.gz` image.

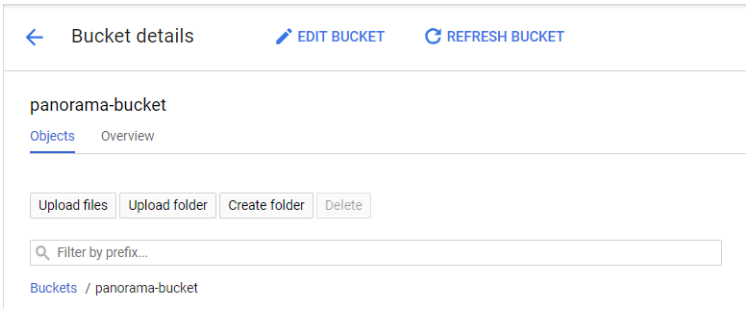
 ©2025 Palo Alto Networks, Inc.

STEP 2 | Upload the Panorama virtual appliance image to the Google Cloud Platform.

1. Log in to the Google Cloud Console.
2. From the **Products and Services** menu, select **Storage**.
3. Click **Create Bucket**, configure the new storage bucket and click **Create.**

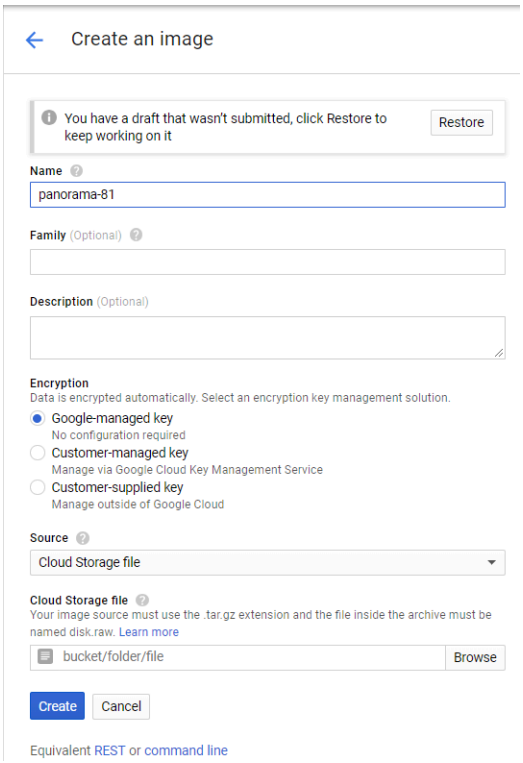


4. Select the storage bucket you created in the previous step, click **Upload files**, and select the Panorama virtual appliance image you downloaded.



5. From the **Products and Services** menu, select **Compute Engine** > **Images**.
6. Click **Create Image** and create the Panorama virtual appliance image:
   1. **Name** the Panorama virtual appliance image.
   2. In the **Source** field, select **Cloud Storage file** from the drop-down menu.
   3. Click **Browse** and navigate to the storage bucket where you uploaded the Panorama virtual appliance image, and **Select** the uploaded image.
   4. **Create** the Panorama virtual appliance image.

 ©2025 Palo Alto Networks, Inc.



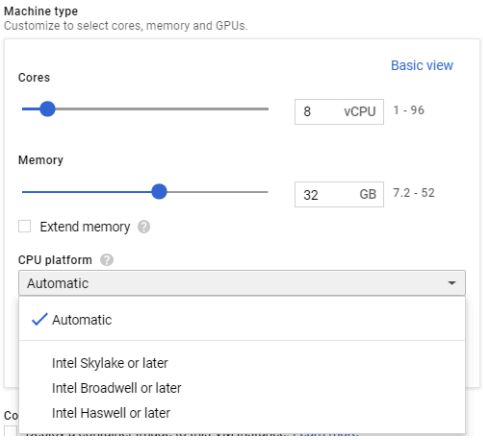
 ©2025 Palo Alto Networks, Inc.

STEP 3 | Configure the Panorama virtual appliance.

1. From the **Products and Services** menu and select the **Compute Engine**.
2. Click **Create Instance** to begin deploying the Panorama virtual appliance.
3. Add a descriptive **Name** to easily identify the Panorama virtual appliance.
4. Select the **Region** and **Zone** where you want to deploy the Panorama virtual appliance.
5. Allocate the **Machine Type** and **Customize** the CPU cores, memory and CPU platform. Review the Setup Prerequisites for the Panorama Virtual Appliance for minimum resource requirements.

   If you plan to use the Panorama virtual appliance as a Dedicated Log Collector, ensure that you configure the appliance with the required resources during initial deployment. The Panorama virtual appliance does not remain in Log Collector mode if you resize the virtual machine after you deploy it, and this results in a loss of log data.

   The GCP zone selection determines the CPU platforms available to you. For more information, refer to Regions and Zones for details.

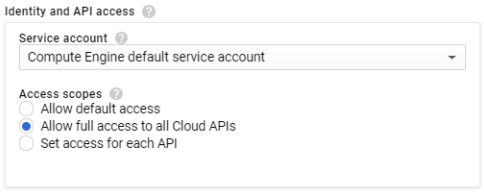


6. Configure the Panorama boot disk.
   1. For the **Boot Disk**, click **Change** > **Custom image** and select the Panorama image file you uploaded in Step 2
   2. Review the boot disk **Size** and verify the system disk is **81**GB.

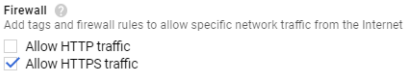
      *The Panorama virtual appliance must be initially installed with the **Default** system disk size. Installing the Panorama virtual appliance with a system disk larger than the **Default** system disk size is not supported and may result in limited utilization. You have the option to increase the system disk size after initial installation*

   3. Click **Select** to save your configuration.
7. Under **Identity and API access**, select **Allow full access to all Cloud APIs**.

 ©2025 Palo Alto Networks, Inc.



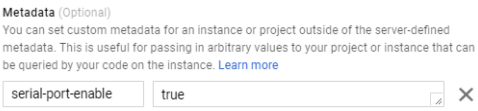
8. Under **Firewall**, select **Allow HTTPS traffic**.



STEP 4 | Expand **Management, security, disks, networking, sole tenancy** Management, security, disks, networking, sole tenancy .

STEP 5 | Enable access to the serial port so you can manage the Panorama virtual appliance.

1. Select **Management**.
2. Enter the following name-value pair as Metadata:
   `serial-port-enable true`



 ©2025 Palo Alto Networks, Inc.

STEP 6 | Reserve a static IP address for the management interface.

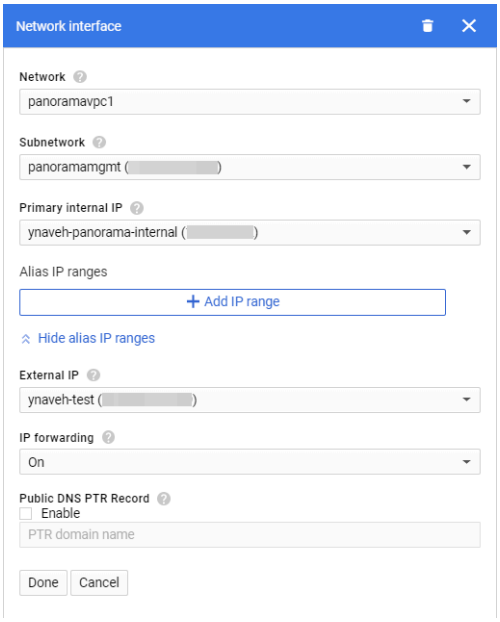
Reserve static internal and external IP addresses for the management interface in the event that if the Panorama virtual appliance is rebooted, your managed devices do not lose connection to the Panorama virtual appliance when the IP addresses are reassigned.

For more information on how to reserve IP addresses, refer to Reserving a Static Internal IP Address and Reserving a Static External IP Address.

1. Select **Networking**.
2. **Edit** the network interface.



3. Select the Panorama virtual appliance **Network**.
4. Select the Panorama virtual appliance **Subnetwork**. Instances in the same subnetwork will communicate with each other using their internal IP addresses.
5. Set the **Primary internal IP** address.
   - **Ephemeral (Automatic)**— Automatically assign a primary internal IP address.
   - **Ephemeral (Custom)**—Configure a custom IP range that GCP uses to assign a primary internal IP address.
   - **Reserve a static internal IP address**—Manually configure a static primary internal IP address.
6. Set the **External IP** address.
   - **Ephemeral**—Automatically assign an external IP address from a shared IP pool.
   - Select an existing reserved external IP address.
   - **Create IP address**—Reserve an external IP address.
7. Set **IP forwarding** to **On** to allow the Panorama virtual appliance to receive packets from non-matching destinations or source IP addresses.



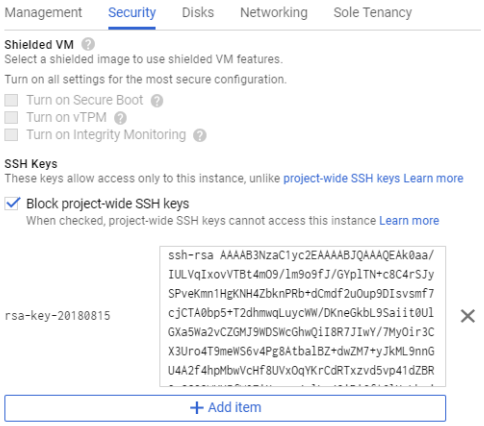
 ©2025 Palo Alto Networks, Inc.

STEP 7 | Configure the SSH key. You need an SSH key to access the Panorama virtual appliance CLI to configure the administrative user password after the initial deployment.

- **PuTTY Users**

  1. Select **Security**.
  2. Select the **Block project-wide SSH keys** box. Only instance keys are currently supported for logging in to the Panorama virtual appliance after initial deployment.
  3. Paste the SSH key in the comment box. For information on the correct SSH key format and how to generate SSH keys for GCP, refer to Managing SSH keys in Metadata.

     When generating the SSH key, save the private key in .ppk format. The private key is required to log in to the Panorama virtual appliance after the initial deployment before you can configure the administrative password.



- **Linux and macOS Users**

  1. Generate the SSH key from the CLI of your Linux device.

     ```
     ssh-keygen -C admin@panorama -f <panorama_key_name>
     ```

     Where **admin@panorama** is a comment GCP requires and **<panorama_key_name>** is the name of the key file being generated.
  2. Create an output file for the SSH key.

     ```
     cat <panorama_key_name>.pub
     ```

     After the output file for the SSH key is created, manually copy the SSH key contents.
  3. Paste the public key into the SSH Keys section of the GCP instance creation.

STEP 8 | (Optional) Add additional storage for log collection. Repeat this step as needed to add additional virtual logging disks.

If you intend to use the Panorama virtual appliance in Panorama mode or as a Dedicated Log Collector, add the virtual logging disks during initial deployment. By default, the Panorama virtual appliance is in Panorama mode for the initial deployment when you meet the Panorama mode resource requirements and have added at least one virtual logging disk. Otherwise, the

 ©2025 Palo Alto Networks, Inc.

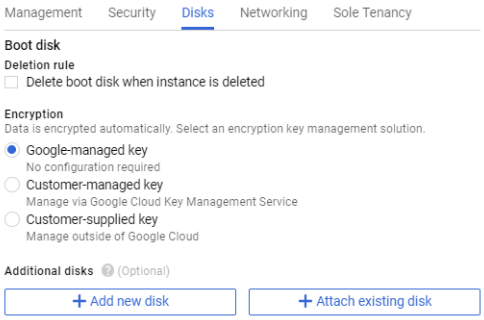
Panorama virtual appliance defaults to Management Only mode in which you can manage devices and Dedicated Log Collectors, and cannot collect logs locally.

The Panorama virtual appliance on GCP only supports 2TB logging disks, and in total supports up to 24TB of log storage. You are unable to add a logging disk smaller than 2TB, or a logging

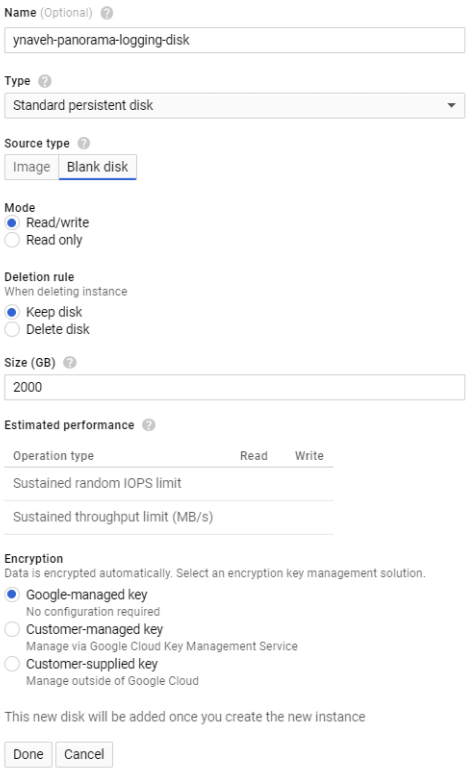
 ©2025 Palo Alto Networks, Inc.

disk with a size not divisible by the 2TB logging disk requirement. The Panorama virtual appliance partitions logging disks larger than 2TB into 2TB partitions.

1. Select **Disks** > **Add new disk**.



2. Enter the **Name**.
3. Expand the **Type** drop-down menu and select the desired type.
4. For the **Source type**, select **Blank disk**.
5. For the **Mode**, select **Read/write**.
6. Select the **Deletion rule** to configure whether to delete the virtual logging disk if the Panorama virtual appliance instance is deleted. To
7. Set the **Size (GB)** of the virtual logging disk.
8. Set your preferred **Encryption** solution for the data on the virtual logging disk.
9. Click **Done**.



 ©2025 Palo Alto Networks, Inc.

STEP 9 | **Create** the Panorama virtual appliance. The Panorama virtual appliances takes roughly 10 minutes to boot up after initial deployment.

STEP 10 | Configure a new administrative password for the Panorama virtual appliance.

You must configure a unique administrative password before you can access the web interface of the Panorama virtual appliance. To access the CLI, use the private key to launch the Panorama virtual appliance.

The new password must be a minimum of eight characters and include a minimum of one lowercase character, one uppercase character, and one number or special character.

- If you have an SSH service installed on your computer:
  1. Enter the following command to log into the Panorama virtual appliance:
     - Windows Devices

       ```
       ssh -i <private_key.ppk> admin@<public-ip_address>
       ```

     - Linux Devices

       ```
       ssh -i <prive_key.ppk> -oHostKeyAlgorithms=+ssh-rsa
        <username>@<public-ip_address>
       ```

       Including *-oHostKeyAlgorithms=+ssh-rsa* is required to specify the host key type. An error is displayed if this is not included in the SSH login command.
  2. Configure a new password using the following commands and follow the onscreen prompts:

     ```
     admin> configure
     ```

 ©2025 Palo Alto Networks, Inc.

```
admin# set mgt-config users admin password
```

3. If you have a BYOL that you need to, set the DNS server IP address so that the Panorama virtual appliance can access the Palo Alto Networks licensing server. Enter the following command to set the DNS server IP address:

```
admin# set deviceconfig system dns-setting servers
 primary <ip_address>
```

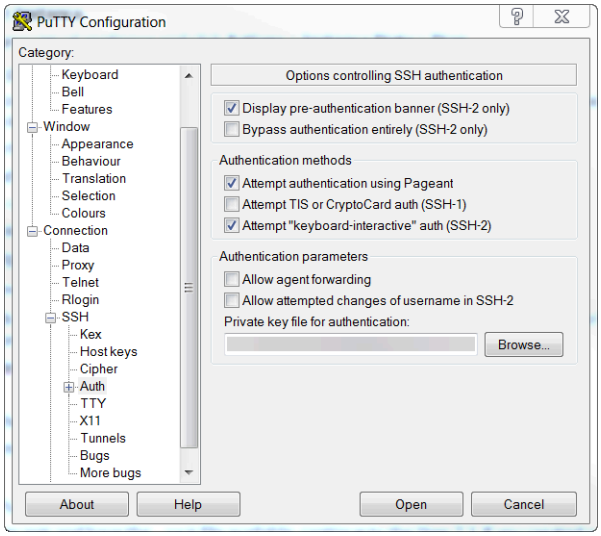
4. Commit your changes:

```
admin# commit
```

5. Terminate the SSH session.

- If you are using PuTTY to SSH into the Panorama virtual appliance:

1. If you are using an existing key pair and have the `.ppk` file available, continue to Step 11.3. If you created a new key pair or only have the `.pem` file of the existing key pair, open PuTTYgen and **Load** the `.pem` file.
2. **Save the private key** to a local accessible destination.
3. Open PuTTY and select **SSH** > **Auth** and **Browse** for the `.ppk` file saved in the previous step.



4. Select **Sessions** and enter the public IP address of the Panorama virtual appliance. Then **Open** and click **Yes** when the security prompt appears.
5. Login as admin when prompted.
6. Configure a new password using the following commands and follow the on screen prompts:

```
admin> configure
```

 ©2025 Palo Alto Networks, Inc.

```
admin# set mgt-config users admin password
```

7. Set the DNS server IP address so that the Panorama virtual appliance can access the Palo Alto Networks licensing server. Enter the following command to set the DNS server IP address:

```
admin# set deviceconfig system dns-setting servers
 primary <ip_address>
```

8. Commit your changes with the command:

```
admin# commit
```

9. Terminate the SSH session.

STEP 11 | Register the Panorama virtual appliance and activate the device management license and support licenses.

1. (VM Flex Licensing Only) Provisioning the Panorama Virtual Appliance Serial Number.

   When leveraging VM Flex licensing, this step is required to generate the Panorama virtual appliance serial number needed to register the Panorama virtual appliance with the Palo Alto Networks Customer Support Portal (CSP).

2. Register Panorama.

   You must register the Panorama virtual appliance using the serial number provided by Palo Alto Networks in the order fulfillment email.

   This step is not required when leveraging VM Flex licensing as the serial number is automatically registered with the CSP when generated.

3. Activate the firewall management license.

   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected.

4. Activate a Panorama Support License.

 ©2025 Palo Alto Networks, Inc.

STEP 12 | Complete configuring the Panorama virtual appliance for your deployment needs.

- For Panorama in Log Collector Mode.
  1. Add a Virtual Disk to Panorama on Google Cloud Platform as needed.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Log Collector mode.
  2. Begin at Step 6 to switch to Log Collector mode.

     *Enter the Public IP address of the Dedicated Log Collector when you add the Log Collector as a managed collector to the Panorama management server. You cannot specify the **IP Address**, **Netmask**, or **Gateway**.*
- For Panorama in Panorama mode.
  1. Add a Virtual Disk to Panorama on Google Cloud Platform.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Panorama mode.
  2. Set up a Panorama Virtual Appliance in Panorama Mode.
  3. Configure a Managed Collector.
- For Panorama in Management Only mode.
  1. Set up a Panorama Virtual Appliance in Management Only Mode.
  2. Configure a Managed Collector to add a Dedicated Log Collector to the Panorama virtual appliance.

     Management Only mode does not support local log collection, and requires a Dedicated Log Collector to store managed device logs.
- For SD-WAN deployments.
  1. Increase the System Disk for Panorama on Google Cloud Platform

     To leverage SD-WAN on Panorama deployed on GCP, you must increase the the system disk to 224GB.

     You cannot migrate back to a 81GB system disk after successfully increasing the system disk to 224GB.
  2. Set up a Panorama Virtual Appliance in Management Only Mode.
  3. Add a Virtual Disk to Panorama on Google Cloud Platform.

     To leverage SD-WAN, you must add a single 2TB logging disk to Panorama in Management Only mode.

Install Panorama on KVM

You can now deploy Panorama™ and a Dedicated Log Collector on KVM. Panorama deployed on KVM is Bring Your Own License (BYOL), supports all deployment modes (Panorama, Log Collector, and Management Only), and shares the same processes and functionality as the M-Series hardware appliances. For more information on Panorama modes, see Panorama Models.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Download the Panorama 10.2 base image QCOW2 file.

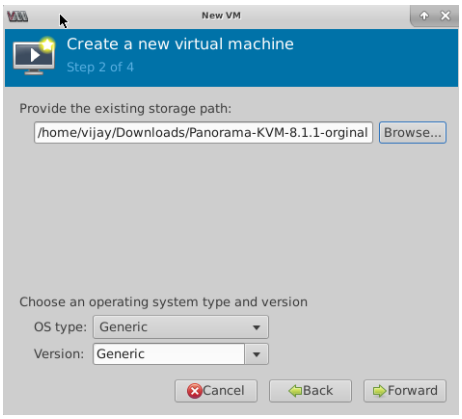
1. Log in to the Palo Alto Networks Support Portal.
2. Select **Updates** > **Software Updates** and filter by **Panorama Base Images** to download the QCOW2 file (`Panorama-KVM-10.2.0.qcow2`).

STEP 2 | Create a new virtual machine image and add the Panorama virtual appliance image for KVM to the Virtual Machine Manager.

1. On the Virtual Machine Manager, select **Create a new virtual machine**.
2. Select **Import Existing disk image** and click **Forward**.



3. **Browse** and select the Panorama virtual appliance image volume and **Choose volume**.
4. Click **Forward**.



 ©2025 Palo Alto Networks, Inc.

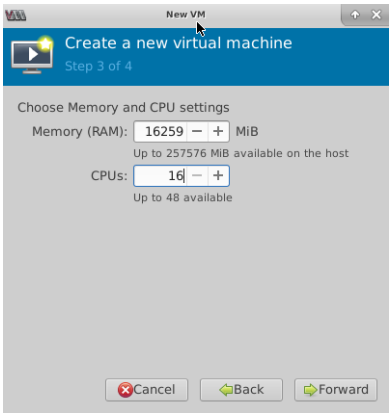
STEP 3 | Configure the memory and CPU settings.

Review the Setup Prerequisites for the Panorama Virtual Appliance for minimum resource requirements.

If you plan to use the Panorama virtual appliance as a Dedicated Log Collector, ensure that you configure the appliance with the required resources during initial deployment. The Panorama virtual appliance does not remain in Log Collector mode if you resize the virtual machine after you deploy it, and this results in a loss of log data.

1. Configure the **Memory** based on the requirements for the desired operational mode.

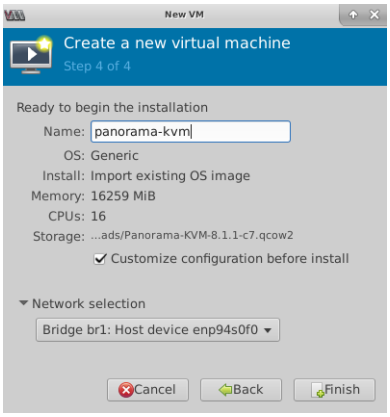
   The Virtual Machine Manager may use MiB (mebibyte) to allocate memory depending on the version you are running. If MiB is used, be sure to correctly convert your required memory allocation to avoid under provisioning the Panorama virtual appliance.
2. Configure the **CPU** based on the requirements for the desired operational mode.
3. Click **Forward**.



 ©2025 Palo Alto Networks, Inc.

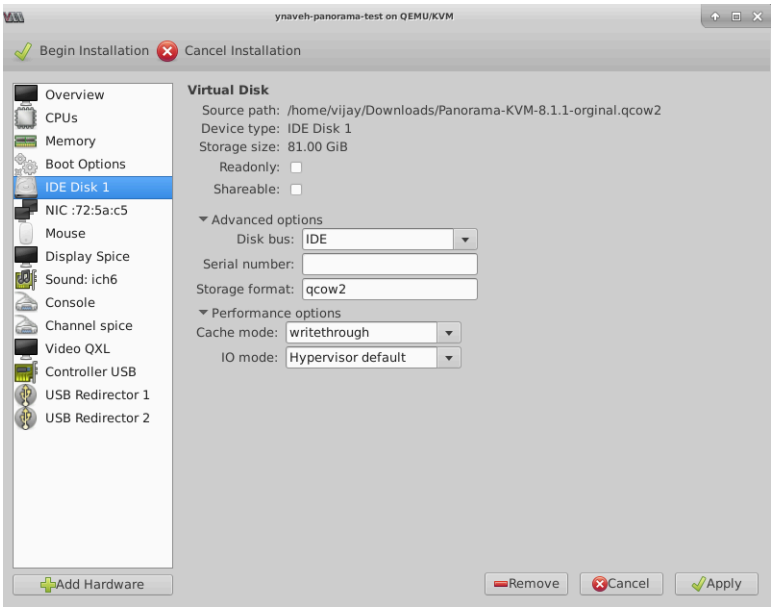
STEP 4 | Name the Panorama virtual appliance, enable configuration customization, and select the management interface bridge.

1. Enter a descriptive **Name** for the Panorama virtual appliance.
2. **Customize configuration before install**.
3. Make a **Network selection**—select the bridge for the management interface and accept the default settings.
4. Click **Finish**.



STEP 5 | Configure the virtual system disk settings.

1. Select **IDE Disk 1**, go to **Advanced options**, and select the following:
   - **Disk Bus**—**VirtIO** or **IDE**, depending on your configuration.
   - **Storage format**—**qcow2**
2. Go to **Performance options** and set **Cache mode** to **writethrough**. This setting improves installation time and execution speed on the Panorama virtual appliance.
3. Click **Apply**.



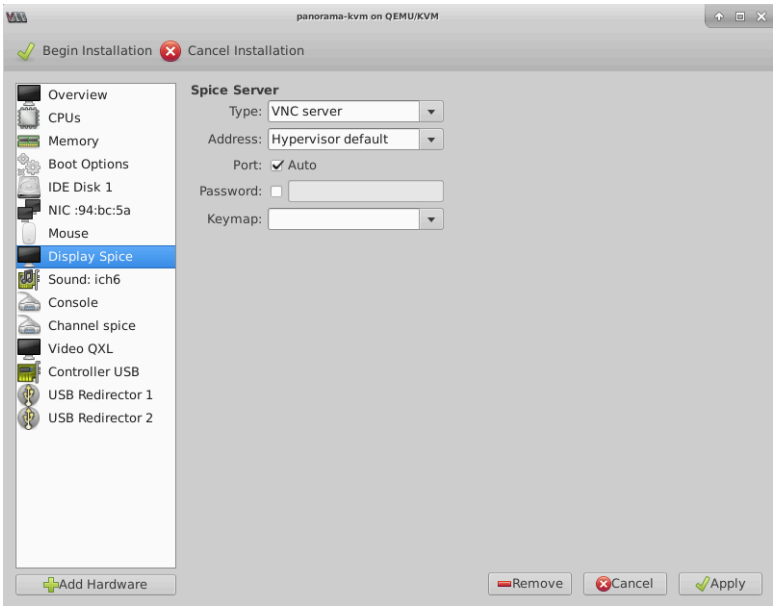
 ©2025 Palo Alto Networks, Inc.

STEP 6 | Configure the virtual machine console display to use the VNC server to interact with the virtual machine.

1. Select **Display Spice**.

   *Continue to the next step if **Display VNC** is listed in the Hardware list because the virtual machine is already configured to use the VNC server for the display.*

2. In the **Type** drop-down, select **VNC server**.
3. Click **Apply**.



STEP 7 | (Optional) Add additional storage for log collection. Repeat this step as needed to add additional virtual logging disks.

If you intend to use the Panorama virtual appliance in Panorama mode or as a Dedicated Log Collector, add the virtual logging disks during the initial deployment. By default, the Panorama virtual appliance is in Panorama mode for the initial deployment when you meet the Panorama mode resource requirements and have added at least one virtual logging disk. Otherwise, the Panorama virtual appliance defaults to Management Only mode. Change the Panorama virtual appliance to Management Only mode if you just want to manage devices and Dedicated Log Collectors, and to not collect logs locally.

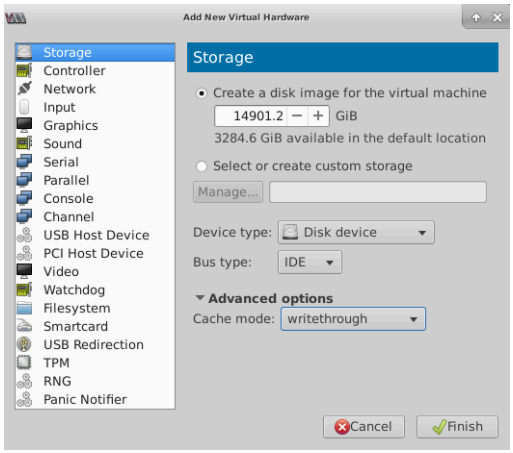
The Panorama virtual appliance on KVM only supports 2TB logging disks, and in total supports up to 24TB of log storage. You are unable to add a logging disk smaller than 2TB, or a logging

 ©2025 Palo Alto Networks, Inc.

disk with a size not divisible by the 2TB logging disk requirement. The Panorama virtual appliance partitions logging disks larger than 2TB into 2TB partitions.

1. **Add Hardware**.
2. Configure the new **Storage** disk:
   1. **Create a disk image for a virtual machine** and configure the virtual disk storage capacity to 14901.2 GiB (this is equivalent to 2TB).

      The Virtual Machine Manager may use GiB (gibibyte) to allocate memory depending on the version you are running. If GiB is used, be sure to correctly convert the required storage capacity to avoid under provisioning the virtual logging disk and sending the Panorama virtual appliance into maintenance mode.
   2. Set the **Device type** to **Disk** device.
   3. Set the **Bus type** to **VirtIO** or **IDE**, depending on your configuration.
   4. Go to **Advanced options** and set **Cache mode** to **writethrough**.
3. Click **Finish**.



STEP 8 | **Begin Installation** (Begin Installation). The Panorama virtual appliances takes approximately 10 minutes to boot.

STEP 9 | Open a connection to the console of the Panorama virtual appliance.

You are prompted to log in to the firewall using the default username and password: **`admin`**/**`admin`**.

STEP 10 | Configure a new administrative password for the Panorama virtual appliance.

You must configure a unique administrative password before you can access the web interface or CLI of the Panorama virtual appliance. The new password must be a minimum of eight characters and include a minimum of one lowercase character, one uppercase character, and one number or special character.

When you first log in to the Panorama CLI, you are prompted to enter the **Old Password** and the **New Password** for the `admin` user before you can continue.

 ©2025 Palo Alto Networks, Inc.

STEP 11 | Configure the network access settings for the management interface.

1. Enter configuration mode using the following command:

   ```
   admin> configure
   ```

2. Use the following commands to configure and enable access to the management interface:

   ```
   admin# set deviceconfig system type static
   admin# set deviceconfig system ip-address <Panorama-IP>
    netmask <netmask> default-gateway <gateway-IP> dns-setting
    servers primary <DNS-IP>
   ```

   where *<Panorama-IP>* is the IP address you want to assign to the management interface, *<netmask>* is the subnet mask, *<gateway-IP>* is the IP address of the network gateway, and *<DNS-IP>* is the IP address of the DNS server.

   ```
   admin# commit
   ```

STEP 12 | Register the Panorama virtual appliance and activate the device management license and support licenses.

1. (VM Flex Licensing Only) Provisioning the Panorama Virtual Appliance Serial Number.

   When leveraging VM Flex licensing, this step is required to generate the Panorama virtual appliance serial number needed to register the Panorama virtual appliance with the Palo Alto Networks Customer Support Portal (CSP).

2. Register Panorama.

   You must register the Panorama virtual appliance using the serial number provided by Palo Alto Networks in the order fulfillment email.

   This step is not required when leveraging VM Flex licensing as the serial number is automatically registered with the CSP when generated.

3. Activate the firewall management license.

   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected.

4. Activate a Panorama Support License.

 ©2025 Palo Alto Networks, Inc.

STEP 13 | Complete configuring the Panorama virtual appliance for your deployment needs.

- For Panorama in Log Collector Mode.
  1. Add a Virtual Disk to Panorama on KVM as needed.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Log Collector mode.
  2. Begin at Step 6 to switch to Log Collector mode.

     *Enter the Public IP address of the Dedicated Log Collector when you add the Log Collector as a managed collector to the Panorama management server. You cannot specify the **IP Address**, **Netmask**, or **Gateway**.*
- For Panorama in Panorama mode.
  1. Add a Virtual Disk to Panorama on KVM.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Panorama mode.
  2. Set up a Panorama Virtual Appliance in Panorama Mode.
  3. Configure a Managed Collector.
- For Panorama in Management Only mode.
  1. Set up a Panorama Virtual Appliance in Management Only Mode.
  2. Configure a Managed Collector to add a Dedicated Log Collector to the Panorama virtual appliance.

     Management Only mode does not support local log collection, and requires a Dedicated Log Collector to store managed device logs.

Install Panorama on Hyper-V

You can now deploy Panorama™ and a Dedicated Log Collector on Hyper-V. Panorama deployed on Hyper-V is Bring Your Own License (BYOL), supports all deployment modes (Panorama, Log Collector, and Management Only), and shares the same processes and functionality as the M-Series hardware appliances. For more information on Panorama modes, see Panorama Models. Panorama virtual appliance and virtual Dedicated Log Collector on Hyper-V is available only on PAN-OS 8.1.3 and later releases.

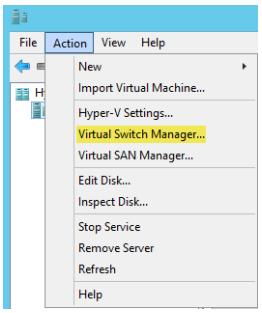
STEP 1 | Download the Panorama 10.2 base image VHDX file.

1. Log in to the Palo Alto Networks Support Portal.
2. Select **Updates** > **Software Updates** and filter by **Panorama Base Images** to download the VHDX file (`Panorama-HPV-10.2.0.vhdx`).

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Set up any vSwitch(es) that you will need. For more information, review the Virtual Switch Types for more information.

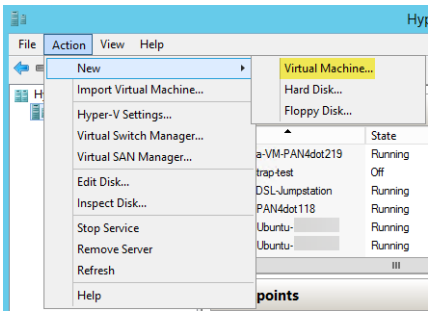
1. From Hyper-V Manager, select the host and select **Action** > **Virtual Switch Manager** to open the Virtual Switch Manager window.



2. Under **Create virtual switch**, select the type of vSwitch to create and click **Create Virtual Switch**.

STEP 3 | Install the Panorama virtual appliance.

1. On the Hyper-V Manager, select the host and select **Action** > **New** > **Virtual Machine.** Configure the following settings in the New Virtual Machine Wizard:



1. Choose a **Name** and **Location** for the Panorama virtual appliance. The Panorama virtual appliance stores the VHDX file at the specified location.
2. Choose **Generation 1**. This is the default option and the only version supported.
3. For **Startup Memory**, assign the memory based on the intended system mode. See the Setup Prerequisites for the Panorama Virtual Appliance for the memory requirements for each mode.

   Do not enable dynamic memory; the Panorama virtual appliance requires static memory allocation.
4. Configure **Networking**. Select an external vSwitch to connect the management interface on the firewall.
5. To connect the **Virtual Hard Disk**, select **Use an existing virtual hard disk** and browse to the VHDX file you downloaded earlier.
6. Review the summary and click **Finish**.

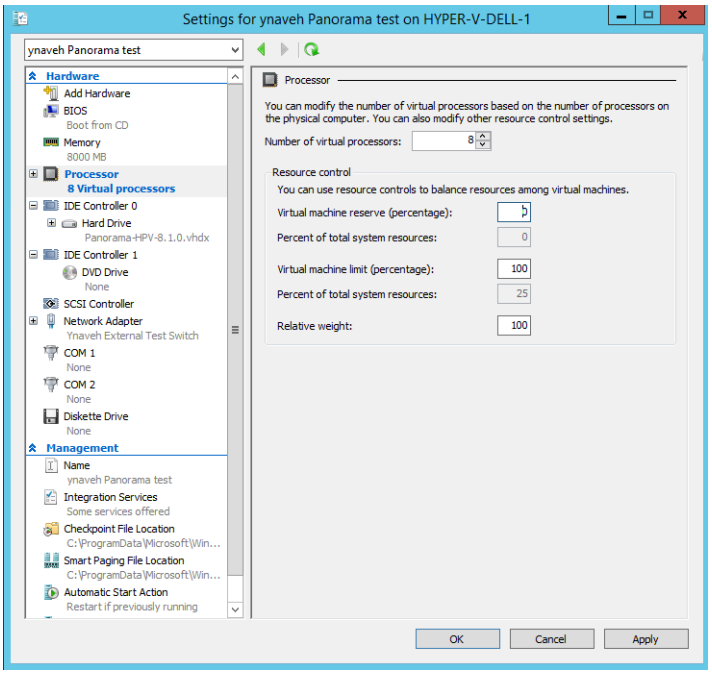
 ©2025 Palo Alto Networks, Inc.

STEP 4 | Allocate the Panorama virtual appliance CPU cores.

Review the Setup Prerequisites for the Panorama Virtual Appliance for minimum resource requirements.

If you plan to use the Panorama virtual appliance as a Dedicated Log Collector, ensure that you configure the appliance with the required resources during initial deployment. The Panorama virtual appliance does not remain in Log Collector mode if you resize the virtual machine after you deploy it, and this results in a loss of log data.

1. In the **Hardware** list, select **Processor**.
2. Edit the currently allocated **Number of virtual processors**.



STEP 5 | Connect at least one network adapter for the dataplane interface on the firewall. Repeat this to create additional network interfaces on the Panorama virtual appliance.

1. Select **Settings** > **Hardware** > **Add Hardware** and select the **Hardware type** for your network adapter.

   Legacy Network Adapter and SR-IOV are not supported. If selected, the VM-Series firewall will boot into maintenance mode.

2. Click **OK**.

STEP 6 | (Optional) Add additional storage for log collection. Repeat this step as needed to add additional virtual logging disks. If you want to deploy the Panorama virtual appliance in Management Only mode, continue to Step 6.

If you intend to use the Panorama virtual appliance in Panorama mode or as a Dedicated Log Collector, add the virtual logging disks during the initial deployment. By default, the Panorama virtual appliance is in Panorama mode for the initial deployment when you meet the Panorama mode resource requirements and have added at least one virtual logging disk. Otherwise, the Panorama virtual appliance defaults to Management Only mode. Change the Panorama virtual

 ©2025 Palo Alto Networks, Inc.

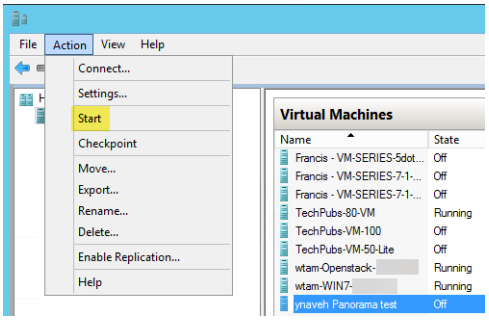
appliance to Management Only mode if you just want to manage devices and Dedicated Log Collectors, and to not collect logs locally.

The Panorama virtual appliance on Hyper-V only supports 2TB logging disks, and in total supports up to 24TB of log storage. You are unable to add a logging disk smaller than 2TB, or a logging disk with a size not divisible by the 2TB logging disk requirement. The Panorama virtual appliance partitions logging disks larger than 2TB into 2TB partitions.

1. On the Hyper-V Manager, select the host and select **Action** > **New** > **Hard Disk.**
2. If you see the Before You Begin prompt, click **Next** to begin adding the virtual logging disk.
3. For the Disk Format, select **VHDX**. Click **Next**. to continue
4. For the Disk Type, select **Fixed Size** or **Dynamically Expanding** based on your needs. Click **Next** to continue.
5. Specify the **Name** and **Location** for the virtual logging disk file. Click **Next** to continue.
6. To configure the disk, select **Create a new virtual hard disk** and enter the disk size. Click **Next** to continue.
7. Review the Summary and **Finish** adding the virtual hard logging disk.

STEP 7 | Power on the Panorama virtual appliance.

1. Select the Panorama virtual appliance instance from the list of **Virtual Machines**.
2. Select **Action** > **Start** to power on the Panorama virtual appliance.



STEP 8 | Connect to the Panorama virtual appliance console from the Hyper-V Manager.

1. In the **Virtual Machines** list, select the Panorama virtual appliance.
2. Select **Actions** > **Connect** and enter the username and password to log in (default is admin for both).

STEP 9 | Configure a new administrative password for the Panorama virtual appliance.

You must configure a unique administrative password before you can access the web interface or CLI of the Panorama virtual appliance. The new password must be a minimum of eight characters and include a minimum of one lowercase character, one uppercase character, and one number or special character.

When you first log in to the Panorama CLI, you are prompted to enter the **Old Password** and the **New Password** for the `admin` user before you can continue.

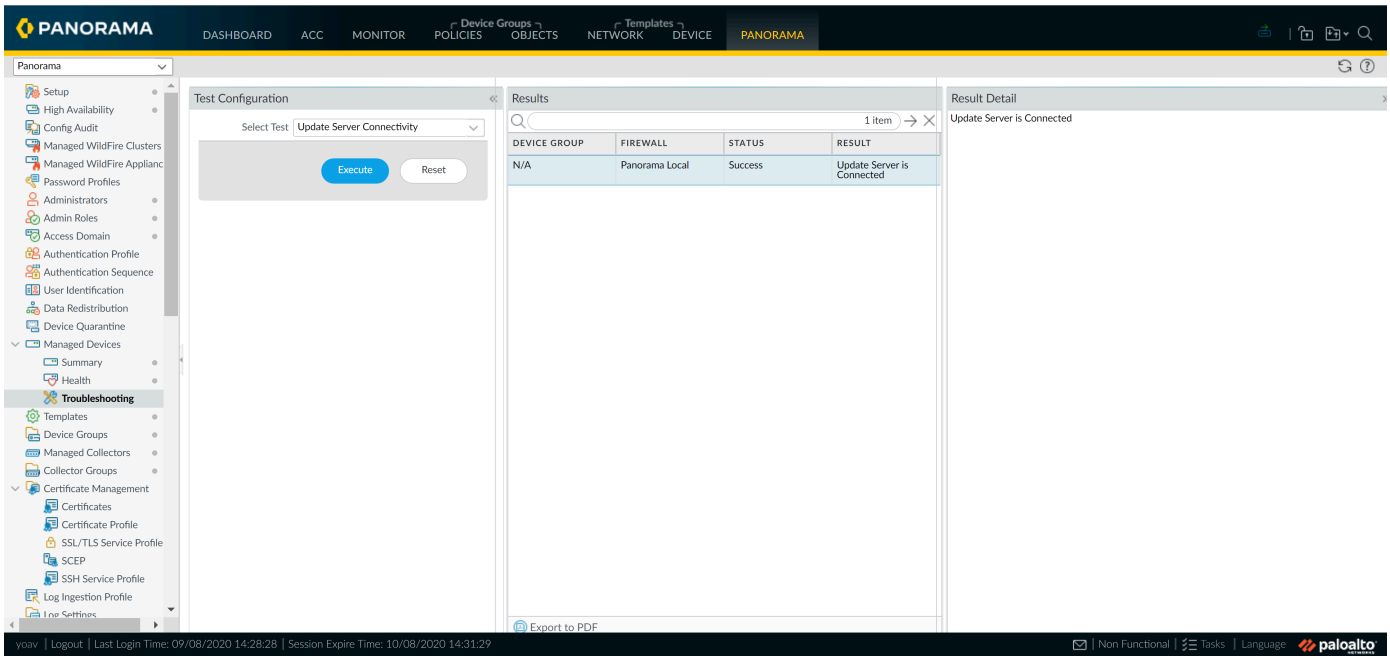
 ©2025 Palo Alto Networks, Inc.

STEP 10 | Configure the IP address of the management interface.

1. Enter the following commands, where *<Panorama-IP>* is the IP address you want to assign to the Panorama management interface, *<netmask>* is the subnet mask, *<gateway-IP>* is the IP address of the network gateway, and *<DNS-IP>* is the IP address of the DNS server:

```
admin> configure
admin# set deviceconfig system ip-address <Panorama-IP>
 netmask <netmask> default-gateway <gateway-IP> dns-setting
 servers primary <DNS-IP>
admin# commit
admin# exit
```

2. Troubleshoot Connectivity to Network Resources to verify network access to external services required for firewall management, such as the default gateway, DNS server, and the Palo Alto Networks Update Server, as shown in the following example:



 ©2025 Palo Alto Networks, Inc.

STEP 11 | Register the Panorama virtual appliance and activate the device management license and support licenses.

1. (VM Flex Licensing Only) Provisioning the Panorama Virtual Appliance Serial Number.

   When leveraging VM Flex licensing, this step is required to generate the Panorama virtual appliance serial number needed to register the Panorama virtual appliance with the Palo Alto Networks Customer Support Portal (CSP).
2. Register Panorama.

   You must register the Panorama virtual appliance using the serial number provided by Palo Alto Networks in the order fulfillment email.

   This step is not required when leveraging VM Flex licensing as the serial number is automatically registered with the CSP when generated.
3. Activate the firewall management license.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected.
4. Activate a Panorama Support License.

STEP 12 | Complete configuring the Panorama virtual appliance for your deployment needs.

- For Panorama in Log Collector Mode.
  1. Add a Virtual Disk to Panorama on Hyper-V as needed.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Log Collector mode.
  2. Begin at Step 6 to switch to Log Collector mode.

     *Enter the Public IP address of the Dedicated Log Collector when you add the Log Collector as a managed collector to the Panorama management server. You cannot specify the **IP Address**, **Netmask**, or **Gateway**.*
- For Panorama in Panorama mode.
  1. Add a Virtual Disk to Panorama on Hyper-V.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Panorama mode.
  2. Set up a Panorama Virtual Appliance in Panorama Mode.
  3. Configure a Managed Collector.
- For Panorama in Management Only mode.
  1. Set up a Panorama Virtual Appliance in Management Only Mode.
  2. Configure a Managed Collector to add a Dedicated Log Collector to the Panorama virtual appliance.

     Management Only mode does not support local log collection, and requires a Dedicated Log Collector to store managed device logs.

 ©2025 Palo Alto Networks, Inc.

Set Up Panorama on Oracle Cloud Infrastructure (OCI)

Set up a Panorama™ virtual appliance on Oracle Cloud Infrastructure (OCI) to centrally managed the configuration of physical and VM-Series firewalls.

- Upload the Panorama Virtual Appliance Image to OCI
- Install Panorama on Oracle Cloud Infrastructure (OCI)
- Generate a SSH Key for Panorama on OCI

Upload the Panorama Virtual Appliance Image to OCI

Complete the following procedure to upload a Panorama qcow2 file for KVM and create a custom image that you need to launch the Panorama virtual appliance. Uploading and creating the image is required only once. You can use the same image for all subsequent deployments of the Panorama virtual appliance.

STEP 1 | Download the Panorama qcow2 file for KVM from the Palo Alto Networks Customer Support Portal (CSP).

1. Log in to the Palo Alto Networks CSP.
2. Select **Updates** > **Software Updates** and select **Panorama Base Images** from the software updates filter drop-down.
3. Download the latest version of the `Panorama-KVM` qcow2 image.

STEP 2 | Log in to the Oracle Cloud Infrastructure console.

STEP 3 | Create a storage bucket for the qcow2 file.

1. Select **Object Storage** > **Object Storage** and **Create Bucket**.
2. Enter a descriptive **Bucket Name**.
3. For the Storage Tier, select **Standard**.
4. **Create Bucket**.

STEP 4 | Upload the qcow2 image to the OCI storage bucket.

1. Click the storage bucket you created in the previous step to view the bucket details.
2. Click **Upload** and select the qcow2 image you downloaded from the Palo Alto Networks CSP.
3. **Upload** the image.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Create a pre-authenticated request for the qcow2 file.

This is required to create the object URL used in the creation of the custom image for the Panorama virtual appliance.

1. Select **Object Storage** > **Object Storage** and click the storage bucket you created in the previous step.
2. Select **Pre-Authenticated Requests** > **Create Pre-Authenticated Request**.
3. Enter a descriptive **Name** for your Pre-Authenticated Request.
4. Select **Object** and enter the qcow2 image name for the **Object Name**.
5. **Create Pre-Authenticated Request**.
6. For the Access Type, select **Permit object reads and writes**.
7. Enter an **Expiration** date and time.
8. **Create Pre-Authenticated Request**.
9. In the Pre-Authenticated Request Details, copy the Pre-Authenticated Request URL.

   The Pre-Authenticated Request URL is required to create the custom image and must be copied when displayed to you.

   The Pre-Authenticated Request URL is only displayed after the request is created and is not shown again.
10. **Close** the Pre-Authenticated Request Details after you copy the URL.

STEP 6 | Import the qcow2 file and create a custom Panorama virtual appliance image.

1. Select **Compute** > **Custom Images** and **Import Image**.
2. Enter a descriptive **Name** for your image.
3. Select **Import from an Object Storage URL** and paste the object storage URL.
4. For the Image type, select **QCOW2**.
5. For the Launch Mode, select **Paravirtualized Mode**.
6. **Import Image**.

Install Panorama on Oracle Cloud Infrastructure (OCI)

Create a Panorama™ virtual appliance instance on Oracle Cloud Infrastructure (OCI). An OCI instance supports a single NIC by default. You must manually upload a Panorama virtual appliance qcow2 image downloaded from the Palo Alto Networks Customer Supported Portal (CSP) to OCI to successfully install the Panorama virtual appliance on OCI.

A Panorama virtual appliance deployed on OCI is Bring Your Own License (BYOL), supports all deployment modes (Panorama, Log Collector, and Management Only), and shares the same processes and functionality as the M-Series hardware appliances. For more information on Panorama modes, see Panorama Models.

A machine running a Linux operating system is required successfully install the Panorama on OCI. To successfully install Panorama on OCI, you must generate a `.pub` key using OpenSSH. Additionally, you can only use a Linux machine to log into the Panorama CLI for the initial network configuration.

 ©2025 Palo Alto Networks, Inc.

Review the Setup Prerequisites for the Panorama Virtual Appliance to determine the virtual resources required for your needs. The virtual resources requirement for the Panorama virtual appliance is based on the total number of firewalls managed by the Panorama virtual appliance and the required Logs Per Second (LPS) for forwarding logs from your managed firewalls to your Log Collector.

Under-provisioning the Panorama virtual appliance will impact management performance. This includes the Panorama virtual appliance becoming slow or unresponsive depending on how under-provisioning the Panorama virtual appliance is.

STEP 1 | Log in to the Oracle Cloud Infrastructure console.

STEP 2 | Set up the Virtual Cloud Network (VCN) for your network needs.

Whether you launch the Panorama virtual appliance in an existing VCN or you create a new VCN, the Panorama virtual appliance must be able to receive traffic from other instances in the VCN and perform inbound and outbound communication between the VCN and the internet as needed.

Refer to the OCI VCN documentation for more information.

1. Configure a VCN or use an existing VCN.
2. Verify that the network and security components are appropriately defined.
   - Create an internet gateway to enable internet access to the subnet of your Panorama virtual appliance. Internet access is required to install software and content updates, activate licenses, and leverage Palo Alto Networks cloud services. Otherwise, you must manually install updates and activate licenses.

     If the Panorama virtual appliance instance is part of a private subnet, you can configure a NAT gateway to enable only outbound internet access for the subnet.
   - Create subnets. Subnets are segments of the IP address range assigned to the VCN in which you can launch OCI instances. It is recommended that the Panorama virtual

 ©2025 Palo Alto Networks, Inc.

appliance belong to the management subnet so that you can configure it to access the internet if needed.

- Add routes to the route table for a private subnet to ensure traffic can be routed across subnets in the VCN and from the internet if applicable.

  Ensure you create routes between subnets to allow communication between:

  - Panorama, managed firewalls, and Log Collectors.
  - (Optional) Panorama and the internet.
- Ensure that the following ingress security rules are allowed for the VCN to manage VCN traffic. The ingress traffic source for each rule is unique to your deployment topology.

  See Ports Used for Panorama for more information.

  - Allow SSH (port **22**) traffic to enable access to the Panorama CLI.
  - Allow HTTPS (port **443** and **28270**) traffic to enable access to the Panorama web interface.
  - Allow traffic on port **3978** to enable communication between Panorama, manage firewalls, and managed Log Collectors. This port is also used by Log Collectors to forward logs to Panorama.
  - Allow traffic on port **28443** to enable managed firewalls to get software and content updates from Panorama.

STEP 3 | Select **Compute** > **Instances** and **Create Instance**.

STEP 4 | Enter a descriptive **Name** for the Panorama virtual appliance image.

STEP 5 | Select the **Availability domain**.

STEP 6 | Select the Palo Alto Networks Panorama image.

*See Upload the Panorama Virtual Appliance Image to OCI to upload and maintain your own Panorama virtual appliance **Custom Image** on OCI.*

1. Under Image and shape, select **Change Image**.
2. For the Image Source, select **Partner Image**.

   If you maintain your own Panorama virtual appliance image, select **Custom Image** instead and select the Panorama virtual appliance image you uploaded to OCI.
3. Search for **`Palo Alto Networks Panorama`** and select (check) the image.

   *Skip this step if you selected **Custom Image** in the previous step.*

   PAN-OS 10.2.0 is the default PAN-OS version.
4. **Select Image**.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Configure the instance resources.

Refer to the Setup Prerequisites for the Panorama Virtual Appliance for more information for the minimum resources required based on your Panorama usage needs.

1. Under Image and shape, select **Change Shape**.
2. Select the shape with number of CPUs, amount of RAM, and number of interfaces you require.
3. **Select Shape**.

STEP 8 | Configure the instance Networking settings.

1. For the Network, **Select existing virtual cloud network** and select the VCN.
2. For the Subnet, **Select existing subnet** and select the subnet.

   It is recommended to deploy the Panorama virtual appliance instance in a management subnet to safely allow internet access if needed.
3. (Optional) For the Public IP Address, select **Assign a public IPv4 address** if you want to make the Panorama virtual appliance accessible from outside the VCN.

STEP 9 | Configure the Panorama virtual appliance instance boot volume.

1. For the Boot volume, **specify a custom boot volume size**.
2. For the Boot volume size, enter **81**.

STEP 10 | **Create** the Panorama virtual appliance image.

STEP 11 | Log in to the Panorama virtual appliance CLI from the OCI console.

1. Generate a SSH Key for Panorama on OCI.
2. In the OCI console, select **Instances** and select the Panorama virtual appliance instance.
3. Select **Console Connection** and **Create Console Connection**.
4. Select **Upload public key files (.pub)** and upload the public SSH key you generated to **Create Console Connection**.
5. In the Instance Details screen, expand the Console Connection options and **Copy Serial Connection for Linux/Mac**.
6. On your Linux machine, open a terminal and paste the serial connection.

STEP 12 | Configure a new administrative password for the Panorama virtual appliance.

You must configure a unique administrative password before you can access the web interface or CLI of the Panorama virtual appliance. The new password must be a minimum of eight characters and include a minimum of one lowercase character, one uppercase character, and one number or special character.

When you first log in to the Panorama CLI, you are prompted to enter the **Old Password** and the **New Password** for the `admin` user before you can continue.

 ©2025 Palo Alto Networks, Inc.

STEP 13 | Configure the system IP address settings for the Panorama virtual appliance.

1. Configure the initial network settings for the Panorama virtual appliance.

```
admin> configure
```

```
admin#  set deviceconfig system type static
```

```
admin#  set deviceconfig system ip-address <instance-private-
IP address> netmask <netmask> default-gateway <default-
gateway-IP>
```

```
admin#  set deviceconfig system dns-setting servers primary
 <primary-dns-IP>
```

```
admin#  set deviceconfig system dns-setting servers secondary
 <secondary-dns-IP>
```

```
admin# commit
```

2. Verify you can log in to the Panorama web interface.

   If you cannot log in to the Panorama web interface, review your route table and VCN security rules to ensure the correct routes and security rules are created.

STEP 14 | Register the Panorama virtual appliance and activate the device management license and support licenses.

1. (VM Flex Licensing Only) Provisioning the Panorama Virtual Appliance Serial Number.

   When leveraging VM Flex licensing, this step is required to generate the Panorama virtual appliance serial number needed to register the Panorama virtual appliance with the Palo Alto Networks Customer Support Portal (CSP).

2. Register Panorama.

   You must register the Panorama virtual appliance using the serial number provided by Palo Alto Networks in the order fulfillment email.

   This step is not required when leveraging VM Flex licensing as the serial number is automatically registered with the CSP when generated.

3. Activate the firewall management license.

   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected.
   - Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected.

4. Activate a Panorama Support License.

 ©2025 Palo Alto Networks, Inc.

STEP 15 | Complete configuring the Panorama virtual appliance for your deployment needs.

- For Panorama in Log Collector Mode.
  1. Add a Virtual Disk to Panorama on Oracle Cloud Infrastructure (OCI) as needed.

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Log Collector mode.
  2. Begin at Step 6 to switch to Log Collector mode.

     *Enter the Public IP address of the Dedicated Log Collector when you add the Log Collector as a managed collector to the Panorama management server. You cannot specify the **IP Address**, **Netmask**, or **Gateway**.*
- For Panorama in Panorama mode.
  1. Add a Virtual Disk to Panorama on Oracle Cloud Infrastructure (OCI).

     Adding at least one virtual logging disk is required before you can change the Panorama virtual appliance to Panorama mode.
  2. Set up a Panorama Virtual Appliance in Panorama Mode.
  3. Configure a Managed Collector.
- For Panorama in Management Only mode.
  1. Set up a Panorama Virtual Appliance in Management Only Mode.
  2. Configure a Managed Collector to add a Dedicated Log Collector to the Panorama virtual appliance.

     Management Only mode does not support local log collection, and requires a Dedicated Log Collector to store managed device logs.

Generate a SSH Key for Panorama on OCI

To connect to the Panorama™ virtual appliance installed on Oracle Cloud Infrastructure (OCI), you must generate a public and private SSH key on a Linux machine. You use the generated SSH key to log in to the Panorama CLI to set up a new administrative password and configure the Panorama network settings.

A Linux machine is required to generate the SSH key and access the Panorama CLI for the initial configuration. Generating a SSH from OCI or third-party applications such as PuTTygen is not supported.

STEP 1 | Open the terminal on your Linux machine.

STEP 2 | Navigate to the hidden `.ssh` directory.

```
admin:~$ cd ~/.ssh
```

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Generate an SSH key in the `.ssh` directory.

```
admin:~/.ssh$ ssh-keygen
```

When prompted, save the key in the default `.ssh` directory. A password for the key is optional.

The default name for the private key is `id_rsa` and the default name for the public key is `id_rsa.pub`.

STEP 4 | Copy the public key from the `.ssh` directory to your home directory.

This step is required to upload the public key to OCI.

```
admin: ~/.ssh$ cp id_rsa.pub ~
```

Perform Initial Configuration of the Panorama Virtual Appliance

Based on your Panorama model, use the Alibaba Cloud Console, AWS, Azure, GCP, or OCI web interface, KVM Virtual Machine Manager, Hyper-V Manager, VMware vSphere Client, or vCloud Air web console to set up network access to the Panorama virtual appliance. By default, the Panorama virtual appliance is deployed in Panorama mode. For unified reporting, consider using Greenwich Mean Time (GMT) or Coordinated Universal Time (UTC) as the uniform time zone across Panorama and all the managed firewalls and Log Collectors.

STEP 1 | Gather the required information from your network administrator.

Collect the following information for the management (MGT) interface:

- IP address for the management (MGT) interface

  The default management interface IP address is 192.168.1.1. if you do not configure the management interface as described when you install the Panorama virtual appliance.

- Netmask
- Default gateway
- DNS server IP address

  To complete the configuration of the MGT interface, you must specify the IP address, netmask (for IPv4) or prefix length (for IPv6), and default gateway. If you omit settings (such as the default gateway), you can access Panorama only through the console port for future configuration changes. As a best practice, always commit a complete MGT interface configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Access the console of the Panorama virtual appliance.

Panorama uses the MGT interface for management traffic, high availability synchronization, log collection, and communication within Collector Groups.

Starting with PAN-OS 9.0.4, the default `admin` credentials are no longer supported. When you first install the Panorama virtual appliance, you are required to log in to the Panorama CLI to configure a unique `admin` password.

*If this is the first time you are logging in to the Panorama CLI, you are prompted to enter the **Old Password** and the **New Password** for the `admin` user before you can continue with the initial configuration of the Panorama virtual appliance.*

1. Access the console.

   On an ESXi server:

   1. Launch the VMware vSphere Client.
   2. Select the **Console** tab for the Panorama virtual appliance and press enter to access the login screen.

   On vCloud Air:

   1. Access the vCloud Air web console and select your **Virtual Private Cloud OnDemand** region.
   2. Select the **Virtual Machines** tab, right-click the Panorama virtual machine, and select **Open In Console**.
2. Enter your username and password to log in (default is admin for both).

   On Alibaba Cloud, AWS, Azure, GCP, KVM, Hyper-V, and OCI:

   - Log in to the Panorama CLI.

STEP 3 | Configure the network access settings for the MGT interface.

Panorama uses the MGT interface for management traffic, high availability synchronization, log collection, and communication within Collector Groups.

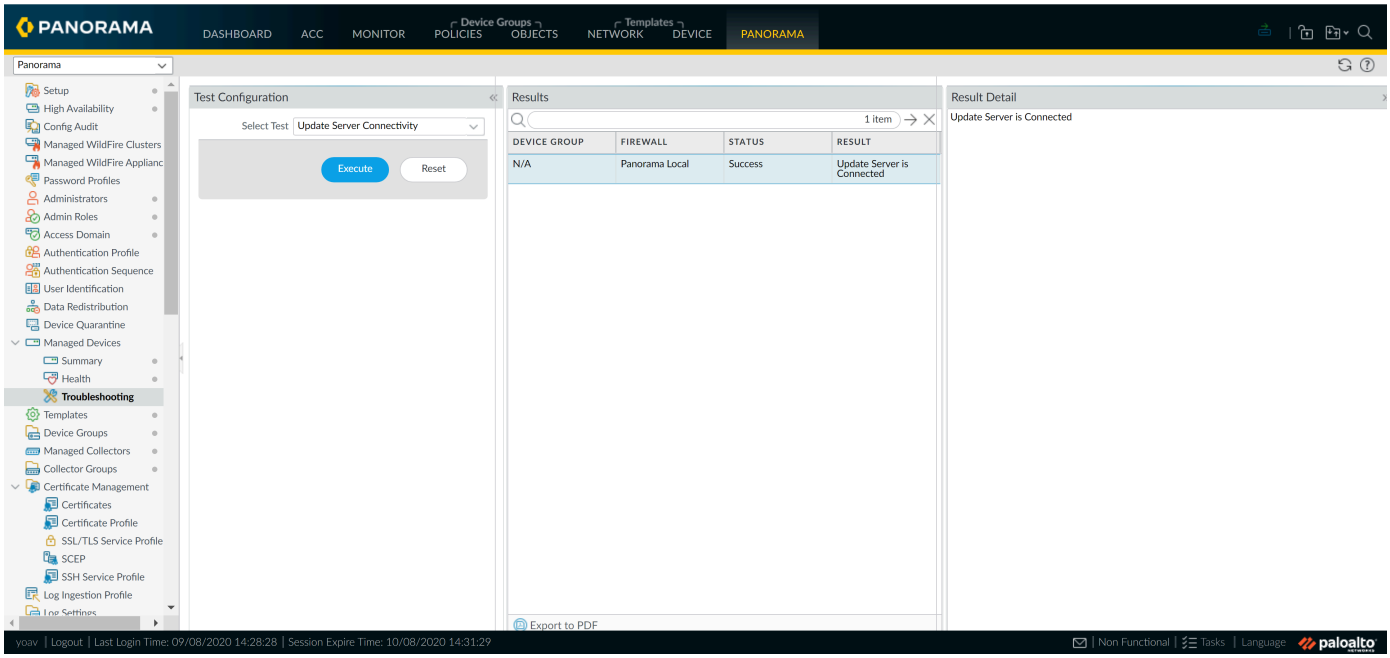
1. Enter the following commands, where <*Panorama-IP*> is the IP address you want to assign to the Panorama management interface, <*netmask*> is the subnet mask, <*gateway-IP*> is the IP address of the network gateway, and <*DNS-IP*> is the IP address of the DNS server:

```
> configure
# set deviceconfig system ip-address <Panorama-IP>
 netmask <netmask> default-gateway <gateway-IP> dns-setting
 servers primary <DNS-IP>
# commit
```

 ©2025 Palo Alto Networks, Inc.

```
# exit
```

2. Troubleshoot Connectivity to Network Resources to verify network access to external services required for firewall management, such as the default gateway, DNS server, and the Palo Alto Networks Update Server, as shown in the following example:



STEP 4 | Configure the general settings.

1. Using a secure connection (HTTPS) from a web browser, log in to the Panorama web interface using the IP address and password you assigned to the management interface (https://<IP address>).
2. Select **Panorama** > **Setup** > **Management** and edit the General Settings.
3. Enter a **Hostname** for the server and enter the network **Domain** name. The domain name is just a label; Panorama doesn't use it to join the domain.
4. Align the clock on Panorama and the managed firewalls to use the same **Time Zone**, for example GMT or UTC. If you plan to use the Strata Logging Service, you must configure NTP so that Panorama can stay in sync with the Strata Logging Service.

   Timestamps are recorded when Panorama receives the logs and the managed firewalls generate the logs. Aligning the time zones on Panorama and the firewalls ensures that the timestamps are synchronized and the process of querying logs and generating reports on Panorama is harmonious.
5. Enter the **Latitude** and **Longitude** to enable accurate placement of the Panorama management server on the world map.
6. Enter the **Serial Number** you received in the order fulfillment email.
7. Click **OK** to save your changes.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | (Optional) Modify the management interface settings.

To configure connectivity to Panorama using an IPv6 IP address, you must configure both an IPv4 and IPv6 to successfully configure Panorama using an IPv6 IP address. Panorama does not support configuring the management interface with only an IPv6 IP address.

1. Select **Panorama** > **Setup** > **Interfaces** and click **Management**.
2. If your firewalls connect to the Panorama management server using a public IP address that is translated to a private IP address (NAT), enter the public IP in the **Public IP Address** field, and the private IP in the **IP Address** field to push both addresses to your firewalls.
3. Select which Network Connectivity Services to allow on the interface (such as **SSH** access).

   *Don't select **Telnet** or **HTTP**. These services use plaintext and are less secure than the other services.*
4. Click **OK** to save your changes to the interface.

STEP 6 | Commit your configuration changes.

Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 7 | Next steps...

1. If necessary, Expand Log Storage Capacity on the Panorama Virtual Appliance.
2. (Best Practice) Replace the default certificate that Panorama uses to secure HTTPS traffic over the management (MGT) interface.
3. Activate a Panorama Support License.
4. Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected.
5. Install Content and Software Updates for Panorama.
6. Set Up Administrative Access to Panorama

Set Up The Panorama Virtual Appliance as a Log Collector

If you want a dedicated virtual appliance for log collection, configure a Panorama virtual appliance on ESXi, Alibaba Cloud, AWS, AWS GovCloud, Azure, Google Cloud Platform, KVM, Hyper-V, or Oracle Cloud Infrastructure (OCI) in Log Collector mode. To do this, you first perform the initial configuration of the virtual appliance in Panorama mode, which includes licensing, installing software and content updates, and configuring the management (MGT) interface. You then switch the Panorama virtual appliance to Log Collector mode and complete the Log Collector configuration. Additionally, if you want to use dedicated M-Series Appliance Interfaces (recommended) instead of the MGT interface for log collection and Collector Group communication, you must first configure the interfaces for the Panorama management server, then configure them for the Log Collector, and then perform a Panorama commit followed by a Collector Group commit.

 ©2025 Palo Alto Networks, Inc.

Perform the following steps to set up a new virtual appliance as a Log Collector or to convert an existing virtual appliance that was previously deployed as a Panorama management server.

Switching the virtual appliance from Panorama mode to Log Collector mode reboots the appliance, deletes the local Log Collector, deletes any existing log data, and deletes all configurations except the management access settings. Switching the mode does not delete licenses, software updates, or content updates.

STEP 1 | Set up the Panorama virtual appliance management server that will manage the Log Collector if you have not already done so.

Perform one of the following tasks:

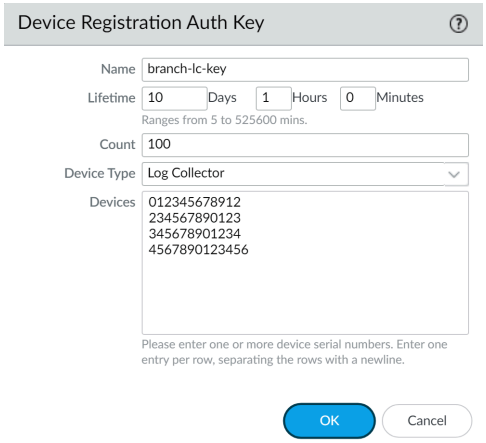
- Set Up the Panorama Virtual Appliance
- Set Up the M-Series Appliance

 ©2025 Palo Alto Networks, Inc.

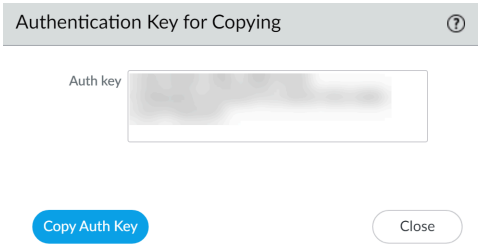
STEP 2 | On the Panorama management server, create a device registration authentication key to securely add the Dedicated Log Collector to Panorama management.

1. Log in to the Panorama Web Interface.
2. Select **Panorama** > **Device Registration Auth Key** and **Add** a new authentication key.
3. Configure the authentication key.
   - **Name**—Add a descriptive name for the authentication key.
   - **Lifetime**—Specify the key lifetime for how long you can use the authentication key to onboard new Log Collectors.
   - **Count**—Specify how many times you can use the authentication key to onboard new Log Collectors.
   - **Device Type**—Specify that this authentication key is used to authenticate only a **Log Collector**.

     *You can select **Any** to use the device registration authentication key to onboard firewalls, Log Collectors, and WildFire appliances.*
   - (Optional) **Devices**—Enter one or more device serial numbers to specify for which Log Collectors the authentication key is valid.
4. Click **OK**.



5. **Copy Auth Key** and **Close**.



 ©2025 Palo Alto Networks, Inc.

STEP 3 | Record the management IP addresses of the Panorama management server.

If you deployed Panorama in a high availability (HA) configuration, you need the IP address of each HA peer.

1. Log in to the web interface of the Panorama management server.
2. Record the **IP Address** of the solitary (non-HA) or active (HA) Panorama by selecting **Panorama** > **Setup** > **Management** and checking the Management Interface Settings.
3. For an HA deployment, record the **Peer HA IP Address** of the passive Panorama by selecting **Panorama** > **High Availability** and checking the Setup section.

STEP 4 | Set up the Panorama virtual appliance that will serve as a Dedicated Log Collector.

If you previously deployed this appliance as a Panorama management server, you can skip this step because the MGT interface is already configured and the licenses and updates are already installed.

The Panorama virtual appliance in Log Collector mode does not have a web interface for configuration tasks, only a CLI. Therefore, before changing the mode on the Panorama virtual appliance, use the web interface in Panorama mode to:

1. Set up the Panorama virtual appliance in one of the following supported hypervisors:
   - Install Panorama on an ESXi Server
   - Install Panorama on Alibaba Cloud
   - Install Panorama on AWS
   - Install Panorama on AWS GovCloud
   - Install Panorama on Azure
   - Install Panorama on Google Cloud Platform
   - Install Panorama on Hyper-V
   - Set Up Panorama on Oracle Cloud Infrastructure (OCI)
2. Perform Initial Configuration of the Panorama Virtual Appliance.
3. Register Panorama and Install Licenses.
4. Install Content and Software Updates for Panorama.

STEP 5 | (Panorama on Azure only) Modify the admin password.

The Dedicated Log Collector supports only the admin Administrator user in order to change in to Log Collector mode. Modify the admin password to allow you to log in using the admin Administrator user.

1. Log in to the Panorama Web Interface.
2. Select **Panorama** > **Administrators** and select **admin**.
3. Enter the **Password**, **Confirm Password** and click **OK**.
4. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | (Panorama on AWS and Azure only) Delete all users, except for the admin user.

1. Log in to the Panorama Web Interface as admin.
2. Select **Panorama** > **Administrators**.
3. Select the existing Administrators, except admin, and **Delete**.
4. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 7 | Log in to the Panorama CLI.

STEP 8 | Switch from Panorama mode to Log Collector mode.

1. Switch to Log Collector mode by entering the following command:

```
> request system system-mode logger
```

2. Enter **Y** to confirm the mode change. The virtual appliance reboots. If the reboot process terminates your terminal emulation software session, reconnect to the virtual appliance to see the Panorama login prompt.

   *If you see a **CMS Login** prompt, this means the Log Collector has not finished rebooting. Press Enter at the prompt without typing a username or password.*

3. Log back in to the CLI.
4. Verify that the switch to Log Collector mode succeeded:

```
> show system info | match system-mode
```

   If the mode change succeeded, the output displays:

```
system-mode: logger
```

STEP 9 | Enable connectivity between the Log Collector and Panorama management server.

Enter the following commands at the Log Collector CLI, where <*IPaddress1*> is for the MGT interface of the solitary (non-HA) or active (HA) Panorama and <*IPaddress2*> is for the MGT interface of the passive (HA) Panorama, if applicable.

```
> configure
# set deviceconfig system panorama-server <IPaddress1> panorama-
server-2 <IPaddress2>
# commit
# exit
```

STEP 10 | Add the device registration authentication key to the Dedicated Log Collector.

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



 ©2025 Palo Alto Networks, Inc.

STEP 11 | Record the serial number of the Log Collector.

You need the serial number to add the Log Collector as a managed collector on the Panorama management server.

1. At the Log Collector CLI, enter the following command to display its serial number.

```
> show system info | match serial
```

2. Record the serial number.

STEP 12 | Add the Log Collector as a managed collector to the Panorama management server.

1. Select **Panorama** > **Managed Collectors** and **Add** a managed collector.
2. In the **General** settings, enter the serial number (**Collector S/N**) you recorded for the Log Collector.
3. In the **Panorama Server IP** field, enter the IP address or FQDN of the solitary (non-HA) or active (HA) Panorama. For HA deployments, enter the IP address or FQDN of the passive Panorama peer in the **Panorama Server IP 2** field.

   These IP addresses must specify a Panorama interface that has **Device Management and Device Log Collection** services enabled. By default, these services are enabled only on the MGT interface. However, you might have enabled the services on other interfaces when you Set Up the M-Series Appliance that is a Panorama management server.
4. Select **Interfaces**, click **Management**, and enter the **Public IP Address** of the Dedicated Log Collector.
5. Click **OK** twice to save your changes to the Log Collector.
6. Select **Commit** > **Commit to Panorama** and **Commit** your changes to the Panorama configuration.
7. Verify that **Panorama** > **Managed Collectors** lists the Log Collector you added. The Connected column displays a check mark to indicate that the Log Collector is connected to Panorama. You might have to wait a few minutes before the page displays the updated connection status.

   At this point, the Configuration Status column displays Out of Sync and the Run Time Status column displays disconnected. The status will change to In Sync and connected after you configure a Collector Group.

STEP 13 | Enable the logging disks.

1. Select **Panorama** > **Managed Collectors** and edit the Log Collector.
2. Select **Disks** and **Add** each disk.
3. Click **OK** to save your changes.
4. Select **Commit** > **Commit to Panorama** and **Commit** your changes to the Panorama configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 14 | (Recommended) Configure the **Ethernet1**, **Ethernet2**, **Ethernet3**, **Ethernet4**, and **Ethernet5** interfaces if the Panorama management server and Log Collector will use them for **Device Log Collection** (receiving logs from firewalls) and **Collector Group Communication**.

If you previously deployed the Log Collector as a Panorama management server and configured these interfaces, you must reconfigure them because switching to Log Collector mode would have deleted all configurations except the management access settings.

1. Configure each interface on the Panorama management server (other than the MGT interface) if you haven't already:
   1. Select **Panorama** > **Setup** > **Interfaces** and click the Interface Name.
   2. Select *<interface-name>* to enable the interface.
   3. Complete one or both of the following field sets based on the IP protocols of your network:
      - For ESXi
        - IPv4—**Public IP Address**, **IP Address**, **Netmask**, and **Default Gateway**

          IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
      - For Alibaba Cloud, AWS, Azure, GCP, and OCI
        - **Public IP address**
   4. Select the Device Management Services that the interface supports:

      Device Management and Device Log Collection—You can assign one or more interfaces.

      Collector Group Communication—You can assign only one interface.

      Device Deployment (software and content updates)—You can assign only one interface.
   5. Click **OK** to save your changes.
2. Configure each interface on the Log Collector (other than the MGT interface):
   1. Select **Panorama** > **Managed Collectors** and edit the Log Collector.
   2. Select **Interfaces** and click the name of the interface.
   3. Select *<interface-name>* to enable the interface.
   4. Complete one or both of the following field sets based on the IP protocols of your network:
      - For ESXi
        - IPv4—**Public IP Address**, **IP Address**, **Netmask**, and **Default Gateway**

          IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
      - For Alibaba Cloud, AWS, Azure, GCP, and OCI
        - **Public IP address**
   5. Select the Device Management Services that the interface supports:

      Device Log Collection—You can assign one or more interfaces.

      Collector Group Communication—You can assign only one interface.

 ©2025 Palo Alto Networks, Inc.

6. Click **OK** to save your changes to the interface.

3. Click **OK** to save your changes to the Log Collector.
4. Select **Commit** > **Commit to Panorama** and **Commit** your changes to the Panorama configuration.

STEP 15 | (Optional) If your deployment is using custom certificates for authentication between Panorama and managed devices, deploy the custom client device certificate. For more information, see Set Up Authentication Using Custom Certificates.

1. Select **Panorama** > **Certificate Management** > **Certificate Profile** and choose the certificate profile from the drop-down or click **New Certificate Profile** to create one.
2. Select **Panorama** > **Managed Collectors** > **Add** > **Communication** for a Log Collector.
3. Select the **Secure Client Communication** check box.
4. Select the type of device certificate the Type drop-down.
   - If you are using a local device certificate, select the **Certificate** and **Certificate Profile** from the respective drop-downs.
   - If you are using SCEP as the device certificate, select the **SCEP Profile** and **Certificate Profile** from the respective drop-downs.
5. Click **OK**.

 ©2025 Palo Alto Networks, Inc.

STEP 16 | (Optional) Configure Secure Server Communication Sn a Log Collector. For more information, see Set Up Authentication Using Custom Certificates.

1. Select **Panorama** > **Managed Collectors** > **Add** > **Communication**.
2. Verify that the **Custom Certificate Only** check box is not selected. This allows you to continue managing all devices while migrating to custom certificates.

   When the Custom Certificate Only check box is selected, the Log Collector does not authenticate and cannot receive logs from devices using predefined certificates.
3. Select the SSL/TLS service profile from the **SSL/TLS Service Profile** drop-down. This SSL/TLS service profile applies to all SSL connections between the Log Collector and devices sending it logs.
4. Select the certificate profile from the **Certificate Profile** drop-down.
5. Select **Authorize Client Based on Serial Number** to have the server check clients against the serial numbers of managed devices. The client certificate must have the special keyword $UDID set as the CN to authorize based on serial numbers.
6. In **Disconnect Wait Time (min)**, enter the number of minutes Panorama should wait before breaking and reestablishing the connection with its managed devices. This field is blank by default and the range is 0 to 44,640 minutes.

   The disconnect wait time does not begin counting down until you commit the new configuration.
7. (Optional) Configure an authorization list.
   1. Click **Add** under Authorization List.
   2. Select the **Subject** or **Subject Alt Name** as the Identifier type.
   3. Enter an identifier of the selected type.
   4. Click **OK**.
   5. Select **Check Authorization List** to enforce the authorization list.
8. Click **OK**.
9. Select **Commit** > **Commit to Panorama**.

STEP 17 | Assign the Log Collector to a Collector Group.

1. Configure a Collector Group. You must perform a Panorama commit and then a Collector Group commit to synchronize the Log Collector configuration with Panorama and to put

 ©2025 Palo Alto Networks, Inc.

the Eth1, Eth2, Eth3, Eth4, and Eth5 interfaces (if you configured them) in an operational state on the Log Collector.

In any single Collector Group, all the Log Collectors must run on the same Panorama model: all M-700 appliances, all M-600 appliances, all M-500 appliances, all M-300 appliances, all M-200 appliances, or all Panorama virtual appliances.

*As a best practice, **Enable log redundancy across collectors** if you add multiple Log Collectors to a single Collector group. This option requires each Log Collector to have the same number of logging disks.*

2. Select **Panorama** > **Managed Collectors** to verify that the Log Collector configuration is synchronized with Panorama.

   The Configuration Status column should display In Sync and the Run Time Status column should display connected.

3. Access the Log Collector CLI and enter the following command to verify that its interfaces are operational:

   ```
   > show interface all
   ```

   The output displays the `state` as `up` for each interface that is operational.

4. If the Collector Group has multiple Log Collectors, Troubleshoot Connectivity to Network Resources to verify they can communicate with each other by performing a Ping connectivity test for each interface that the Log Collectors use. For the `source` IP address, specify the interface of one of the Log Collectors. For the `host` IP address, specify the matching interface of another Log Collector in the same Collector Group.

STEP 18 | Next steps...

To enable the Log Collector to receive firewall logs:

1. Configure Log Forwarding to Panorama.
2. Verify Log Forwarding to Panorama.

Set Up the Panorama Virtual Appliance with Local Log Collector

If the Panorama virtual appliance is in Legacy mode after you upgrade from a Panorama 8.0 or earlier release to a Panorama 8.1 (or later) release, switch to Panorama mode in order to create a local Log Collector, add multiple logging disks without losing existing logs. increase log storage up to 24TB, and enable faster report generation.

Once you change from Legacy mode to Panorama mode, Legacy mode will no longer be available.

After upgrading to Panorama 8.1, the first step is to increase the system resources on the virtual appliance to the minimum required for Panorama mode. Panorama reboots when you increase resources, so perform this procedure during a maintenance window. You must install a larger system disk (81GB), increase CPUs and memory based on the log storage capacity, and add a virtual logging disk. The new logging disk must have at least as much capacity as the appliance

 ©2025 Palo Alto Networks, Inc.

currently uses in Legacy mode and cannot be less than 2TB. Adding a virtual disk enables you to migrate existing logs to the Log Collector and enables the Log Collector to store new logs.

If Panorama is deployed in an HA configuration, perform the following steps on the secondary peer first and then on the primary peer.

STEP 1 | Determine which system resources you need to increase before the virtual appliance can operate in Panorama mode.

You must run the command specified in this step even if you have determined that Panorama already has adequate resources.

1. Access the Panorama CLI:
   1. Use terminal emulation software such as PuTTY to open an SSH session to the IP address that you specified for the Panorama MGT interface.
   2. Log in to the CLI when prompted.
2. Check the resources you must increase by running the following command:

```
> request system system-mode panorama
```

Enter **y** when prompted to continue. The output specifies the resources you must increase. For example:

```
Panorama mode not supported on current system disk of size
 52.0 GB.
Please attach a disk of size 81.0 GB, then use 'request system
 clone-system-disk' to migrate the current system disk
Please add a new virtual logging disk with more than 50.00 GB
 of storage capacity.
Not enough CPU cores: Found 4 cores, need 8 cores
```

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Increase the CPUs and memory, and replace the system disk with a larger disk.

1. Access the VMware ESXi vSphere Client, select **Virtual Machines**, right-click the Panorama virtual appliance, and select **Power** > **Power Off**.
2. Right-click the Panorama virtual appliance and **Edit Settings**.
3. Select **Memory** and enter the new **Memory Size**.
4. Select **CPUs** and specify the number of CPUs (the **Number of virtual sockets** multiplied by the **Number of cores per socket**).
5. Add a virtual disk.

   You will use this disk to replace the existing system disk.

   1. In the **Hardware** settings, **Add** a disk, select **Hard Disk** as the hardware type, and click **Next**.
   2. **Create a new virtual disk** and click **Next**.
   3. Set the **Disk Size** to exactly 81GB and select the **Thick Provision Lazy Zeroed** disk format.
   4. Select **Specify a datastore or datastore structure** as the location, **Browse** to a datastore of at least 81GB, click **OK**, and click **Next**.
   5. Select a SCSI **Virtual Device Node** (you can use the default selection) and click **Next**.

      Panorama will fail to boot if you select a format other than SCSI.

   6. Verify that the settings are correct and then click **Finish** and **OK**.
6. Right-click the Panorama virtual appliance and select **Power** > **Power On**. Wait for Panorama to reboot before continuing.
7. Return to the Panorama CLI and copy the data from the original system disk to the new system disk:

   ```
   > request system clone-system-disk target sdb
   ```

   Enter **y** when prompted to continue.

   The copying process takes around 20 to 25 minutes, during which Panorama reboots. When the process finishes, the output tells you to shut down Panorama.
8. Return to the vSphere Client console, right-click the Panorama virtual appliance, and select **Power** > **Power Off**.
9. Right-click the Panorama virtual appliance and **Edit Settings**.
10. Select the original system disk, click **Remove**, select **Remove from virtual machine**, and click **OK**.
11. Right-click the Panorama virtual appliance and **Edit Settings**.
12. Select the new system disk, set the **Virtual Device Node** to **SCSI (0:0)**, and click **OK**.
13. Right-click the Panorama virtual appliance and select **Power** > **Power On**. Before proceeding, wait for Panorama to reboot on the new system disk (around 15 minutes).

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Add a virtual logging disk.

This is the disk to which you will migrate existing logs.

1. In the VMware ESXi vSphere Client, right-click the Panorama virtual appliance and select **Power** > **Power Off**.
2. Right-click the Panorama virtual appliance and **Edit Settings**.
3. Repeat the steps to Add a virtual disk. Set the **Disk Size** to a multiple of 2TB based on the amount of log storage you need. The capacity must be at least as large as the existing virtual disk or NFS storage that Panorama currently uses for logs. The disk capacity must be a multiple of 2TB and can be up to 24TB. For example, if the existing disk has 5TB of log storage, you must add a new disk of at least 6TB.

   After you switch to Panorama mode, Panorama will automatically divide the new disk into 2TB partitions, each of which will function as a separate virtual disk.
4. Right-click the Panorama virtual appliance and select **Power** > **Power On**. Wait for Panorama to reboot before continuing.

STEP 4 | Switch from Legacy mode to Panorama mode.

After switching the mode, the appliance reboots again and then automatically creates a local Log Collector and Collector Group. The existing logs won't be available for querying or reporting until you migrate them later in this procedure.

1. Return to the Panorama CLI and run the following command.

   ```
   > request system system-mode panorama
   ```

   Enter **y** when prompted to continue. After rebooting, Panorama automatically creates a local Log Collector (named Panorama) and creates a Collector Group (named default) to contain it. Panorama also configures the virtual logging disk you added and divides it into separate 2TB disks. Wait for the process to finish and for Panorama to reboot (around five minutes) before continuing.
2. Log in to the Panorama web interface.
3. In the **Dashboard**, **General Information** settings, verify that the **Mode** is now **panorama**.

   In an HA deployment, the secondary peer is in a suspended state at this point because its mode (Panorama) does not match the mode on the primary peer (Legacy). You will un-

 ©2025 Palo Alto Networks, Inc.

suspend the secondary peer after switching the primary peer to Panorama mode later in this procedure.

4. Select **Panorama** > **Collector Groups**to verify that the **default** collector group has been created, and that the local Log Collector is part of the default collector group.
5. Push the configuration to the managed devices.
   - If there are no pending changes:
     1. Select **Commit** > **Push to Devices** and **Edit Selections**.
     2. Select **Collector Group** and make sure the **default** collector group is selected.
     3. Click **OK** and **Push**.
   - If you have pending changes:
     1. Select **Commit** > **Commit and Push** and **Edit Selections**.
     2. Verify that your **Device Group** devices and **Templates** are included.
     3. Select **Collector Group** and make sure the **default** collector group is selected.
     4. Click **OK** and **Commit and Push**.
6. Select **Panorama** > **Managed Collectors** and verify that the columns display the following information for the local Log Collector:
   - Collector Name—This defaults to the Panorama hostname. It should be listed under the **default** Collector Group.
   - Connected—Check mark
   - Configuration Status—In sync
   - Run Time Status—connected

STEP 5 | (HA only) Switch the primary Panorama from Legacy mode to Panorama mode.

This step triggers failover.

1. Repeat Step 1 through Step 4 on the primary Panorama.

   Wait for the primary Panorama to reboot and return to an active HA state. If preemption is not enabled, you must manually fail back: select **Panorama** > **High Availability** and, in the Operational Commands section, **Make local Panorama functional**.
2. On the primary Panorama, select **Dashboard** and, in the High Availability section, **Sync to peer**, click **Yes**, and wait for the **Running Config** to display **Synchronized** status.
3. On the secondary Panorama, select **Panorama** > **High Availability** and, in the Operational Commands section, **Make local Panorama functional**.

   This step is necessary to bring the secondary Panorama out of its suspended HA state.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Migrate existing logs to the new virtual logging disks.

If you deployed Panorama in an HA configuration, perform this only on the primary peer.

Palo Alto Networks recommends migrating existing logs to the new virtual logging disks during your maintenance window. The log migration requires a large number of the Panorama virtual appliance CPU cores to execute and impacts Panorama operational performance.

1. Return to the Panorama CLI.
2. Start the log migration:

   ```
   > request logdb migrate vm start
   ```

   The process duration varies by the volume of log data you are migrating. To check the status of the migration, run the following command:

   ```
   > request logdb migrate vm status
   ```

   When the migration finishes, the output displays: `migrationhas been done`.
3. Verify that the existing logs are available.
   1. Log in to the Panorama web interface.
   2. Select **Panorama** > **Monitor**, select a log type that you know matches some existing logs (for example, **Panorama** > **Monitor** > **System**), and verify that the logs display.

STEP 7 | Next steps...

Configure log forwarding to Panorama so that the Log Collector receives new logs from firewalls.

Set up a Panorama Virtual Appliance in Panorama Mode

Panorama mode allows the Panorama™ virtual appliance to operate as a Panorama management server with local log collection capabilities. By default, the Panorama virtual appliance is deployed in Panorama mode when at least one virtual logging disk is attached to a Panorama virtual appliance.

While still supported, switching from Legacy mode with a 50GB logging disk to Panorama mode is not recommended for production environments. If you switch to Panorama mode with a 50GB logging disk, you are unable to add additional logging disks.

STEP 1 | Log in to the Panorama CLI.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Switch to Panorama mode.

1. Change to Panorama mode:

```
> request system system-mode panorama
```

2. Enter **Y** to confirm the mode change. The Panorama virtual appliance reboots. If the reboot process terminates your terminal emulation software session, reconnect to the Panorama virtual appliance to see the Panorama login prompt.

   If you see a `CMS Login` prompt, this means the Panorama virtual appliance has not finished rebooting. Press Enter at the prompt without typing a username or password.

STEP 3 | Verify that the switch to Panorama mode succeeded.

1. Log back in to the CLI.
2. Verify that the switch to Panorama mode succeeded:

```
> show system info | match system-mode
```

   If the mode change succeeded, the output displays:

```
> system mode:panorama
```

Set up a Panorama Virtual Appliance in Management Only Mode

Management Only mode allows the Panorama virtual appliance to operate strictly as a Panorama management server without local log collection capabilities. By default, the Panorama virtual appliance is in Panorama mode for the initial deployment. It is recommended to change the Panorama virtual appliance to Management Only immediately after the initial deployment because changing to Management Only mode requires that there are no logs being forwarded to the Panorama management server because the Panorama virtual appliance in Management Only mode does not support log collection. After you change to Management Only mode, any existing log data stored on the Panorama virtual appliance becomes inaccessible, and the ACC and reporting features cannot query the logs stored on the Panorama virtual appliance.

(Panorama in Legacy mode) There is no impact to the Panorama virtual appliance when you change the Panorama virtual appliance from Legacy mode to Management Only mode. As a precaution, Palo Alto Networks recommends taking a virtual machine snapshot of your Panorama virtual appliance that you can use to restore Panorama in the event of unexpected impact.

*If you configured a local Log Collector, the local Log Collector still exists on Panorama when you change to Management Only mode despite having no log collection capabilities. Deleting the local Log Collector (**Panorama** > **Managed Collectors**) deletes the Eth1/1 interface configuration the local Log Collector uses by default. If you decide to delete the local Log Collector, you must reconfigure the Eth1/1 interface.*

STEP 1 | Log in to the Panorama CLI.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Switch to Management Only mode.

1. Change to Management Only mode:

```
> request system system-mode management-only
```

2. Enter **Y** to confirm the mode change. The Panorama virtual appliance reboots. If the reboot process terminates your terminal emulation software session, reconnect to the Panorama virtual appliance to see the Panorama login prompt.

   If you see a `CMS Login` prompt, this means the Panorama virtual appliance has not finished rebooting. Press Enter at the prompt without typing a username or password.

STEP 3 | Verify that the switch to Management Only mode succeeded.

1. Log back in to the CLI.
2. Verify that the switch to Management Only mode succeeded:

```
> show system info | match system-mode
```

   If the mode change succeeded, the output displays:

```
> system mode:management-only
```

Expand Log Storage Capacity on the Panorama Virtual Appliance

After you Perform Initial Configuration of the Panorama Virtual Appliance, the available log storage capacity and the options for expanding it depend on the virtual platform (VMware ESXi, vCloud Air, Alibaba Cloud, AWS, AWS GovCloud, Azure, Google Cloud Platform, KVM, Hyper-V, or OCI) and mode (Legacy, Panorama, or Log Collector mode): see Panorama Models for details.

To expand the log storage capacity on the Panorama virtual appliance, you must add additional logging disks. Expanding the log storage capacity of an existing logging disk is not supported, and Panorama does not recognize the additional storage capacity. For example; if you added a 2TB logging disk, and then expanded that existing logging disk to 4TB, Panorama continues to recognize the logging disk as having 2TB of storage capacity and ignores the additional 2TB of storage capacity.

For additional log storage, you can also forward firewall logs to Dedicated Log Collectors (see Configure a Managed Collector) or Configure Log Forwarding from Panorama to External Destinations.

Before expanding log storage capacity on Panorama, Determine Panorama Log Storage Requirements.

- Preserve Existing Logs When Adding Storage on Panorama Virtual Appliance in Legacy Mode
- Add a Virtual Disk to Panorama on an ESXi Server
- Add a Virtual Disk to Panorama on vCloud Air
- Add a Virtual Disk to Panorama on Alibaba Cloud
- Add a Virtual Disk to Panorama on AWS

 ©2025 Palo Alto Networks, Inc.

- Add a Virtual Disk to Panorama on Azure
- Add a Virtual Disk to Panorama on Google Cloud Platform
- Add a Virtual Disk to Panorama on KVM
- Add a Virtual Disk to Panorama on Hyper-V
- Add a Virtual Disk to Panorama on Oracle Cloud Infrastructure (OCI)
- Mount the Panorama ESXi Server to an NFS Datastore

Preserve Existing Logs When Adding Storage on Panorama Virtual Appliance in Legacy Mode

The Panorama virtual appliance in Legacy mode can use only one virtual disk for logging. Therefore, if you add a virtual disk that is dedicated for logging, Panorama stops using the default 11GB log storage on the system disk and automatically copies any existing logs to the new logging disk. (Panorama continues using the system disk for data other than logs.)

If you replace an existing dedicated logging disk of up to 2TB storage capacity with a disk of up to 8TB, you will lose the logs on the existing disk. To preserve the logs, your choices are:

- Configure log forwarding to external destinations before you replace the virtual disk.
- Set up a new Panorama virtual appliance for the new 8TB disk and maintain access to the Panorama containing the old disk for as long as you need the logs. To forward firewall logs to the new Panorama virtual appliance, one option is to reconfigure the firewalls to connect with the new Panorama IP address (select **Device** > **Setup** > **Management** and edit the Panorama Settings), add the firewalls as managed devices to the new Panorama, and Configure Log Forwarding to Panorama. To reuse the old Panorama IP address on the new Panorama, another option is to export the configuration of the old Panorama and then import and load the configuration on the new Panorama.
- Copy logs from the old disk to the new disk. Copying can take several hours, depending on how many logs the disk currently stores, and Panorama cannot collect logs during the process. Contact Palo Alto Networks Customer Support for instructions.

Add a Virtual Disk to Panorama on an ESXi Server

To expand log storage capacity on the Panorama virtual appliance, you can add virtual logging disks. If the appliance is in Panorama mode, you can add 1 to 12 virtual logging disks of 2TB each or 1 24TB logging disk, for a maximum total of 24TB. If the appliance is in Legacy mode, you can add one virtual logging disk of up to 8TB on ESXi 5.5 and later versions or one disk of up to 2TB on earlier ESXi versions. Additionally, it is recommended to add logging disks with the same disk provisioning format to avoid any unexpected performance that may arise from having multiple disk with different provisioning formats.

If Panorama loses connectivity to the new virtual disk, Panorama might lose logs during the failure interval.

To allow for redundancy, use the virtual disk in a RAID configuration. RAID10 provides the best write performance for applications with high logging characteristics.

If necessary, you can Replace the Virtual Disk on an ESXi Server.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Add additional disks to Panorama

In all modes, the first logging disk on the Panorama VM must be at least 2TB in order to add additional disks. If the first logging disk is smaller than 2TB, you will be unable to add additional disk space.

1. Access the VMware vSphere Client and select **Virtual Machines**.
2. Right-click the Panorama virtual appliance and select **Power** > **Power off**.
3. Right-click the Panorama virtual appliance and select **Edit Settings**.
4. Click **Add** in the **Hardware** tab to launch the Add Hardware wizard.
5. Select **Hard Disk** as the hardware type and click **Next**.
6. **Create a new virtual disk** and click **Next**.
7. Set the **Disk Size**. If the Panorama virtual appliance is in Panorama mode, set the size to at least 2TB. If the appliance is in Legacy mode, you can set the size to as much as 8TB.

   In Panorama mode, you can add disk sizes larger than 2TB and Panorama will automatically create as many 2TB partitions as possible. For example, if disk sdc was 24TB, it will create 12 2TB partitions. These disks will be named sdc1-12.
8. Select the **Disk Provisioning** format and click **Next**.
9. **Specify a datastore or datastore structure**, **Browse** to a datastore with enough space for the specified **Disk Size**, click **OK**, and click **Next**.
10. Select a SCSI **Virtual Device Node** (you can use the default selection) and click **Next**.

    The selected node must be in SCSI format; Panorama will fail to boot if you select another format.
11. Verify that the settings are correct and then click **Finish** and **OK**.

    The new disk appears in the list of devices for the virtual appliance.
12. Repeat Step 4 through Step 11 to add additional disks to the Panorama virtual appliance if necessary.
13. Right click the Panorama virtual appliance and select **Power** > **Power On**. The virtual disk initializes for first-time use. The size of the new disk determines how long initialization takes.

STEP 2 | Configure each disk.

The following example uses the sdc virtual disk.

1. Log in to the Panorama CLI.
2. Enter the following command to view the disks on the Panorama virtual appliance:

   `show system disk details`

   The user will see the following response:

```
 Name
: sdb
 State : Present
 Size : 2048000 MB
```

 ©2025 Palo Alto Networks, Inc.

```
Status : Available
Reason : Admin enabled
Name : sdc
State : Present
Size : 2048000 MB
Status : Available
Reason : Admin disabled
```

3. Enter the following command and confirm the request when prompted for all disks with the `Reason : Admin disabled` response:

   `request system disk add sdc`

   *The **`request system disk add`** command is not available on a Panorama management server in Management Only mode because logging is not supported in this mode. If you do not see the command, Set up a Panorama Virtual Appliance in Panorama Mode to enable the logging disks. Once in Panorama mode, Log in to the Panorama CLI and continue to Step 4 to verify the disk addition.*

4. Enter the **`show system disk details`** command to verify the status of the disk addition. Continue to Step 3 when all newly added disk responses display `Reason : Admin enabled`.

STEP 3 | Make disks available for logging.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Managed Collectors** and edit the Log Collector.
3. Select **Disks** and Add each newly added disk.
4. Click **OK**.
5. Select **Commit** > **Commit to Panorama**.

   For Panorama in an Active/Passive high availability (HA) configuration, wait for HA sync to complete before continuing.

6. Select **Commit** > **Push to Devices** and push the changes to the Collector Group the Log Collector belongs to.

STEP 4 | Configure Panorama to receive logs.

This step is intended for new Panorama deployments in Panorama mode. If you are adding logging disks to an existing Panorama virtual appliance, continue to Step 5.

1. Configure a Managed Collector.
2. Configure a Collector Group.
3. Configure Log Forwarding to Panorama.

STEP 5 | Verify that the Panorama Log Storage capacity has been increased.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Collector Groups** and select the Collector Group that the Panorama virtual appliance belongs to.
3. Verify that the **Log Storage** capacity accurately displays the disk capacity.

 ©2025 Palo Alto Networks, Inc.

Add a Virtual Disk to Panorama on vCloud Air

You can add virtual logging disks to expand log storage capacity on the Panorama™ virtual appliance. If the appliance is in Panorama mode, you can add 1 to 12 virtual logging disks of 2TB each or 1 24TB logging disk, for a maximum total of 24TB. If the appliance is in Legacy mode, you can add one virtual logging disk of up to 8TB.

If Panorama loses connectivity to the new virtual disk, Panorama might lose logs for the duration of the failure.

If necessary, you can Replace the Virtual Disk on vCloud Air.

STEP 1 | Add additional disks to Panorama.

In all modes, the first logging disk on the Panorama VM must be at least 2TB to add additional disks. If the first logging disk is less than 2TB, you will be unable to add additional disk space.

1. Access the vCloud Air web console and select your **Virtual Private Cloud On Demand** region.
2. Select the Panorama virtual appliance in the **Virtual Machines** tab.
3. **Add another disk** (**Actions** > **Edit Resources**).
4. Set the **Storage** size. If the Panorama virtual appliance is in Panorama mode, set the size to at least 2TB. If the appliance is in Legacy mode, you can set the size to as much as 8TB.

   In Panorama mode, you can add disk sizes larger than 2TB and Panorama will automatically create as many 2TB partitions as possible. For example, if disk sdc was 24TB, Panorama will create 12 2TB partitions. These disks will be named sdc1 through sdc12.
5. Set the storage tier to **Standard** or **SSD-Accelerated**.
6. Repeat the previous steps to add additional disks to the Panorama virtual appliance as needed.
7. **Save** your changes.

STEP 2 | Configure each disk.

The following example uses the sdc virtual disk.

1. Log in to the Panorama CLI.
2. Enter the following command to view the disks on the Panorama virtual appliance:

   `show system disk details`

   The user will see the following response:

```
 Name
: sdb
 State : Present
 Size : 2048000 MB
 Status : Available
```

 ©2025 Palo Alto Networks, Inc.

```
Reason : Admin enabled
Name : sdc
State : Present
Size : 2048000 MB
Status : Available
Reason : Admin disabled
```

3. Enter the following command and confirm the request when prompted for all disks with the `Reason : Admin disabled` response:

   `request system disk add sdc`

   > *The **`request system disk add`** command is not available on a Panorama management server in Management Only mode because logging is not supported in this mode. If you do not see the command, Set up a Panorama Virtual Appliance in Panorama Mode to enable the logging disks. Once in Panorama mode, Log in to the Panorama CLI and continue to Step 4 to verify the disk addition.*

4. Enter the **`show system disk details`** command to verify the status of the disk addition. Continue to the next step when all newly added disk responses display `Reason : Admin enabled`.

STEP 3 | Make disks available for logging.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Managed Collectors** and edit the Log Collector.
3. Select **Disks** and **Add** each new disk.
4. Click **OK**.
5. Select **Commit** > **Commit to Panorama**.

   > *For Panorama in an Active/Passive high availability (HA) configuration, wait for HA sync to complete before continuing.*

6. Select **Commit** > **Push to Devices** and push the changes to the Collector Group the Log Collector belongs to.

STEP 4 | Configure Panorama to receive logs.

This step is intended for new Panorama deployments in Panorama mode. If you are adding logging disks to an existing virtual Panorama appliance, continue to the next step.

1. Configure a Managed Collector.
2. Configure a Collector Group.
3. Configure Log Forwarding to Panorama.

STEP 5 | Verify that the Panorama Log Storage capacity has been increased.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Collector Groups** and select the Collector Group to which the virtual Panorama appliance belongs.
3. Verify that the **Log Storage** capacity accurately displays your new disk capacity.

 ©2025 Palo Alto Networks, Inc.

Add a Virtual Disk to Panorama on Alibaba Cloud

After you Install Panorama on Alibaba Cloud, add additional virtual logging disks to expand log storage capacity on the Panorama™ virtual appliance for logs generated by managed firewalls. You can add virtual disks to a local Log Collector for a Panorama virtual appliance in Panorama mode or for a Dedicated Log Collector. To add virtual disks, you must have access to the Alibaba Cloud Console, the Panorama command-line interface (CLI), and the Panorama web interface.

The Panorama virtual appliance on Alibaba Cloud supports only 2TB logging disks and, in total, supports up to 24TB of log storage. You cannot add a logging disk smaller than 2TB or a logging disk of a size that is not evenly divisible by 2TB because the Panorama virtual appliance partitions logging disks in to 2TB partitions. For example, if you attach a 4TB logging disk, Panorama will create two 2TB partitions. However, you cannot add a 5TB logging disk because the leftover 1TB is not supported as a partition.

STEP 1 | Log in to the Alibaba Cloud Console.

STEP 2 | Select **Elastic Compute Service** > **Instances & Images** > **Instances** and navigate to the Panorama virtual appliance instance.

STEP 3 | Add a virtual logging disk to Panorama.

In all modes, the first logging disk on the Panorama VM must be at least 2TB in order to add additional disks. If the first logging disk is smaller than 2TB, you will be unable to add additional disk space.

1. In the Actions column, select **Manage**.
2. Select **Cloud Disk** and **Create Disk**.
3. Configure the virtual logging disk.
   - **Attach**—Select **Attach to ECS Instance**.
   - **ECS Instance**—Select the region and the Panorama virtual appliance instance.
   - **Storage**—Select type of virtual disk and enter the disk capacity.
   - (Optional) **Quantity**—Specify how many virtual disks to create. By default, 1 virtual disk is created. When creating multiple logging disks, be sure that the sum of all virtual disks does not exceed 24TB.
   - **Terms of Service**—Review the Alibaba Cloud Terms of Service and check after you have reviewed.
4. **Preview** the virtual disk creation.
5. **Create** the new virtual disk.

   A status window displays after you create the new virtual disk. After the virtual disk is successfully created, **Go to the Disk List** to confirm the disk is successfully created.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure each disk.

The following example uses the sdc virtual disk.

1. Log in to the Panorama CLI.
2. Enter the following command to view the disks on the Panorama virtual appliance:
   `show system disk details`

   The user will see the following response:

```
 Name : sdb
State : Present
Size : 2048000 MB
Status : Available
Reason : Admin disabled
```

3. Enter the following command and confirm the request when prompted for all disks with the `Reason : Admin disabled` response:
   `request system disk add sdc`

   *The **`request system disk add`** command is not available on a Panorama management server in Management Only mode because logging is not supported in this mode. If you do not see the command, Set up a Panorama Virtual Appliance in Panorama Mode to enable the logging disks. Once in Panorama mode, log in to the Panorama CLI and continue to the next step to verify the disk addition.*
4. Enter the **`show system disk details`** command to verify the status of the disk addition. Continue to the next step when all newly added disk responses display `Reason : Admin enabled`.

STEP 5 | Make disks available for logging.

1. Log in to the Panorama web interface.
2. Edit a Log Collector (**Panorama** > **Managed Collectors**).
3. Select **Disks** and **Add** each newly added disk.
4. Click **OK**.
5. Select **Commit** > **Commit to Panorama**.

   For Panorama in an Active/Passive high availability (HA) configuration, wait for HA sync to complete before continuing.
6. Select **Commit** > **Push to Devices** and push the changes to the Collector Group the Log Collector belongs to.

STEP 6 | (New Panorama deployments in Panorama mode only) Configure Panorama to receive logs.

If you are adding logging disks to an existing Panorama virtual appliance, skip to step 6.

1. Configure a Collector Group.
2. Configure Log Forwarding to Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Verify that the Panorama Log Storage capacity is increased.

1. Log in to the Panorama web interface.
2. Select the Collector Group to which the Panorama virtual appliance belongs (**Panorama** > **Collector Groups**).
3. Verify that the **Log Storage** capacity accurately displays the disk capacity.

Add a Virtual Disk to Panorama on AWS

After you Install Panorama on AWS or Install Panorama on AWS GovCloud, add virtual logging disks to the Panorama™ virtual appliance instance to provide storage for logs generated by managed firewalls. You can add virtual disks to a local log Collector for a Panorama virtual appliance in Panorama mode or for a Dedicated Log Collector. To add virtual disks, you must have access to the Amazon Web Service Console, the Panorama command-line interface (CLI), and the Panorama web interface.

The Panorama virtual appliance on AWS supports only 2TB logging disks and, in total, supports up to 24TB of log storage. You cannot add a logging disk smaller than 2TB or a logging disk of a size that is not evenly divisible by 2TB because the Panorama virtual appliance partitions logging disks in to 2TB partitions. For example, if you attach a 4TB logging disk, Panorama will create two 2TB partitions. However, you cannot add a 5TB logging disk because the leftover 1TB is not supported as a partition.

STEP 1 | Log in to AWS Web Service console and select the EC2 Dashboard.

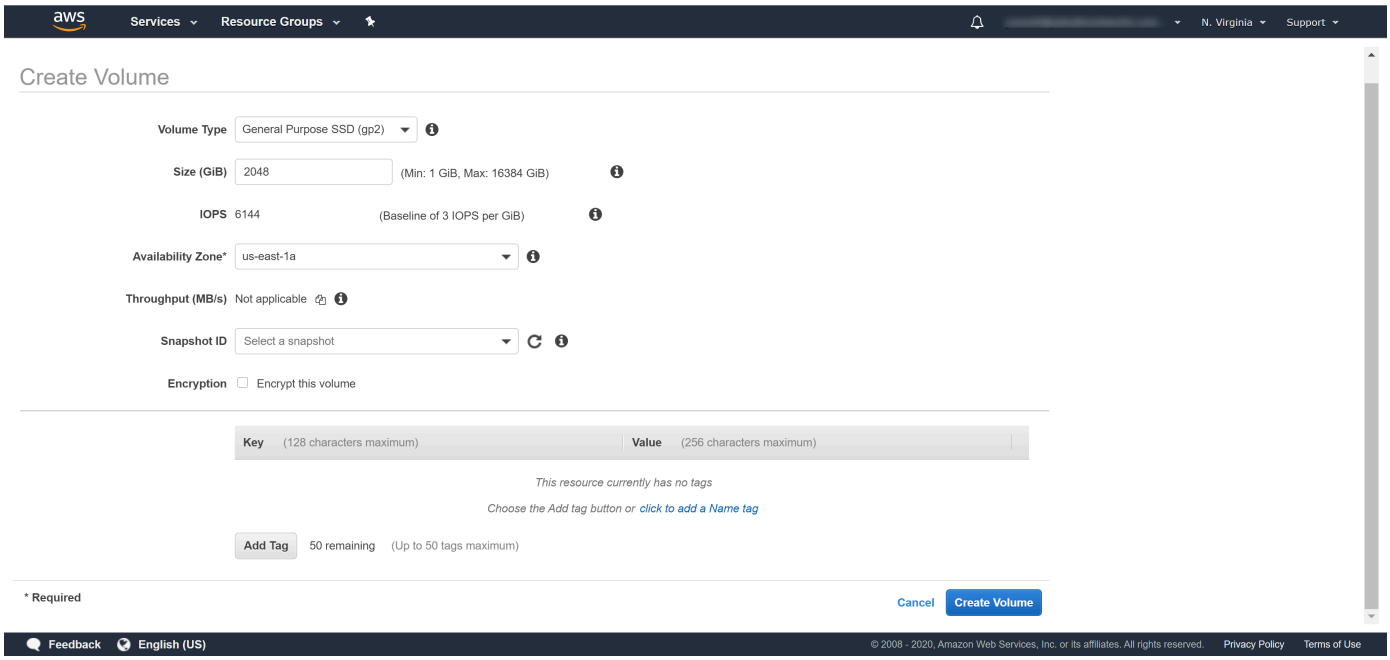
- Amazon Web Service Console
- AWS GovCloud Web Service Console

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Add a virtual logging disk to Panorama.

In all modes, the first logging disk on the Panorama VM must be at least 2TB in order to add additional disks. If the first logging disk is smaller than 2TB, you will be unable to add additional disk space.

1. On the EC2 Dashboard, select **Volumes** and **Create Volume**:
   - Select your preferred Volume Type. For general purpose use, select **General Purpose SSD (GP2)**.
   - Configure the **Size** of the volume as 2048 GiB.
   - Select the same Availability Zone that your Panorama virtual appliance instance is located in.
   - (Optional) Encrypt the volume.
   - (Optional) Add tags to your volume.
2. Click **Create Volume**.



3. In the Volumes page, select the volume you, and select **Actions** > **Attach Volume**.
4. Attach the Panorama virtual appliance Instance.
   1. Select your Panorama **Instance**.
   2. Specify the **Device name** for the logging disk volume you created.

STEP 3 | Configure each disk.

The following example uses the sdc virtual disk.

1. Log in to the Panorama CLI.
2. Enter the following command to view the disks on the Panorama virtual appliance:
   `show system disk details`

   The user will see the following response:

```
Name : nvme1n1
State : Present
Size : 2048000 MB
Status : Available
Reason : Admin enabled
Name : nvme2n1
State : Present
Size : 2048000 MB
Status : Available
Reason : Admin disabled
```

3. Enter the following command and confirm the request when prompted for all disks with the `Reason : Admin disabled` response:
   `request system disk add nvme2n1`

   *The **`request system disk add`** command is not available on a Panorama management server in Management Only mode because logging is not supported in this mode. If you do not see the command, Set up a Panorama Virtual Appliance in Panorama Mode to enable the logging disks. Once in Panorama mode, Log in to the Panorama CLI and continue to Step 4 to verify the disk addition.*
4. Enter the **`show system disk details`** command to verify the status of the disk addition. Continue to the next step when all newly added disk responses display `Reason : Admin enabled`.

STEP 4 | Make disks available for logging.

1. Log in to the Panorama web interface.
2. Edit a Log Collector (**Panorama** > **Managed Collectors**).
3. Select **Disks** and **Add** each newly added disk.
4. Click **OK**.
5. Select **Commit** > **Commit to Panorama**.

   For Panorama in an Active/Passive high availability (HA) configuration, wait for HA sync to complete before continuing.
6. Select **Commit** > **Push to Devices** and push the changes to the Collector Group the Log Collector belongs to.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | (New Panorama deployments in Panorama mode only) Configure Panorama to receive logs.

If you are adding logging disks to an existing Panorama virtual appliance, skip to step 6.

1. Configure a Collector Group.
2. Configure Log Forwarding to Panorama.

STEP 6 | Verify that the Panorama Log Storage capacity is increased.

1. Log in to the Panorama web interface.
2. Select the Collector Group to which the Panorama virtual appliance belongs (**Panorama** > **Collector Groups**).
3. Verify that the **Log Storage** capacity accurately displays the disk capacity.

Add a Virtual Disk to Panorama on Azure

After you Install Panorama on Azure, add virtual logging disks to the Panorama™ virtual appliance instance to provide storage for logs generated by managed firewalls. You can add virtual disks to a local log Collector for a Panorama virtual appliance in Panorama mode or for a Dedicated Log Collector. To add virtual disks, you must have access to the Microsoft Azure portal, the Panorama command-line interface (CLI), and the Panorama web interface.

The Panorama virtual appliance on Azure supports only 2TB logging disks and, in total, supports up to 24TB of log storage. You cannot add a logging disk smaller than 2TB or a logging disk of a size that is not evenly divisible by 2TB because the Panorama virtual appliance partition logging disks in to 2TB partitions. For example, if you attach a 4TB logging disk, Panorama will create two 2TB partitions. However, you cannot add a 5TB logging disk because the leftover 1TB is not supported as a partition.

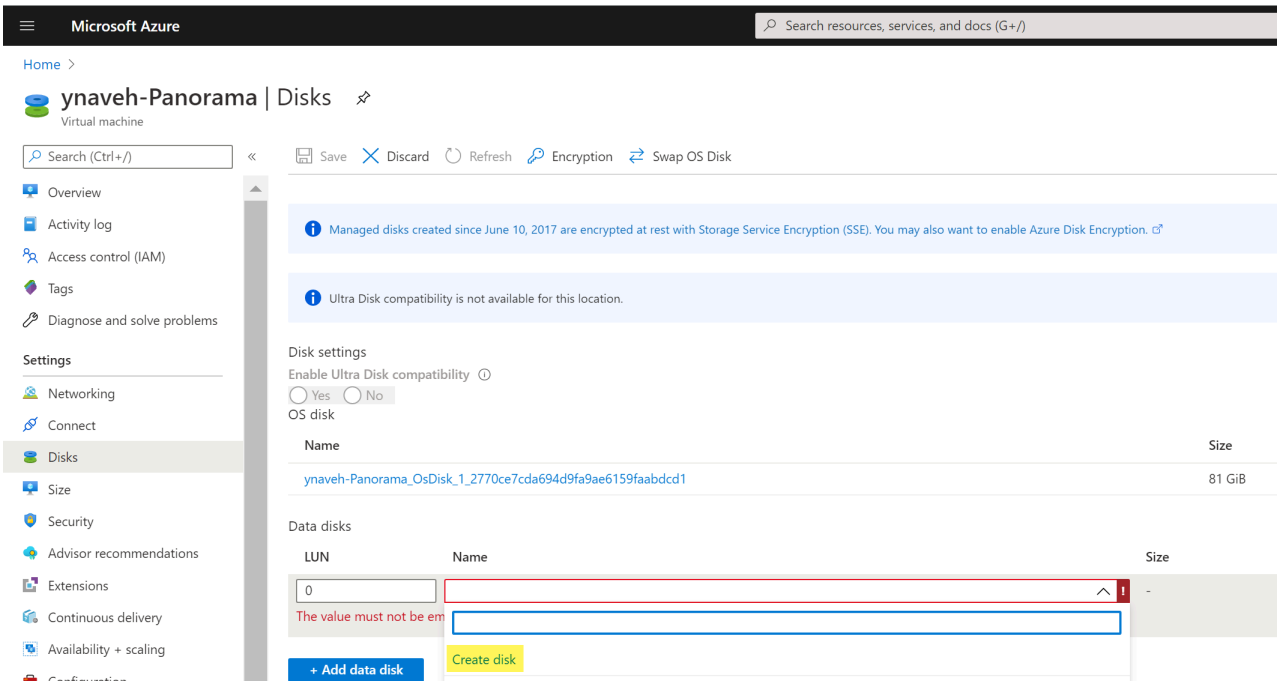
STEP 1 | Log in to the Microsoft Azure portal.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Add a virtual logging disk to Panorama.

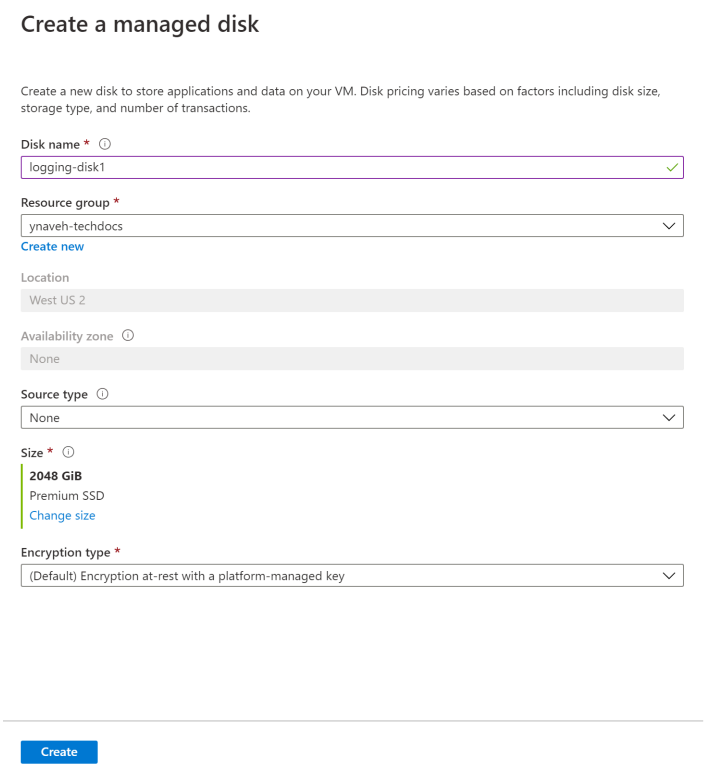
In all modes, the first logging disk on the Panorama VM must be at least 2TB in order to add additional disks. If the first logging disk is smaller than 2TB, you will be unable to add additional disk space.

1. In the Azure Dashboard, select the Panorama **Virtual Machines** to which you want to add a logging disk.
2. Select **Disks**.
3. **+Add data disk**.
4. In the drop-down for the new disk, **Create disk**.

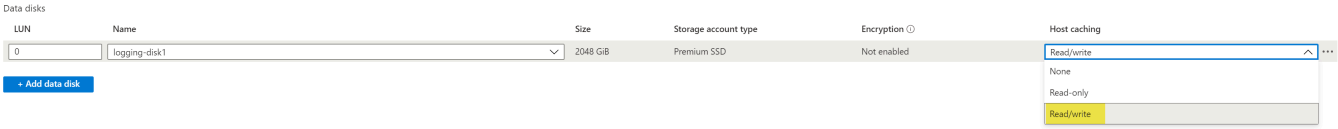


5. Configure the logging disk.
   1. Enter the disk **Name**.
   2. Select the Resource group. If you **Create new** resource groups, enter the group name.
   3. Verify the **Account type** (this field is automatically populated).
   4. In the **Source type** drop-down, select **None**.
   5. Select **Change Size** and select a 2048 GiB logging disk.
   6. **Create** the new logging disk.

 ©2025 Palo Alto Networks, Inc.



7. For the **Host caching**, select **Read/write**.



STEP 3 | Enable each disk.

The following example uses the sdc virtual disk.

1. Log in to the Panorama CLI.
2. Enter the following command to view the disks on the Panorama virtual appliance:

   `show system disk details`

   The user will see the following response:

```
 Name
: sdb
 State : Present
 Size : 2048000 MB
 Status : Available
 Reason : Admin enabled
 Name : sdc
 State : Present
 Size : 2048000 MB
 Status : Available
```

 ©2025 Palo Alto Networks, Inc.

```
Reason : Admin disabled
```

3. Enter the following command and confirm the request when prompted for all disks with the `Reason : Admin disabled` response:
   `request system disk add sdc`

   > *The **`request system disk add`** command is not available on a Panorama management server in Management Only mode because logging is not supported in this mode. If you do not see the command, Set up a Panorama Virtual Appliance in Panorama Mode to enable the logging disks. Once in Panorama mode, Log in to the Panorama CLI and continue to Step 4 to verify the disk addition.*

4. Enter the **`show system disk details`** command to verify the status of the disk addition. Continue to the next step when all newly added disk responses display `Reason : Admin enabled`.

STEP 4 | Make disks available for logging.

1. Log in to the Panorama web interface.
2. Edit a Log Collector (**Panorama** > **Managed Collectors**)
3. Select **Disks** and **Add** each newly added disk.
4. Click **OK**.
5. Select **Commit** > **Commit to Panorama**.

   > *For Panorama in an Active/Passive high availability (HA) configuration, wait for HA sync to complete before continuing.*

6. Select **Commit** > **Push to Devices** and push the changes to the Collector Group the Log Collector belongs to.

STEP 5 | (New Panorama deployments in Panorama mode only) Configure Panorama to receive logs.

If you are adding logging disks to an existing Panorama virtual appliance, skip to step 6.

1. Configure a Collector Group.
2. Configure Log Forwarding to Panorama.

STEP 6 | Verify that the Panorama Log Storage capacity is increased.

1. Log in to the Panorama web interface.
2. Select the Collector Group to which the Panorama virtual appliance belongs (**Panorama** > **Collector Groups**).
3. Verify that the **Log Storage** capacity accurately displays the disk capacity.

Add a Virtual Disk to Panorama on Google Cloud Platform

After you Install Panorama on Google Cloud Platform, add virtual logging disks to the Panorama™ virtual appliance instance to provide storage for logs generated by managed firewalls. You can add virtual disks to a local log Collector for a Panorama virtual appliance in Panorama mode or for a Dedicated Log Collector. The Panorama virtual appliance on Google Cloud Platform supports only 2TB logging disks and, in total, supports up to 24TB of log storage. You cannot add a logging

 ©2025 Palo Alto Networks, Inc.

disk smaller than 2TB or a logging disk of a size that is not evenly divisible by 2TB because the Panorama virtual appliance partitions logging disks in to 2TB partitions. For example, if you attach a 4TB logging disk, Panorama will create two 2TB partitions. However, you cannot add a 5TB logging disk because the leftover 1TB is not supported as a partition.

STEP 1 | Log in to the Google Cloud Console.

STEP 2 | Add the virtual logging disk.



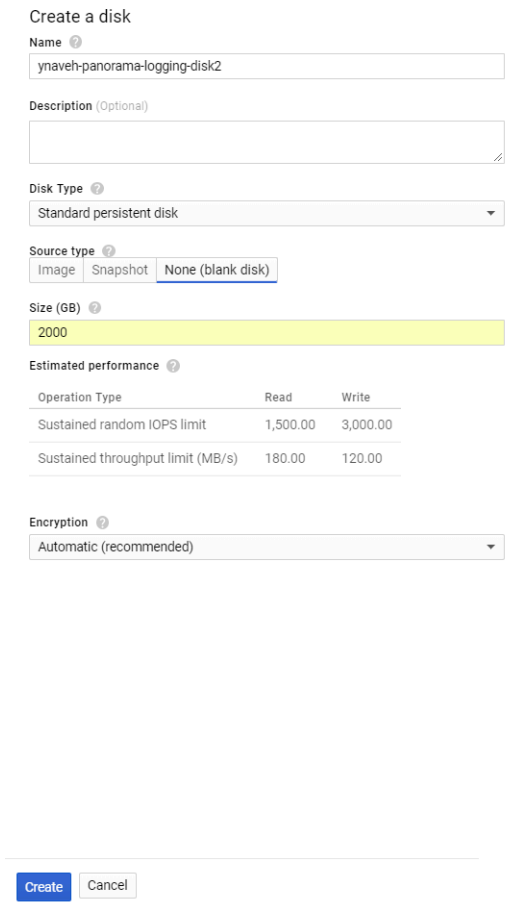
In all modes, the first logging disk on the Panorama VM must be at least 2TB in order to add additional disks. If the first logging disk is smaller than 2TB, you will be unable to add additional disk space.

1. In the Products & Services menu, select and then **Edit** the Panorama virtual appliance instance (**Compute Engine** > **VM Instances**).
2. In the Additional Disks section, **Add Item**.
3. **Create disk** (**Name** drop-down).

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Configure the virtual logging disks.

1. Enter the **Name**.
2. Expand the **Disk Type** drop-down menu and select the desired type.
3. For the **Source type**, select **None (blank disk)**.
4. Set the **Size (GB)** of the virtual logging disk.
5. Click **Create**.



6. **Save** the changes to update the Panorama virtual appliance instance.

STEP 4 | Configure each disk.

The following example uses the sdc virtual disk.

1. Log in to the Panorama CLI.
2. Enter the following command to view the disks on the Panorama virtual appliance:
   `show system disk details`

   The user will see the following response:

```
 Name
: sdb
 State : Present
 Size : 2048000 MB
```

 ©2025 Palo Alto Networks, Inc.

```
Status : Available
Reason : Admin enabled
Name : sdc
State : Present
Size : 2048000 MB
Status : Available
Reason : Admin disabled
```

3. Enter the following command and confirm the request when prompted for all disks with the `Reason : Admin disabled` response:

   `request system disk add sdc`

   *The **`request system disk add`** command is not available on a Panorama management server in Management Only mode because logging is not supported in this mode. If you do not see the command, Set up a Panorama Virtual Appliance in Panorama Mode to enable the logging disks. Once in Panorama mode, Log in to the Panorama CLI and continue to Step 4 to verify the disk addition.*

4. Enter the **`show system disk details`** command to verify the status of the disk addition. Continue to the next step when all newly added disk responses display `Reason : Admin enabled`.

STEP 5 | Make disks available for logging.

1. Log in to the Panorama web interface.
2. Edit a Log Collector (**Panorama** > **Managed Collectors**).
3. Select **Disks** and **Add** each newly added disk.
4. Click **OK**.
5. Select **Commit** > **Commit to Panorama**.

   For Panorama in an Active/Passive high availability (HA) configuration, wait for HA sync to complete before continuing.

6. Select **Commit** > **Push to Devices** and push the changes to the Collector Group the Log Collector belongs to.

STEP 6 | (New Panorama deployments in Panorama mode only) Configure Panorama to receive logs.

If you are adding logging disks to an existing Panorama virtual appliance, skip to step 7.

1. Configure a Collector Group.
2. Configure Log Forwarding to Panorama.

STEP 7 | Verify that the Panorama Log Storage capacity is increased.

1. Log in to the Panorama web interface.
2. Select the Collector Group to which the Panorama virtual appliance belongs (**Panorama** > **Collector Groups**).
3. Verify that the **Log Storage** capacity accurately displays the disk capacity.

 ©2025 Palo Alto Networks, Inc.

Add a Virtual Disk to Panorama on KVM

After you Install Panorama on KVM, add virtual logging disks to the Panorama™ virtual appliance instance to provide storage for logs generated by managed firewalls. You can add virtual disks to a local log Collector for a Panorama virtual appliance in Panorama mode or for a Dedicated Log Collector. The Panorama virtual appliance on KVM supports only 2TB logging disks and, in total, supports up to 24TB of log storage. You cannot add a logging disk smaller than 2TB or a logging disk of a size that is not evenly divisible by 2TB because the Panorama virtual appliance partitions logging disks in to 2TB partitions. For example, if you attach a 4TB logging disk, Panorama will create two 2TB partitions. However, you cannot add a 5TB logging disk because the leftover 1TB is not supported as a partition.

STEP 1 | **Shutdown** the Panorama virtual appliance instance on the Virtual Machine Manager.

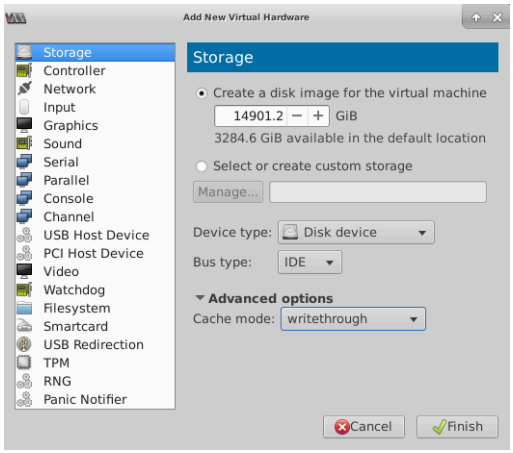
STEP 2 | Double-click the Panorama virtual appliance instance in the Virtual Machine Manager and **Show virtual hardware details**.

STEP 3 | Add the virtual logging disk. Repeat this step as many times as needed.

In all modes, the first logging disk on the Panorama VM must be at least 2TB in order to add additional disks. If the first logging disk is smaller than 2TB, you will be unable to add additional disk space.

1. **Create a disk image for a virtual image** (**Add Hardware** > **Storage**) and configure the virtual disk storage capacity to the appropriate 2TB value:2000GB or 14901.2GiB depending on your Virtual Machine Manager.

   Depending on the version, some Virtual Machine Managers use GiB (gibibyte) to allocate memory. Be sure you correctly convert the required storage capacity to avoid under provisioning the virtual logging disk and sending the Panorama virtual appliance into maintenance mode.
2. In the **Device type** drop-down, select **Disk device**.
3. In the **Bus type** drop-down, select **VirtIO** or **IDE** based on your configuration.
4. Expand **Advanced options** and, in the **Cache mode** drop-down, select **writethrough**.
5. Click **Finish**.



STEP 4 | **Power on** the Panorama virtual appliance instance.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Configure each disk.

The following example uses the sdc virtual disk.

1. Log in to the Panorama CLI.
2. Enter the following command to view the disks on the Panorama virtual appliance:
   `show system disk details`

   The user will see the following response:

```
 Name
: sdb
 State : Present
 Size : 2048000 MB
 Status : Available
 Reason : Admin enabled
 Name : sdc
 State : Present
 Size : 2048 MB
 Status : Available
 Reason : Admin disabled
```

3. Enter the following command and confirm the request when prompted for all disks with the `Reason : Admin disabled` response:
   `request system disk add sdc`

   *The **`request system disk add`** command is not available on a Panorama management server in Management Only mode because logging is not supported in this mode. If you do not see the command, Set up a Panorama Virtual Appliance in Panorama Mode to enable the logging disks. Once in Panorama mode, Log in to the Panorama CLI and continue to Step 4 to verify the disk addition.*
4. Enter the **`show system disk details`** command to verify the status of the disk addition. Continue to the next step when all newly added disk responses display `Reason : Admin enabled`.

STEP 6 | Make disks available for logging.

1. Log in to the Panorama web interface.
2. Edit a Log Collector (**Panorama** > **Managed Collectors**).
3. Select **Disks** and **Add** each newly added disk.
4. Click **OK**.
5. Select **Commit** > **Commit to Panorama**.

   For Panorama in an Active/Passive high availability (HA) configuration, wait for HA sync to complete before continuing.
6. Select **Commit** > **Push to Devices** and push the changes to the Collector Group the Log Collector belongs to.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | (New Panorama deployments in Panorama mode only) Configure Panorama to receive logs.

If you are adding logging disks to an existing Panorama virtual appliance, skip to step 8.

1. Configure a Collector Group.
2. Configure Log Forwarding to Panorama.

STEP 8 | Verify that the Panorama Log Storage capacity is increased.

1. Log in to the Panorama web interface.
2. Select the Collector Group to which the Panorama virtual appliance belongs (**Panorama** > **Collector Groups**).
3. Verify that the **Log Storage** capacity accurately displays the disk capacity.

Add a Virtual Disk to Panorama on Hyper-V

After you Install Panorama on Hyper-V, add virtual logging disks to the Panorama™ virtual appliance instance to provide storage for logs generated by managed firewalls. You can add virtual disks to a local log Collector for a Panorama virtual appliance in Panorama mode or for a Dedicated Log Collector. The Panorama virtual appliance on Hyper-V supports only 2TB logging disks and, in total, supports up to 24TB of log storage. You cannot add a logging disk smaller than 2TB or a logging disk of a size that is not evenly divisible by 2TB because the Panorama virtual appliance partitions logging disks in to 2TB partitions. For example, if you attach a 4TB logging disk, Panorama will create two 2TB partitions. However, you cannot add a 5TB logging disk because the leftover 1TB is not supported as a partition.

STEP 1 | Power off the Panorama virtual appliance.

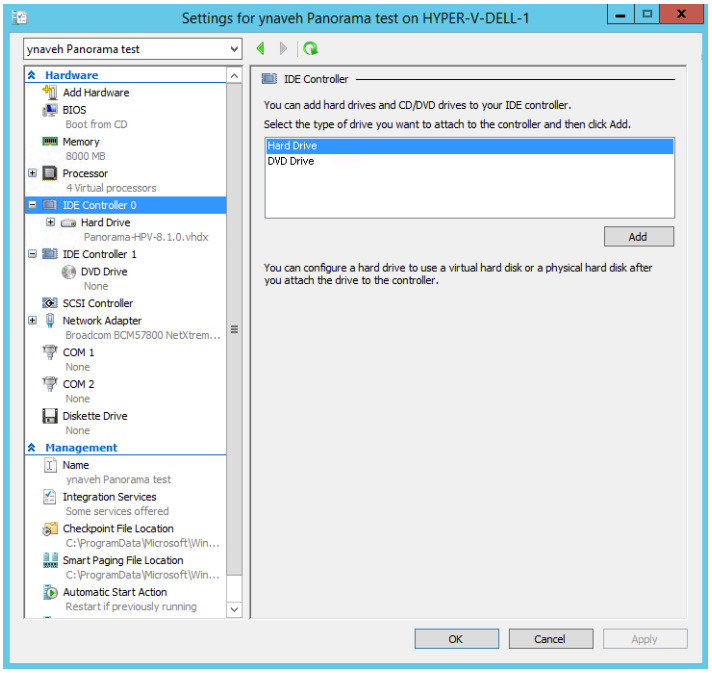
1. On the Hyper-V Manager, select the Panorama virtual appliance instance from the list of **Virtual Machines**.
2. Select **Action** > **Turn Off** to power off the Panorama virtual appliance.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Add the virtual logging disk. Repeat this step as many times as needed.

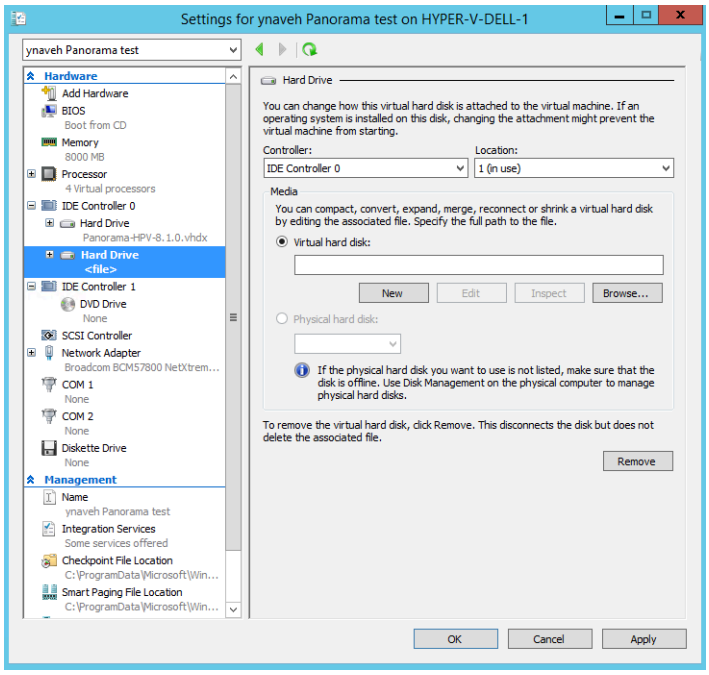
In all modes, the first logging disk on the Panorama VM must be at least 2TB in order to add additional disks. If the first logging disk is smaller than 2TB, you will be unable to add additional disk space.

1. Select the Panorama virtual appliance from the list of **Virtual Machines**, and select **Action** > **Settings**.
2. In the **Hardware** list, select **IDE Controller 0**.
3. From the **IDE Controller** drives list, select **Hard Drive** and **Add** the new virtual logging disk.



4. Select the new **Hard Drive** created under **IDE Controller 0**.
5. Under **Media**, add a **New** hard disk.

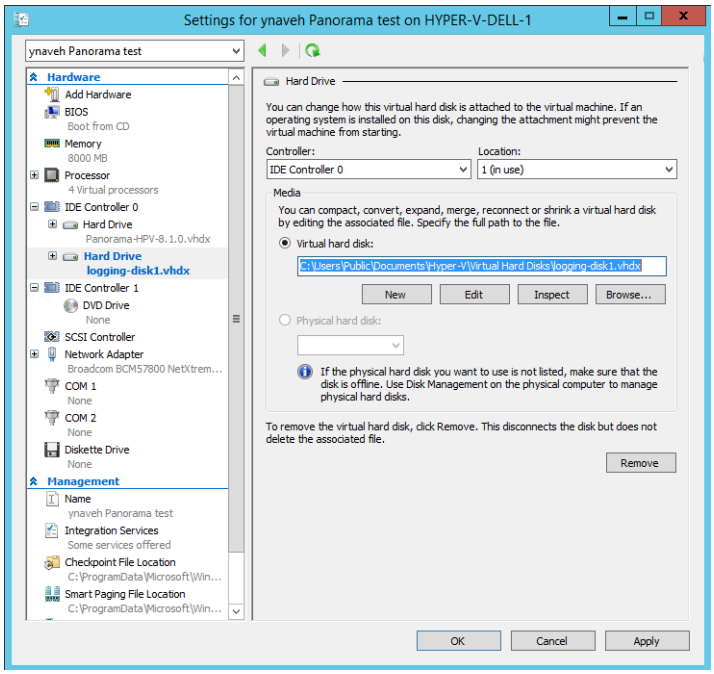
 ©2025 Palo Alto Networks, Inc.



 ©2025 Palo Alto Networks, Inc.

STEP 3 | Configure the new virtual logging disk.

1. If you see the Before You Begin prompt, click **Next** to begin adding the virtual logging disk
2. For the Disk Format, select **VHDX**. Click **Next** to continue
3. For the Disk Type, select **Fixed Size** or **Dynamically Expanding** based on your needs. Click **Next** to continue.
4. Specify the **Name** and **Location** for the virtual logging disk file. Click **Next** to continue.
5. To configure the disk, select **Create a new virtual hard disk** and enter the disk size. Click **Next** to continue.
6. Review the Summary and **Finish** adding the virtual hard logging disk.
7. **Apply** the new hard disk addition.



STEP 4 | Power on the Panorama virtual appliance.

1. Select the Panorama virtual appliance instance from the list of **Virtual Machines**.
2. Select **Action** > **Start** to power on the Panorama virtual appliance.

STEP 5 | Configure each disk.

The following example uses the sdc virtual disk.

1. Log in to the Panorama CLI.
2. Enter the following command to view the disks on the Panorama virtual appliance:
   `show system disk details`

   The user will see the following response:

```
 Name
: sdb
 State : Present
```

 ©2025 Palo Alto Networks, Inc.

```
Size : 2048000 MB
Status : Available
Reason : Admin enabled
Name : sdc
State : Present
Size : 2048 MB
Status : Available
Reason : Admin disabled
```

3. Enter the following command and confirm the request when prompted for all disks with the `Reason : Admindisabled` response:

   `request system disk add sdc`

   *The **`request system disk add`** command is not available on a Panorama management server in Management Only mode because logging is not supported in this mode. If you do not see the command, Set up a Panorama Virtual Appliance in Panorama Mode to enable the logging disks. Once in Panorama mode, Log in to the Panorama CLI and continue to Step 4 to verify the disk addition.*

4. Enter the **`show system disk details`** command to verify the status of the disk addition. Continue to the next step when all newly added disk responses display `Reason : Admin enabled`.

STEP 6 | Make disks available for logging.

1. Log in to the Panorama web interface.
2. Edit a Log Collector (**Panorama** > **Managed Collectors**).
3. Select **Disks** and **Add** each newly added disk.
4. Click **OK**.
5. Select **Commit** > **Commit to Panorama**.

   For Panorama in an Active/Passive high availability (HA) configuration, wait for HA sync to complete before continuing.

6. Select **Commit** > **Push to Devices** and push the changes to the Collector Group the Log Collector belongs to.

STEP 7 | (New Panorama deployments in Panorama mode only) Configure Panorama to receive logs.

If you are adding logging disks to an existing Panorama virtual appliance, skip to Step 8.

1. Configure a Collector Group.
2. Configure Log Forwarding to Panorama.

STEP 8 | Verify that the Panorama Log Storage capacity is increased.

1. Log in to the Panorama web interface.
2. Select the Collector Group to which the Panorama virtual appliance belongs (**Panorama** > **Collector Groups**).
3. Verify that the **Log Storage** capacity accurately displays the disk capacity.

 ©2025 Palo Alto Networks, Inc.

Add a Virtual Disk to Panorama on Oracle Cloud Infrastructure (OCI)

After you Install Panorama on Oracle Cloud Infrastructure (OCI), add additional virtual logging disks to expand log storage capacity on the Panorama™ virtual appliance for logs generated by managed firewalls. You can add virtual disks to a local Log Collector for a Panorama virtual appliance in Panorama mode or for a Dedicated Log Collector. To add virtual disks, you must have access to the OCI console, the Panorama command-line interface (CLI), and the Panorama web interface.

The Panorama virtual appliance on OCI supports only 2TB logging disks and, in total, supports up to 24TB of log storage. You cannot add a logging disk smaller than 2TB or a logging disk of a size that is not evenly divisible by 2TB because the Panorama virtual appliance partitions logging disks in to 2TB partitions. For example, if you attach a 4TB logging disk, Panorama will create two 2TB partitions. However, you cannot add a 5TB logging disk because the leftover 1TB is not supported as a partition.

STEP 1 | Log in to the Oracle Cloud Infrastructure console.

STEP 2 | Create a 2TB block volume.

1. Select **Block Storage** > **Block Volumes** and **Create Block Volume**.
2. Enter a descriptive **Name** for the volume.
3. Select the same **Availability Domain** as the Panorama virtual appliance instance.
4. Select the **Custom** volume size.
5. For the Volume Size, enter **`2000`**.
6. **Create Block Volume**.

STEP 3 | Attach a virtual logging disk to the Panorama virtual appliance instance.

In all modes, the first logging disk on the Panorama VM must be at least 2TB in order to add additional disks. If the first logging disk is smaller than 2TB, you will be unable to add additional disk space.

1. Select **Compute** > **Instances** and click the name of the Panorama virtual appliance instance.
2. Under resources, select **Attached Block Volumes** and **Attach Block Volume**.
3. For the Volume, **Select volume** and select the virtual logging disk.
4. For the Attachment type, select **Paravirtualized**.
   This is required for the Panorama virtual appliance to recognize the virtual logging disk.
5. For the Access, select **Read/Write**.
6. **Attach** the virtual logging disk.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure each disk.

The following example uses the sdc virtual disk.

1. Log in to the Panorama CLI.
2. Enter the following command to view the disks on the Panorama virtual appliance:
   `show system disk details`

   The user will see the following response:

   ```
    Name : sdb
   State : Present
   Size : 2048000 MB
   Status : Unavailable
   Reason : Admin disabled
   ```
3. Enter the following command and confirm the request when prompted for all disks with the `Reason : Admin disabled` response:
   `request system disk add sdc`

   *The **`request system disk add`** command is not available on a Panorama management server in Management Only mode because logging is not supported in this mode. If you do not see the command, Set up a Panorama Virtual Appliance in Panorama Mode to enable the logging disks. Once in Panorama mode, Log in to the Panorama CLI and continue to the next step to verify the disk addition.*
4. Enter the **`show system disk details`** command to verify the status of the disk addition. Continue to the next step when all newly added disk responses display `Reason : Admin enabled`.

STEP 5 | Make disks available for logging.

1. Log in to the Panorama web interface.
2. Edit a Log Collector (**Panorama** > **Managed Collectors**).
3. Select **Disks** and **Add** each newly added disk.
4. Click **OK**.
5. Select **Commit** > **Commit to Panorama**.

   For Panorama in an Active/Passive high availability (HA) configuration, wait for HA sync to complete before continuing.
6. Select **Commit** > **Push to Devices** and push the changes to the Collector Group the Log Collector belongs to.

STEP 6 | (New Panorama deployments in Panorama mode only) Configure Panorama to receive logs.

If you are adding logging disks to an existing Panorama virtual appliance, skip to step 6.

1. Configure a Collector Group.
2. Configure Log Forwarding to Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Verify that the Panorama Log Storage capacity is increased.

1. Log in to the Panorama web interface.
2. Select the Collector Group to which the Panorama virtual appliance belongs (**Panorama** > **Collector Groups**).
3. Verify that the **Log Storage** capacity accurately displays the disk capacity.

Mount the Panorama ESXi Server to an NFS Datastore

When the Panorama virtual appliance in Legacy mode runs on an ESXi server, mounting to a Network File System (NFS) datastore enables logging to a centralized location and expanding the log storage capacity beyond what a virtual disk supports. (ESXi 5.5 and later versions can support a virtual disk of up to 8TB. Earlier ESXi versions support a virtual disk of up to 2TB.) Before setting up an NFS datastore in a Panorama high availability (HA) configuration, see Logging Considerations in Panorama HA.

The Panorama virtual appliance in Panorama mode does not support NFS.

STEP 1 | Select **Panorama** > **Setup** > **Operations** and, in the Miscellaneous section, click **Storage Partition Setup**.

STEP 2 | Set the **Storage Partition** type to **NFS V3**.

STEP 3 | Enter the IP address of the NFS **Server**.

STEP 4 | Enter the **Log Directory** path for storing the log files. For example, export/panorama.

STEP 5 | For the **Protocol**, select **TCP** or **UDP**, and enter the **Port** for accessing the NFS server.

To use NFS over TCP, the NFS server must support it. Common NFS ports are UDP/TCP 111 for RPC and UDP/TCP 2049 for NFS.

STEP 6 | For optimal NFS performance, in the **Read Size** and **Write Size** fields, specify the maximum size of the chunks of data that the client and server pass back and forth to each other. Defining a read/write size optimizes the data volume and speed in transferring data between Panorama and the NFS datastore.

STEP 7 | (Optional) Select **Copy On Setup** to copy the existing logs stored on Panorama to the NFS volume. If Panorama has a lot of logs, this option might initiate the transfer of a large volume of data.

STEP 8 | Click **Test Logging Partition** to verify that Panorama can access the NFS **Server** and **Log Directory**.

STEP 9 | Click **OK** to save your changes.

STEP 10 | Select **Commit** > **Commit to Panorama** and **Commit** your changes. Until you reboot, the Panorama virtual appliance writes logs to the local storage disk.

STEP 11 | Select **Panorama** > **Setup** > **Operations** and select **Reboot Panorama** in the Device Operations section. After rebooting, Panorama starts writing logs to the NFS datastore.

 ©2025 Palo Alto Networks, Inc.

Increase CPUs and Memory on the Panorama Virtual Appliance

When you Perform Initial Configuration of the Panorama Virtual Appliance, you specify the memory and number of CPUs based on whether the appliance is in Panorama mode or Management Only mode and based on the log storage capacity or number of managed firewalls. If you later add storage capacity or managed firewalls, you must also increase the memory and CPUs. A Panorama virtual appliance in Log Collector mode must meet the system requirements, and does not need to have the CPU and memory increased beyond the minimum requirement. Review the Setup Prerequisites for the Panorama Virtual Appliance for the CPU and memory requirements for each Panorama mode.

- Increase CPUs and Memory for Panorama on an ESXi Server
- Increase CPUs and Memory for Panorama on vCloud Air
- Increase CPUs and Memory for Panorama on Alibaba Cloud
- Increase CPUs and Memory for Panorama on AWS
- Increase CPUs and Memory for Panorama on Azure
- Increase CPUs and Memory for Panorama on Google Cloud Platform
- Increase CPUs and Memory for Panorama on KVM
- Increase CPUs and Memory for Panorama on Hyper-V
- Increase the CPUs and Memory for Panorama on Oracle Cloud Infrastructure (OCI)

Increase CPUs and Memory for Panorama on an ESXi Server

For the minimum CPUs and memory that Panorama requires, see Increase CPUs and Memory on the Panorama Virtual Appliance.

STEP 1 | Access the VMware vSphere Client and select **Virtual Machines**.

STEP 2 | Right-click the Panorama virtual appliance and select **Power** > **Power Off**.

STEP 3 | Right-click the Panorama virtual appliance and select **Edit Settings**.

STEP 4 | Select **Memory** and enter the new **Memory Size**.

STEP 5 | Select **CPUs** and specify the number of CPUs (the **Number of virtual sockets** multiplied by the **Number of cores per socket**).

STEP 6 | Click **OK** to save your changes.

STEP 7 | Right-click the Panorama virtual appliance and select **Power** > **Power On**.

Increase CPUs and Memory for Panorama on vCloud Air

For the minimum CPUs and memory that Panorama requires, see Increase CPUs and Memory on the Panorama Virtual Appliance.

STEP 1 | Access the vCloud Air web console and select your **Virtual Private Cloud OnDemand** region.

STEP 2 | In the **Virtual Machines** tab, select the Panorama virtual machine and **Power Off**.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Select **Actions** > **Edit Resources**.

STEP 4 | Set the **CPU** and **Memory**.

STEP 5 | **Save** your changes.

STEP 6 | Select the Panorama virtual machine and **Power On**.

Increase CPUs and Memory for Panorama on Alibaba Cloud

You can change the instance type of the Panorama™ virtual appliance to increase the CPUs and memory allocated to the Panorama virtual appliance instance. Be sure to review the supported Alibaba Cloud instance types and the Setup Prerequisites for the Panorama Virtual Appliance before changing the instance type.

STEP 1 | Log in to the Alibaba Cloud Console.

STEP 2 | Select **Elastic Compute Service** > **Instances & Images** > **Instances** and navigate to the Panorama virtual appliance instance.

STEP 3 | In the Actions column, select **More** > **Instance Status** > **Stop.**

STEP 4 | Change the Panorama virtual appliance instance type.

1. Select the Panorama virtual appliance if not already selected.
2. In the Actions column, select **Change Instance Type**.
3. Select the desired instance type and **Change** the instance type.
4. When prompted, select **Console** to view your Panorama virtual appliance instance.

STEP 5 | In the Actions column for the Panorama virtual appliance instance, select **More** > **Instance Status** > **Start.**

STEP 6 | Verify the increased CPU and memory.

1. Log in to the Panorama CLI.
2. View the Panorma virtual appliance system information.

```
admin> show system info
```

3. Verify that the `num-cpus` and `ram-in-gb` display the correct number of CPUs and amount of memory as per the instance type you selected.

Increase CPUs and Memory for Panorama on AWS

For the minimum CPUs and memory that Panorama™ requires, see Increase CPUs and Memory on the Panorama Virtual Appliance.

STEP 1 | Log in to AWS Web Service console and select the EC2 Dashboard.

- Amazon Web Service Console
- AWS GovCloud Web Service Console

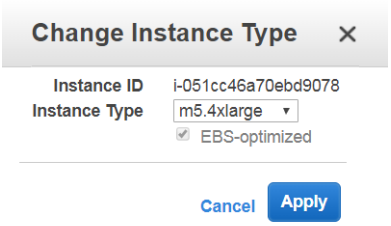
STEP 2 | On the EC2 Dashboard, select **Instances** and select the Panorama virtual appliance instance.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Select **Actions** > **Instance State** > **Stop** to power off the Panorama virtual appliance instance.

STEP 4 | Select **Actions** > **Instance Settings** > **Change Instance Type** to change the Panorama virtual appliance instance type.

STEP 5 | Select the **Instance Type** to which you want to upgrade and **Apply** it.



STEP 6 | Select **Actions** > **Instance State** > **Start** to power on the Panorama virtual appliance instance.

Increase CPUs and Memory for Panorama on Azure

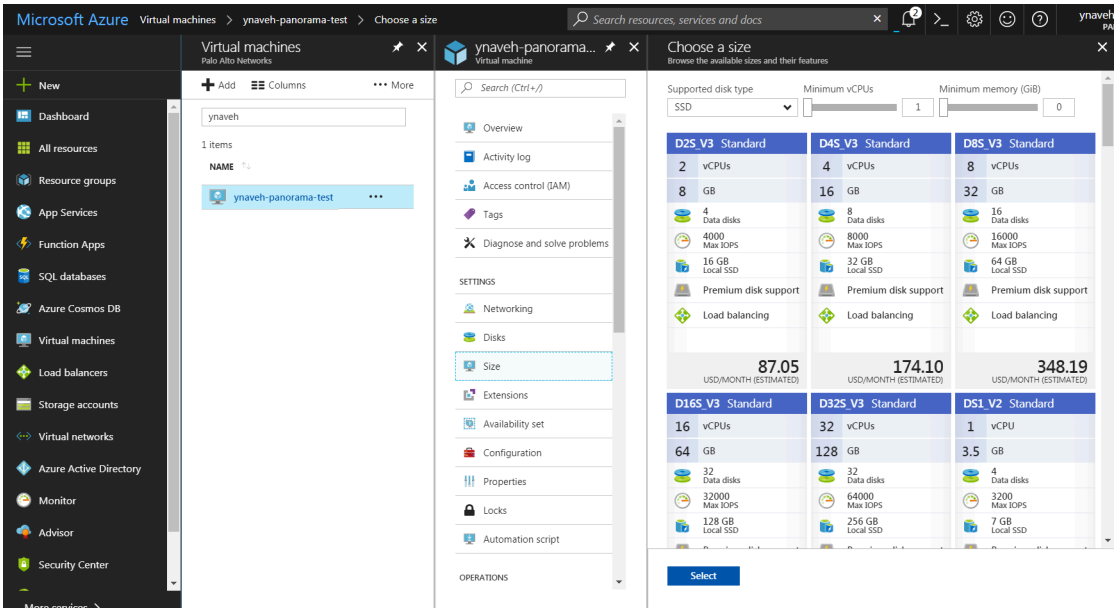
For the minimum CPUs and memory that Panorama™ requires, see Increase CPUs and Memory on the Panorama Virtual Appliance.

STEP 1 | Log in to the Microsoft Azure portal.

STEP 2 | On the Azure Dashboard, under **Virtual machines**, select the Panorama virtual appliance.

STEP 3 | Select **Overview** and **Stop** the Panorama virtual appliance.

STEP 4 | Choose the new virtual machine **Size** and then **Select** it.



STEP 5 | Select **Overview** and **Start** the Panorama virtual appliance.

 ©2025 Palo Alto Networks, Inc.

Increase CPUs and Memory for Panorama on Google Cloud Platform

For the minimum CPUs and memory that Panorama™ requires, see Increase CPUs and Memory on the Panorama Virtual Appliance.

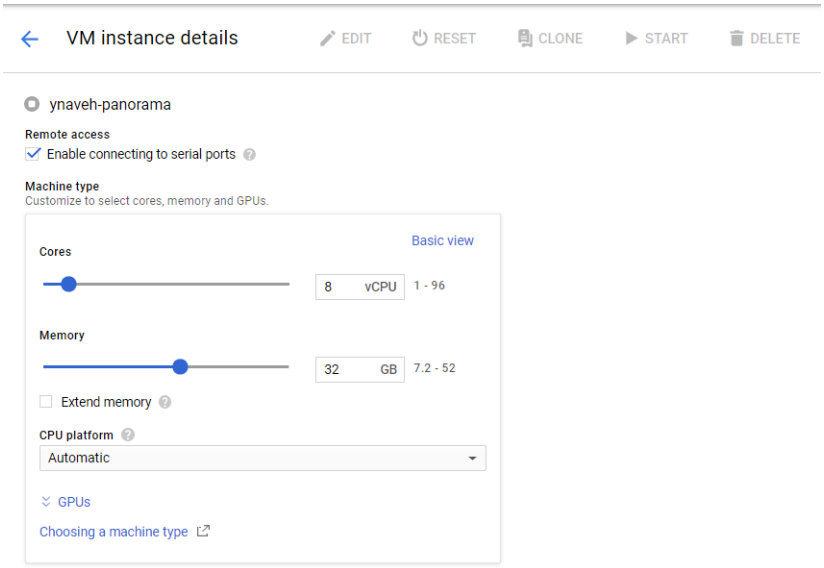
STEP 1 | Log in to the Google Cloud Console.

STEP 2 | Stop the Panorama virtual appliance instance.

1. Select the Panorama virtual appliance instance in the Products & Services menu (**Compute Engine** > **VM Instances**).
2. **Stop** the Panorama virtual appliance instance. It can take 2 to 3 minutes for the Panorama virtual appliance to completely shut down.

STEP 3 | Reconfigure the Panorama virtual appliance resources.

1. **Edit** the Panorama virtual appliance instance details.
2. Under Machine Type, **Customize** the Panorama virtual appliance CPU cores and memory.



STEP 4 | **Save** the changes to update the Panorama virtual appliance instance.

STEP 5 | **Start** the Panorama virtual appliance.

Increase CPUs and Memory for Panorama on KVM

For the minimum CPUs and memory that Panorama™ requires, see Increase CPUs and Memory on the Panorama Virtual Appliance.

STEP 1 | **Shutdown** the Panorama virtual appliance instance on the Virtual Machine Manager.

STEP 2 | Double-click the Panorama virtual appliance instance in the Virtual Machine Manager and **Show virtual hardware details**.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Edit the allocated Panorama virtual appliance CPU cores.

1. Edit the currently allocated **CPUs**.
2. **Apply** the reconfigured CPU core allocation.

STEP 4 | Edit the allocated Panorama virtual appliance memory.

1. Edit the currently allocated **Memory**.
2. **Apply** the reconfigured memory allocation.

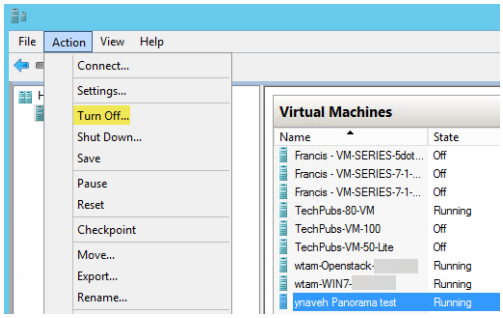
STEP 5 | **Power on** the Panorama virtual appliance instance.

Increase CPUs and Memory for Panorama on Hyper-V

For the minimum CPUs and memory that Panorama™ requires, see Increase CPUs and Memory on the Panorama Virtual Appliance.

STEP 1 | Power off the Panorama virtual appliance.

1. On the Hyper-V Manager, select the Panorama virtual appliance instance from the list of **Virtual Machines**.
2. Select **Action** > **Turn Off** to power off the Panorama virtual appliance.

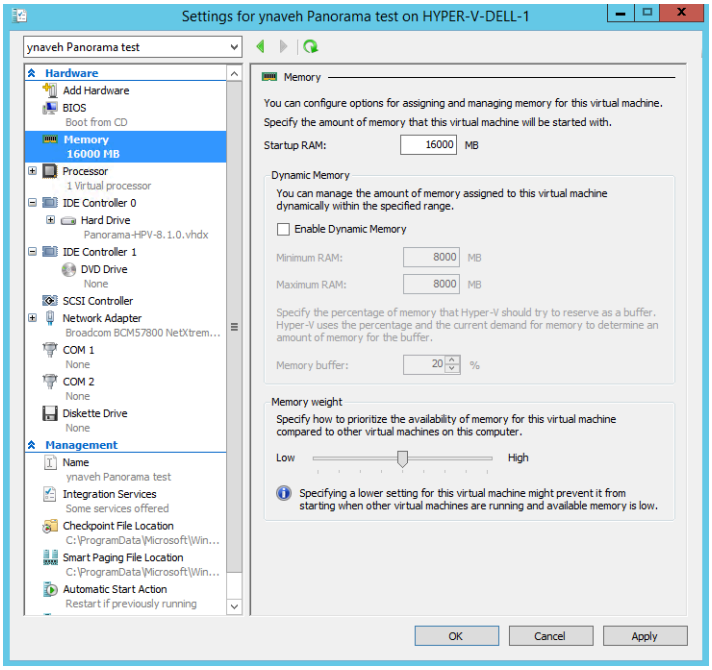


STEP 2 | On the Hyper-V Manager, select the Panorama virtual appliance instance form the list of **Virtual Machines**, and select **Action** > **Settings** to edit the Panorama virtual appliance resources.

 ©2025 Palo Alto Networks, Inc.

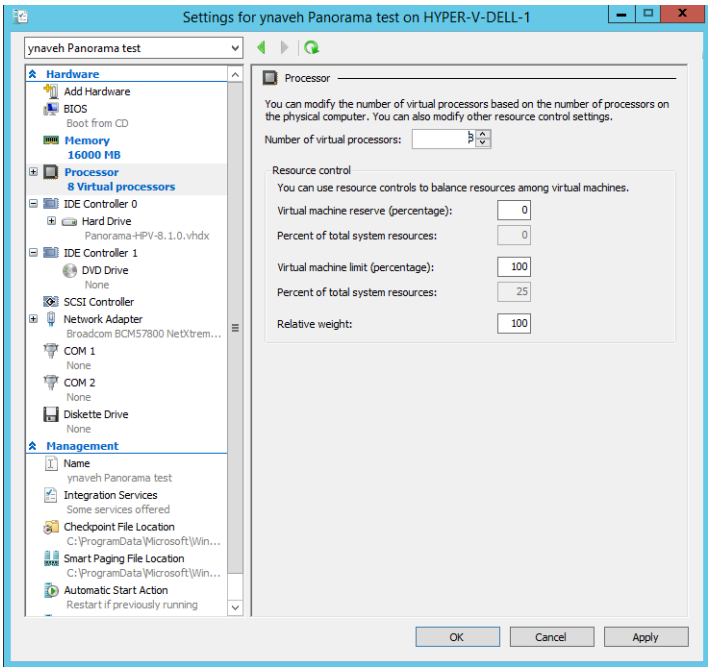
STEP 3 | Edit the allocated Panorama virtual appliance memory.

1. In the **Hardware** list, select **Memory**.
2. Edit the currently allocated **Startup RAM**.



STEP 4 | Edit the allocated Panorama virtual appliance CPU cores.

1. In the **Hardware** list, select **Processor**.
2. Edit the currently allocated **Number of virtual processors**.

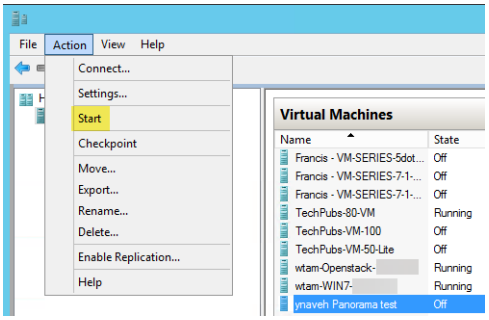


STEP 5 | **Apply** the reallocated memory and CPU cores.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Power on the Panorama virtual appliance.

1. Select the Panorama virtual appliance instance from the list of **Virtual Machines**.
2. Select **Action** > **Start** to power on the Panorama virtual appliance.



Increase the CPUs and Memory for Panorama on Oracle Cloud Infrastructure (OCI)

You can change the instance type of the Panorama™ virtual appliance to increase the CPUs and memory allocated to the Panorama virtual appliance instance. Be sure to review the Setup Prerequisites for the Panorama Virtual Appliance before modifying the Panorama virtual appliance instance CPUs and memory.

STEP 1 | Log in to the Oracle Cloud Infrastructure console.

STEP 2 | Power off the Panorama virtual appliance instance.

1. Select **Compute** > **Instances** and click the name of the Panorama virtual appliance instance.
2. **Stop** the Panorama virtual appliance instance.

STEP 3 | Increase the CPUs and memory.

1. In the instance details, select **Edit** > **Edit Shape**.
2. Increase the number of CPUs and memory allocated to the instance.
3. **Save Changes**.

STEP 4 | In instance details, **Start** the Panorama virtual appliance.

STEP 5 | Verify the increased CPU and memory.

1. Log in to the Panorama CLI.
2. View the Panorma virtual appliance system information.

```
admin> show system info
```

3. Verify that the `num-cpus` and `ram-in-gb` display the correct number of CPUs and amount of memory as per the instance type you selected.

Increase the System Disk on the Panorama Virtual Appliance

Expand the system disk capacity to 224GB for the Panorama virtual appliance to support large datasets to allow for sufficient disk space for things such as dynamic updates when you Manage Large-Scale Firewall Deployments. Additionally, a 224GB system disk expands storage for

 ©2025 Palo Alto Networks, Inc.

monitoring and reporting data for managed firewall health if you intended to use the Panorama virtual appliance in Panorama mode to manage your SD-WAN deployment

- Increase the System Disk for Panorama on an ESXi Server
- Increase the System Disk for Panorama on Google Cloud Platform

Increase the System Disk for Panorama on an ESXi Server

Add a 224GB system disk to replace the default 81GB system disk. For the minimum resource requirements for the Panorama virtual appliance, see Setup Prerequisites for the Panorama Virtual Appliance.

Decreasing the Panorama virtual appliance system disk back to 81GB is not supported.

STEP 1 | (Best Practice) Save and Export Panorama and Firewall Configurations.

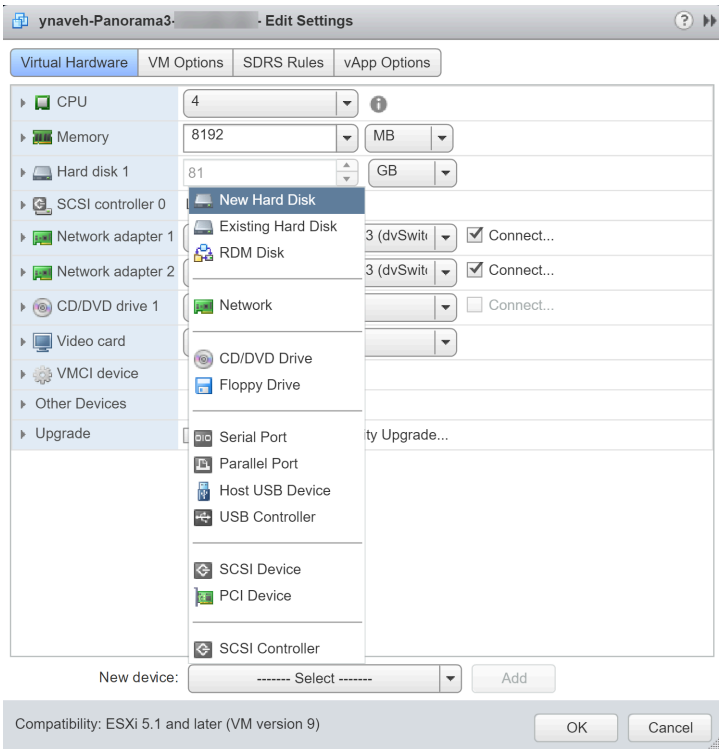
Save and export your Panorama and firewall configuration to ensure you can recover Panorama if you encounter any issues.

STEP 2 | Access the VMware vSphere Client and navigate to your Panorama virtual appliance.

STEP 3 | Right-click the Panorama virtual appliance and select **Power** > **Power Off**.

STEP 4 | Add the new 224GB system disk.

1. Right-click the Panorama virtual appliance and **Edit Settings**.
2. Select **New Hard Disk** as the **New Device** and **Add** the new device.
3. Configure the new hard disk with 224GB and click **OK**.



 ©2025 Palo Alto Networks, Inc.

STEP 5 | Right-click the Panorama virtual appliance and select **Power** > **Power On**.

Panorama may take up to 30 minutes to initialize the new system disk. During this time the Panorama web interface and CLI are unavailable.

STEP 6 | Migrate disk data from the old system disk to the new system disk.

In this example, we are migrating to the newly added system disk labeled sdb.

1. Log in to the Panorama CLI.
2. Enter the following command to view the available system disks for migration:

```
admin> request system clone-system-disk target ?
```

3. Migrate the disk data to the new system disk using the following command:

```
admin> request system clone-system-disk target sdb
```

Enter **Y** when prompted to begin the disk migration.

To begin the migration, Panorama reboots and takes at least 20 minutes to complete the disk migration. During this time the Panorama web interface and CLI are unavailable.

4. Monitor the disk migration from the web Console. Continue to the next step only after Panorama displays the following message to indicate the disk migration is complete.

```
==================================
Disk Cloning Utility (Version 1.0)
==================================
SOURCE - Disk sda (82944 MB)
TARGET - Disk sdb (229376 MB)

Gathering disks info
Finished gathering disks info

Preparing disks
Finished preparing disks

Copying data
Finished copying data

Making disk bootable
Finished making disk bootable

Disk cloning procedure completed. Please shutdown the sytem and switch disks..._
```

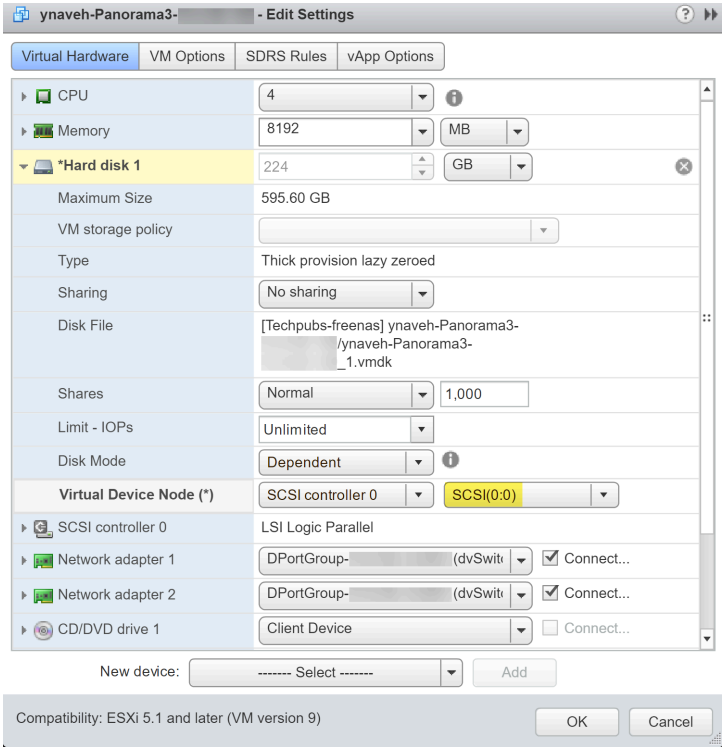
STEP 7 | Delete the old system disk.

1. Access the VMware vSphere Client and navigate to your Panorama virtual appliance.
2. Right-click the Panorama virtual appliance and select **Power** > **Power Off**.
3. Right-click the Panorama virtual appliance and **Edit Settings**.
4. Delete the old 81GB system disk and click **OK**.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Modify the Virtual Device Node for the new system disk.

1. Expand the settings options for the new system disk.
2. Select **SCSI(0:0)** as the **Virtual Device Node**.
3. Click **OK** to save your configuration changes.



STEP 9 | Right-click the Panorama virtual appliance and select **Power** > **Power On**.

STEP 10 | Verify that you successfully migrated to the new system disk.

1. Log in to the Panorama CLI.
2. Enter the following command to view the system disk partitions.

   You must examine the `/dev/root`, `/dev/sda5`, `/dev/sda6`, and `/dev/sda8` partitions to confirm the disk size is increased.

```
admin> show system disk-space
```

```
admin@Panorama-Ynaveh> show system disk-space

Filesystem      Size  Used Avail Use% Mounted on
/dev/root        16G  3.4G   12G  23% /
none            4.0G   60K  4.0G   1% /dev
/dev/sda5        76G  1.8G   71G   3% /opt/pancfg
/dev/sda6        23G  5.0G   17G  24% /opt/panrepo
tmpfs           4.0G  110M  3.8G   3% /dev/shm
cgroup_root     4.0G     0  4.0G   0% /cgroup
/dev/sda8        92G   52G   35G  60% /opt/panlogs
/dev/loop0       50G  7.4G   40G  16% /opt/mongobuffer
tmpfs            12M     0   12M   0% /opt/pancfg/mgmt/ssl/private
```

 ©2025 Palo Alto Networks, Inc.

Increase the System Disk for Panorama on Google Cloud Platform

Add a 224GB system disk to replace the default 81GB system disk. For the minimum resource requirements for the Panorama virtual appliance, see Setup Prerequisites for the Panorama Virtual Appliance.

STEP 1 | (Best Practice) Save and Export Panorama and Firewall Configurations.

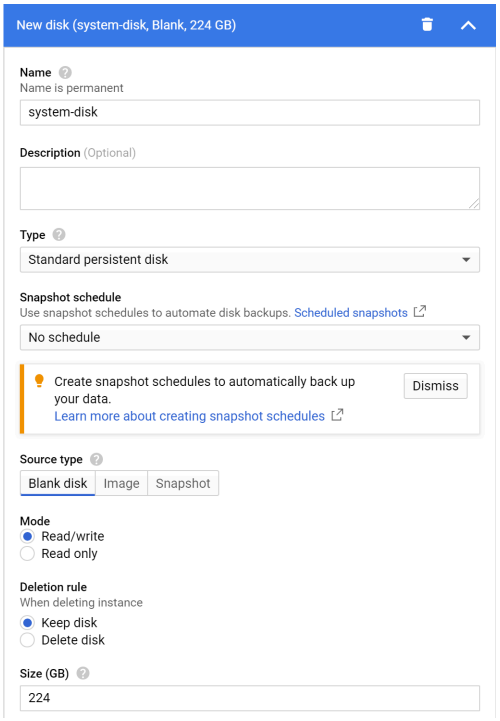
Save and export your Panorama and firewall configuration to ensure you can recover Panorama if you encounter any issues.

STEP 2 | Log in to the Google Cloud Console.

STEP 3 | In **VM Instances**, **Stop** the Panorama VM instance.

STEP 4 | Add the new 224GB system disk.

1. Select the Panorama VM instance and select **Edit**.
2. In the **Additional disks** section **Add new disk**.
3. Configure the new disk with 224GB and click **OK**.



STEP 5 | In **VM Instances**, **Start** the Panorama VM instance.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Migrate disk data from the old system disk to the new system disk.

In this example, we are migrating to the newly added system disk labeled sdb.

1. Log in to the Panorama CLI.
2. Enter the following command to view the available system disks for migration:

```
admin> request system clone-system-disk target ?
```

3. Migrate the disk data to the new system disk using the following command:

```
admin> request system clone-system-disk target sdb
```

Enter **Y** when prompted to begin the disk migration.

To begin the migration, Panorama reboots and takes at least 20 minutes to complete the disk migration. During this time the Panorama web interface and CLI are unavailable.

4. Monitor the disk migration by attempting to log in to the Panorama CLI. The Panorama management server is in maintenance mode after the system disk migration is completed and will allow you to log in to the Panorama CLI while in maintenance mode.

STEP 7 | Attach the new 224GB system disk.

1. In **VM Instances**, **Stop** the Panorama VM instance.
2. Select the Panorama VM instance and select **Edit**.
3. In the **Additional disks** section, detach the new 224GB system disk.
4. In the **Boot Disk** section, detach the old 81GB system disk.
5. In the **Boot Disk** section, **Add item** and select the new 224GB system disk.
6. **Save** your configuration changes.

STEP 8 | In **VM Instances**, **Start** the Panorama VM instance.

 ©2025 Palo Alto Networks, Inc.

STEP 9 | Verify that you successfully migrated to the new system disk.

1. Log in to the Panorama CLI.
2. Enter the following command to view the system disk partitions.

   You must examine the `/dev/root`, `/dev/sda5`, `/dev/sda6`, and `/dev/sda8` partitions to confirm the disk size is increased.

```
admin> show system disk-space
```

```
admin@Panorama-Ynaveh> show system disk-space

Filesystem      Size  Used Avail Use% Mounted on
/dev/root        16G  3.4G   12G  23% /
none            4.0G   60K  4.0G   1% /dev
/dev/sda5        76G  1.8G   71G   3% /opt/pancfg
/dev/sda6        23G  5.0G   17G  24% /opt/panrepo
tmpfs           4.0G  110M  3.8G   3% /dev/shm
cgroup_root     4.0G     0  4.0G   0% /cgroup
/dev/sda8        92G   52G   35G  60% /opt/panlogs
/dev/loop0       50G  7.4G   40G  16% /opt/mongobuffer
tmpfs            12M     0   12M   0% /opt/pancfg/mgmt/ssl/private
```

Complete the Panorama Virtual Appliance Setup

After you Perform Initial Configuration of the Panorama Virtual Appliance, continue with the following tasks for additional configuration:

- Activate a Panorama Support License
- Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected
- Install Content and Software Updates for Panorama
- Access and Navigate Panorama Management Interfaces
- Set Up Administrative Access to Panorama
- Manage Firewalls

Convert Your Panorama Virtual Appliance

You can convert your evaluation Panorama™ virtual appliance to a production Panorama virtual appliance to preserve its existing configuration and begin leveraging the management platform.

If you are utilizing Enterprise License Agreement (ELA) licensing, you can convert an existing production Panorama virtual appliance to leverage the benefits of ELA licensing.

- Convert Your Evaluation Panorama to a Production Panorama with Local Log Collector
- Convert Your Evaluation Panorama to a Production Panorama without Local Log Collector
- Convert Your Evaluation Panorama to VM-Flex Licensing with Local Log Collector
- Convert Your Evaluation Panorama to VM-Flex Licensing without Local Log Collector
- Convert Your Production Panorama to an ELA Panorama

 ©2025 Palo Alto Networks, Inc.

Convert Your Evaluation Panorama to a Production Panorama with Local Log Collector

If you have an evaluation Panorama™ virtual appliance in Panorama mode configured with a local Log Collector, you can convert it to a production Panorama by migrating the configuration from the evaluation Panorama to the production Panorama and modifying as needed.

Logs ingested by the Log Collector on a Panorama virtual appliance cannot be migrated.

If you need to maintain access to the logs stored on your evaluation Panorama virtual appliance, after you migrate the evaluation Panorama configuration to the production Panorama, keep your evaluation Panorama powered on to access the logs locally for the remainder of the evaluation license lifetime. Adding the evaluation Panorama to the production Panorama as a managed collector is not supported.

STEP 1 | Plan the migration.

- ☐ Upgrade the software on the Panorama virtual appliance before you convert your evaluation Panorama virtual appliance to a production Panorama virtual appliance. Review the Compatibility Matrix for the minimum PAN-OS version required for your hypervisor. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
- ☐ Schedule a maintenance window for the migration.

STEP 2 | Set up your production Panorama virtual appliance.

1. Set Up the Panorama Virtual Appliance.
2. Register the Panorama virtual appliance with the Palo Alto Networks Customer Support Portal (CSP).

   The Panorama serial number and authorization code are found in the Order Summary email from Palo Alto Networks.
3. Install Content and Software Updates for Panorama.

STEP 3 | Activate the device management license on the Palo Alto Networks Custer Support Portal (CSP) for the production Panorama virtual appliance.

1. Log in to the Palo Alto Networks CSP.
2. Select **Assets** > **Devices** and locate your Panorama virtual appliance.
3. In the **Action** column, click the pencil icon to edit the device licenses.
4. Select **Activate Auth-Code** and enter the **Authorization Code**.
5. Select **Agree and Submit** to activate the device management license.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Export the Panorama configuration from the evaluation Panorama virtual appliance.

1. Log in to the Panorama Web Interface.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Export named Panorama configuration snapshot**, select `running-config.xml` and click **OK**. Panorama exports the configuration to your client system as an XML file.
4. Locate the `running-config.xml` file you exported and rename the XML file. This is required to import the configuration as Panorama does not support importing an XML file with the name `running-config.xml`.

STEP 5 | Load the Panorama configuration snapshot that you exported from the evaluation Panorama virtual appliance into the production Panorama virtual appliance.

1. Log in to the Panorama Web Interface of the production Panorama virtual appliance.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Import named Panorama configuration snapshot**, **Browse** to the Panorama configuration file you exported from the Panorama virtual appliance, and click **OK**.
4. Click **Load named Panorama configuration snapshot**, select the **Name** of the configuration you just imported, leave the **Decryption Key** blank (empty), and click **OK**. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file.
5. If errors occurred, save them to a local file. Resolve each error to ensure the migrated configuration is valid.

STEP 6 | Modify the configuration on the production Panorama virtual appliance.

1. Select **Panorama** > **Setup** > **Management**.
2. Edit the General Settings, modify the **Hostname**, and click **OK**.
3. Edit the Management Interface Settings to configure the management IP address and click **OK**.

   The most efficient approach is to assign a new IP address to the evaluation Panorama virtual appliance and reuse its old IP address for the production Panorama virtual appliance. This ensures that the evaluation Panorama virtual appliance remains accessible and that firewalls can point to the production Panorama virtual appliance without you reconfiguring the Panorama IP address on each firewall.
4. Remove the Log Collector configuration imported from the evaluation Panorama.
   1. Select **Panorama** > **Collector Group** and **Delete** all configured collector groups.
   2. Select **Panorama** > **Managed Collectors** and **Delete** all configured Log Collectors.
5. Select **Commit** > **Commit to Panorama** and **Commit** your changes to the Panorama configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Configure your Log Collectors and collector groups.

You must add the managed collectors, collector group configuration, and log forwarding configurations you deleted in the previous step, as well as add the local Log Collector.

1. Configure a Managed Collector.
2. Configure a Collector Group.
3. Configure Log Forwarding to Panorama.

STEP 8 | Verify that the support and device management licenses are successfully activated.

1. Select **Panorama** > **Licenses** and **Retrieve license keys from license server**.
2. Verify the **Device Management License** displays the correct number of devices.
3. Select **Panorama** > **Support** and verify that the correct support **Level** and **Expiry Date** are displayed.

STEP 9 | Synchronize the production Panorama virtual appliance with the firewalls to resume firewall management.

Complete this step during a maintenance window to minimize network disruption.

1. On the production Panorama virtual appliance, select **Panorama** > **Managed Devices** and verify that the Device State column displays **Connected** for the firewalls.

   At this point, the Shared Policy (device groups) and Template columns display **Out of sync** for the firewalls.
2. Push your changes to device groups and templates:
   1. Select **Commit** > **Push to Devices** and **Edit Selections**.
   2. Select **Device Groups**, select every device group, **Include Device and Network Templates**, and click **OK**.
   3. **Push** your changes.
3. In the **Panorama** > **Managed Devices** page, verify that the Shared Policy and Template columns display **In sync** for the firewalls.

Convert Your Evaluation Panorama to a Production Panorama without Local Log Collector

Change the serial number of your evaluation Panorama virtual appliance in Management Only mode or in Panorama mode with no local Log Collector configured to convert it to a production Panorama virtual appliance.

If a local Log Collector is configured, see Convert Your Evaluation Panorama to a Production Panorama with Local Log Collector.

STEP 1 | Log in to the Panorama web interface.

STEP 2 | Select **Panorama** > **Setup** > **Management** and edit the General Settings.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Enter the **Serial Number** provided by Palo Alto Networks.

The Panorama serial number and authorization code are obtained from the deployment profile you created in the previous step.

STEP 4 | Click **OK**.

STEP 5 | Select **Commit** and **Commit to Panorama**.

STEP 6 | Restart management server on the Panorama virtual appliance.

1. Log in to the Panorama CLI.
2. Restart the management server.

```
admin> debug software restart process management-server
```

All administrators are logged out of the Panorama web interface and CLI when you restart the management server.

STEP 7 | Verify that the support and device management licenses are successfully activated.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Licenses** and **Retrieve license keys from license server**.
3. Verify the **Device Management License** displays the correct number of devices.
4. Select **Panorama** > **Support** and verify that the correct support **Level** and **Expiry Date** are displayed.

STEP 8 | Synchronize the production Panorama virtual appliance with the firewalls to resume firewall management.

Complete this step during a maintenance window to minimize network disruption.

1. On the production Panorama virtual appliance, select **Panorama** > **Managed Devices** and verify that the Device State column displays **Connected** for the firewalls.

   At this point, the Shared Policy (device groups) and Template columns display **Out of sync** for the firewalls.
2. Push your changes to device groups and templates:
   1. Select **Commit** > **Push to Devices** and **Edit Selections**.
   2. Select **Device Groups**, select every device group, **Include Device and Network Templates**, and click **OK**.
   3. **Push** your changes.
3. In the **Panorama** > **Managed Devices** page, verify that the Shared Policy and Template columns display **In sync** for the firewalls.

Convert Your Evaluation Panorama to VM-Flex Licensing with Local Log Collector

If you have an evaluation Panorama™ virtual appliance in Panorama mode configured with a local Log Collector, you can convert it to a production Panorama with VM Flex licensing by migrating

 ©2025 Palo Alto Networks, Inc.

the configuration from the evaluation Panorama to the production Panorama and modifying as needed.

If a local Log Collector is not configured, see Convert Your Evaluation Panorama to VM-Flex Licensing without Local Log Collector.

Logs ingested by the Log Collector on a Panorama virtual appliance cannot be migrated.

If you need to maintain access to the logs stored on your evaluation Panorama virtual appliance, after you migrate the evaluation Panorama configuration to the production Panorama, keep your evaluation Panorama powered on to access the logs locally for the remainder of the evaluation license lifetime. Adding the evaluation Panorama to the production Panorama as a managed collector is not supported.

STEP 1 | Plan the migration.

- ❑ Upgrade the software on the Panorama virtual appliance before you convert your evaluation Panorama virtual appliance to a production Panorama virtual appliance. Review the Compatibility Matrix for the minimum PAN-OS version required for your hypervisor. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
- ❑ Schedule a maintenance window for the migration.

STEP 2 | Obtain the Panorama serial number and auth code from your flexible VM-Series licensing deployment profile.

1. Log in to the Palo Alto Networks Customer Support Portal (CSP).
2. Create a deployment profile that enables a Panorama virtual appliance.
3. Provision Panorama to generate the a serial number for Panorama.
4. Copy the **Serial Number** and **Auth Code**.

STEP 3 | Set up your production Panorama virtual appliance.

1. Log in to the Palo Alto Networks CSP.
2. Set Up the Panorama Virtual Appliance.
3. Register the Panorama virtual appliance with the Palo Alto Networks Customer Support Portal (CSP).

   The Panorama serial number and authorization code you generated in the previous step.
4. Install Content and Software Updates for Panorama.

STEP 4 | Activate the device management license on the Palo Alto Networks CSP for the production Panorama virtual appliance.

1. Select **Assets** > **Devices** and locate your Panorama virtual appliance.
2. In the **Action** column, click the pencil icon to edit the device licenses.
3. Select **Activate Auth-Code** and enter the **Authorization Code**.
4. Select **Agree and Submit** to activate the device management license.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Export the Panorama configuration from the evaluation Panorama virtual appliance.

1. Log in to the Panorama Web Interface.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Export named Panorama configuration snapshot**, select `running-config.xml` and click **OK**. Panorama exports the configuration to your client system as an XML file.
4. Locate the `running-config.xml` file you exported and rename the XML file. This is required to import the configuration as Panorama does not support importing an XML file with the name `running-config.xml`.

STEP 6 | Load the Panorama configuration snapshot that you exported from the evaluation Panorama virtual appliance into the production Panorama virtual appliance.

1. Log in to the Panorama Web Interface of the production Panorama virtual appliance.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Import named Panorama configuration snapshot**, **Browse** to the Panorama configuration file you exported from the Panorama virtual appliance, and click **OK**.
4. Click **Load named Panorama configuration snapshot**, select the **Name** of the configuration you just imported, leave the **Decryption Key** blank (empty), and click **OK**. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file.
5. If errors occurred, save them to a local file. Resolve each error to ensure the migrated configuration is valid.

STEP 7 | Modify the configuration on the production Panorama virtual appliance.

1. Select **Panorama** > **Setup** > **Management**.
2. Edit the General Settings, modify the **Hostname**, and click **OK**.
3. Edit the Management Interface Settings to configure the management IP address and click **OK**.

   The most efficient approach is to assign a new IP address to the evaluation Panorama virtual appliance and reuse its old IP address for the production Panorama virtual appliance. This ensures that the evaluation Panorama virtual appliance remains accessible and that firewalls can point to the production Panorama virtual appliance without you reconfiguring the Panorama IP address on each firewall.
4. Remove the Log Collector configuration imported from the evaluation Panorama.
   1. Select **Panorama** > **Collector Group** and **Delete** all configured collector groups.
   2. Select **Panorama** > **Managed Collectors** and **Delete** all configured Log Collectors.
5. Select **Commit** > **Commit to Panorama** and **Commit** your changes to the Panorama configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Reconfigure your Log Collectors and collector groups.

You must add the managed collectors, collector group configuration, and log forwarding configurations you deleted in the previous step, as well as add the local Log Collector.

1. Configure a Managed Collector.
2. Configure a Collector Group.
3. Configure Log Forwarding to Panorama.

STEP 9 | Verify that the support and device management licenses are successfully activated.

1. Select **Panorama** > **Licenses** and **Retrieve license keys from license server**.
2. Verify the **Device Management License** displays the correct number of devices.
3. Select **Panorama** > **Support** and verify that the correct support **Level** and **Expiry Date** are displayed.

STEP 10 | Synchronize the production Panorama virtual appliance with the firewalls to resume firewall management.

Complete this step during a maintenance window to minimize network disruption.

1. On the production Panorama virtual appliance, select **Panorama** > **Managed Devices** and verify that the Device State column displays **Connected** for the firewalls.

   At this point, the Shared Policy (device groups) and Template columns display **Out of sync** for the firewalls.
2. Push your changes to device groups and templates:
   1. Select **Commit** > **Push to Devices** and **Edit Selections**.
   2. Select **Device Groups**, select every device group, **Include Device and Network Templates**, and click **OK**.
   3. **Push** your changes.
3. In the **Panorama** > **Managed Devices** page, verify that the Shared Policy and Template columns display **In sync** for the firewalls.

Convert Your Evaluation Panorama to VM-Flex Licensing without Local Log Collector

Change the serial number of your evaluation Panorama virtual appliance in Management Only mode or in Panorama mode with no local Log Collector configured to convert it to a production Panorama virtual appliance.

If a local Log Collector is configured, see Convert Your Evaluation Panorama to VM-Flex Licensing with Local Log Collector.

STEP 1 | Obtain the Panorama serial number and auth code from your flexible VM-Series licensing deployment profile.

1. Log in to the Palo Alto Networks Customer Support Portal (CSP).
2. Create a deployment profile that enables a Panorama virtual appliance.
3. Provision Panorama to generate the a serial number for Panorama.
4. Copy the **Serial Number** and **Auth Code**.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Log in to the Panorama web interface.

STEP 3 | Select **Panorama** > **Setup** > **Management** and edit the General Settings.

STEP 4 | Enter the **Serial Number** provided by Palo Alto Networks.

The Panorama serial number and authorization code you generated in the previous step.

STEP 5 | Click **OK**.

STEP 6 | Select **Commit** and **Commit to Panorama**.

STEP 7 | Restart management server on the Panorama virtual appliance.

1. Log in to the Panorama CLI.
2. Restart the management server.

```
admin> debug software restart process management-server
```

All administrators are logged out of the Panorama web interface and CLI when you restart the management server.

STEP 8 | Verify that the support and device management licenses are successfully activated.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Licenses** and **Retrieve license keys from license server**.
3. Verify the **Device Management License** displays the correct number of devices.
4. Select **Panorama** > **Support** and verify that the correct support **Level** and **Expiry Date** are displayed.

STEP 9 | Synchronize the production Panorama virtual appliance with the firewalls to resume firewall management.

Complete this step during a maintenance window to minimize network disruption.

1. On the production Panorama virtual appliance, select **Panorama** > **Managed Devices** and verify that the Device State column displays **Connected** for the firewalls.

   At this point, the Shared Policy (device groups) and Template columns display **Out of sync** for the firewalls.
2. Push your changes to device groups and templates:
   1. Select **Commit** > **Push to Devices** and **Edit Selections**.
   2. Select **Device Groups**, select every device group, **Include Device and Network Templates**, and click **OK**.
   3. **Push** your changes.
3. In the **Panorama** > **Managed Devices** page, verify that the Shared Policy and Template columns display **In sync** for the firewalls.

 ©2025 Palo Alto Networks, Inc.

Convert Your Production Panorama to an ELA Panorama

You can convert your production Panorama™ virtual appliance to continue leveraging your Panorama with the benefits of ELA licensing. To convert your production deployment, Panorama must have out-bound Internet access.

Converting your production Panorama to ELA licensing is supported in Management Only and Panorama mode with or without a local Log Collector configured. If your Panorama has a local Log Collector configured, you must submit a support ticket with Palo Alto Networks to convert your Panorama to ELA licensing.

During conversion from a production Panorama to ELA licensing, do not change the Panorama serial number if a local Log Collector is configured.

The log on the local Log collector become inaccessible and other Log Collectors in the Collector Group may become inaccessible and no longer ingest logs if the serial number of a Log Collector is changed.

STEP 1 | Covert your Panorama to ELA licensing.

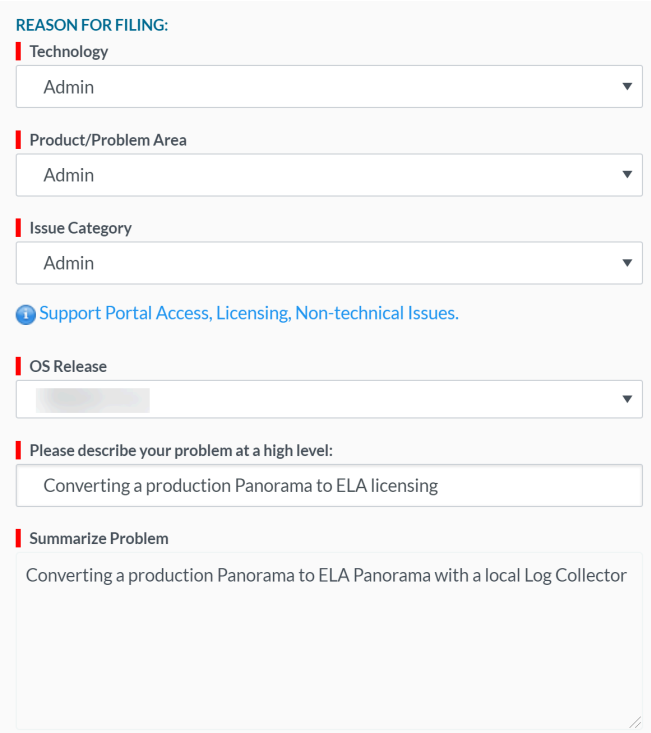
- **Panorama virtual appliance in Panorama mode with a local Log Collector.**

  Submit support ticket with Palo Alto Networks to convert your Panorama to ELA licensing. This is required in order to preserve all existing logs on the local Log Collector when converting a Panorama with a local Log Collector to ELA licensing. An example is provided

 ©2025 Palo Alto Networks, Inc.

below to assist in filing the support ticket. Create the ticket exactly as displayed below, and select the **OS Release** your Panorama is running.

Continue to the next step only after Palo Alto Networks support successfully resolves your support ticket.



- **Panorama virtual appliance in Management Only mode or Panorama mode with no local Log Collector.**
  1. Generate a serial number from your ELA licensing pool.
     1. Log in to the Palo Alto Networks CSP.
     2. Select **Assets** > **VM-Series Auth-Codes** and locate your ELA licensing pool.
     3. In the Actions column, select **Panorama** and **Provision** a new serial number.

        Confirm the new serial number provision when prompted.
     4. Copy the newly provisioned serial number.
  2. Log in to the Panorama web interface.
  3. Select **Panorama** > **Setup** > **Management** and edit the General Settings.
  4. Enter the **Serial Number** you provisioned.
  5. Click **OK**.
  6. Select **Commit** and **Commit to Panorama**.

STEP 2 | Log in to the Panorama web interface if not already logged in.

STEP 3 | Select **Panorama** > **Licenses** and **Retrieve new licenses from the license server**.

STEP 4 | Verify that Panorama retrieved the new licenses as per your ELA agreement.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Verify that the support and device management licenses are successfully activated.

1. Select **Panorama** > **Licenses** and verify that the correct licenses are activated.
2. Select **Panorama** > **Support** and verify that the correct support **Level** and **Expiry Date** are displayed.

 ©2025 Palo Alto Networks, Inc.

Set Up the M-Series Appliance

The M-700, M-600, M-500, M-300, and M-200 appliances are high performance hardware appliances that you can deploy in Management Only mode (as Panorama management servers with no local log collection), Panorama mode (as Panorama management servers with local log collection) or in Log Collector mode (as Dedicated Log Collectors). The appliances provide multiple interfaces that you can assign to various Panorama services such as firewall management and log collection. Before setting up the appliance, consider how you can configure the interfaces to optimize security, enable network segmentation (in large-scale deployments), and load balance the traffic for Panorama services.

- M-Series Appliance Interfaces
- Perform Initial Configuration of the M-Series Appliance
- Perform Initial Configuration of an Air Gapped M-Series Appliance
- M-Series Setup Overview
- Set Up the M-Series Appliance as a Log Collector
- Increase Storage on the M-Series Appliance
- Configure Panorama to Use Multiple Interfaces

M-Series Appliance Interfaces

The Panorama M-700, M-600, M-300, and M-200 appliances have several interfaces for communicating with other systems such as managed firewalls and the client systems of Panorama administrators. Panorama communicates with these systems to perform various services, including managing devices (firewalls, Log Collectors, and WildFire appliances and appliance clusters), collecting logs, communicating with Collector Groups, deploying software and content updates to devices, and providing administrative access to Panorama. By default, Panorama uses its management (MGT) interface for all these services. However, you can improve security by reserving the MGT interface for administrative access and dedicating separate interfaces for the other services. In a large-scale network with multiple subnetworks and heavy log traffic, using multiple interfaces for device management and log collection also enables network segmentation and load balancing (see Configure Panorama to Use Multiple Interfaces).

When assigning Panorama services to various interfaces, keep in mind that only the MGT interface allows administrative access to Panorama for configuration and monitoring tasks. You can assign any interface to the other services when you Perform Initial Configuration of the M-Series Appliance. The M-Series Appliance Hardware Reference Guides explain where to attach cables for the interfaces.

The M-Series appliances, with the exception of the M-700, do not support Link Aggregation Control Protocol (LACP) for aggregating interfaces. The M-700 supports LACP for aggregate interface bond1.

Supported Interfaces

Interfaces can be used for device management, log collection, Collector Group communication, licensing and software updates. See Configure Panorama to Use Multiple Interfaces for more information on network segmentation.

 ©2025 Palo Alto Networks, Inc.

Interface	Maximum Speed	M-700 Appliance	M-600 Appliance	M-300 Appliance	M-200 Appliance
Management (MGT)	1Gbps	—	✓	—	✓
	10Gbps	✓	—	✓	—
Ethernet 1 (Eth1)	1Gbps	—	✓	—	✓
	10Gbps	✓	—	✓	—
Ethernet 2 (Eth2)	1Gbps	—	✓	—	✓
Ethernet 3 (Eth3)	1Gbps	—	✓	—	✓
Ethernet 4 (Eth4)	10Gbps	—	✓	—	—
Ethernet 5 (Eth5)	10Gbps	—	✓	—	—

The M-700 Appliance has two ports on its back panel labeled Ethernet 1/2 and Ethernet 1/3; however, the appliance uses a 20Gb aggregate software interface called ***bond1*** *instead of separate Eth2 and Eth3 subinterfaces.*

Logging Rates

Review the logging rates for the all M-Series appliance models. To achieve the logging rates listed below, the M-Series appliance must be a single log collector in a collector group and you must install all the logging disks for your M-Series model.

Model Capacities and Features	M-700 Appliance	M-600 Appliance	M-300 Appliance	M-200 Appliance
Maximum Logging Rate for Panorama in Management Only mode	Local log storage is not supported			
Maximum Logging Rate for Panorama in Panorama Mode	36,500 logs/ second	25,000 logs/ second	16,500 logs/ second	10,000 logs/ second
Maximum Logging Rate for Panorama	73,000 logs/ second	50,000 logs/ second	33,000 logs/ second	28,000 logs/ second

 ©2025 Palo Alto Networks, Inc.

Model Capacities and Features	M-700 Appliance	M-600 Appliance	M-300 Appliance	M-200 Appliance
in Log Collector Mode				
Maximum Log Storage on Appliance	48TB (12x8TB RAID disk)	48TB (12x8TB RAID disk)	16TB (4x8TB RAID disk)	16TB (4x8TB RAID disk)
Default Log Storage on Appliance	16TB (4x8TB RAID disks)	16TB (4x8TB RAID disks)	16TB (4x8TB RAID disks)	16TB (4x8TB RAID disks)
SSD Storage on Appliance (for logs that M-Series appliances generate)	240GB	240GB	240GB	240GB
NFS Attached Log Storage	Not available			

Perform Initial Configuration of the M-Series Appliance

By default, Panorama has an IP address of 192.168.1.1 and a username/password of admin/admin. For security reasons, you must change these settings before continuing with other configuration tasks. You must perform these initial configuration tasks either from the Management (MGT) interface or using a direct serial port connection to the console port on the M-700, M-600, M-500, M-300, or M-200 appliance.

If you are configuring an M-Series appliance in Log Collector mode with 10GB interfaces, you must complete this entire configuration procedure for the 10GB interfaces to display as Up.

STEP 1 | Gather the required interface and server information from your network administrator.

- Gather the IP address, netmask (for IPv4) or prefix length (for IPv6), and default gateway for each interface that you plan to configure (MGT, Eth1, Eth2, Eth3, Eth4, Eth5). Only the MGT interface is mandatory.

  Palo Alto Networks recommends that you specify all these settings for the MGT interface. If you omit values for some of these settings (such as the default gateway), you can access Panorama only through the console port for future configuration changes. You cannot commit the configurations for other interfaces unless you specify all these settings.

  If you plan to use the appliance as a Panorama management server, Palo Alto Networks recommends using the MGT interface only for managing Panorama and using other

 ©2025 Palo Alto Networks, Inc.

interfaces for managing devices, collecting logs, communicating with Collector Groups, and deploying updates to devices (see M-Series Appliance Interfaces).

- Gather the IP addresses of the DNS servers.

STEP 2 | Access the M-Series appliance from your computer.

1. Connect to the M-Series appliance in one of the following ways:
   - Attach a serial cable from a computer to the Console port on the M-Series appliance and connect using terminal emulation software (9600-8-N-1).
   - Attach an RJ-45 Ethernet cable from a computer to the MGT port on the M-Series appliance. From a browser, go to https://192.168.1.1. Enabling access to this URL might require changing the IP address on the computer to an address in the 192.168.1.0 network (for example, 192.168.1.2).
2. When prompted, log in to the appliance using the default username and password (admin/admin). The appliance starts initializing.

STEP 3 | Change the default admin password.

Starting with PAN-OS 9.0.4, the predefined, default administrator password (admin/admin) must be changed on the first login on a device. The new password must be a minimum of eight characters and include a minimum of one lowercase and one uppercase character, as well as one number or special character.

Be sure to use the best practices for password strength to ensure a strict password and review the password complexity settings.

1. Click the **admin** link in the lower left of the web interface.
2. Enter the **Old Password**, **New Password**, and **Confirm New Password**, and then click **OK**. Store the new password in a safe location.

*To ensure that the MGT interface remains secure, configure Minimum Password Complexity settings (select **Panorama** > **Setup** > **Management**) and specify the interval at which administrators must change their passwords.*

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the network access settings for each interface that you will use to manage Panorama, manage devices, collect logs, communicate with Collector Groups, and deploy updates to devices.

To configure connectivity to Panorama using an IPv6 IP address, you must configure both an IPv4 and IPv6 to successfully configure Panorama using an IPv6 IP address. Panorama does not support configuring the management interface with only an IPv6 IP address.

1. Select **Panorama** > **Setup** > **Interfaces** and click the Interface Name.
2. (Non-MGT interfaces only) **Enable** the interface.
3. Edit the network access settings of each interface that Panorama will use. Only the MGT interface is required. The Eth1, Eth2, Eth3, Eth4, and Eth5 interfaces are optional and apply only if you plan to use the M-Series appliance as a Panorama management server.
   1. Complete one or both of the following field sets based on the IP protocols of your network:

      IPv4—**Public IP Address**, **IP Address**, **Netmask**, and **Default Gateway**

      *If your firewalls connect to the Panorama management server using a public IP address that is translated to a private IP address (NAT), enter the public IP in the **Public IP Address** field, and the private IP in the **IP Address** field to push both addresses to your firewalls.*

      IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
   2. Select the Device Management Services that the interface supports:

      Device Management and Device Log Collection—You can assign one or more interfaces.

      Collector Group Communication—You can assign only one interface.

      Device Deployment (software and content updates)—You can assign only one interface.
   3. (Optional) Select the Network Connectivity Services that the interface supports.

      *(MGT interface only) Disable **Telnet** and **HTTP**; these services use plaintext and so are less secure than other services.*
   4. Click **OK** to save your changes.

STEP 5 | Configure the hostname, time zone, and general settings.

1. Select **Panorama** > **Setup** > **Management** and edit the General Settings.
2. Align the clock on Panorama and the managed firewalls to use the same **Time Zone**, for example GMT or UTC. If you plan to use the Strata Logging Service, you must configure NTP so that Panorama can stay in sync with the Strata Logging Service.

   The firewall records timestamps when it generate logs and Panorama records timestamps upon receiving the logs. Aligning the time zones ensures that the timestamps

 ©2025 Palo Alto Networks, Inc.

are synchronized and that the process of querying logs and generating reports on Panorama is harmonious.

3. Enter a **Hostname** for the server. Panorama uses this as the display name/label for the appliance. For example, this is the name that appears at the CLI prompt. It also appears in the Collector Name field if you add the appliance as a managed collector on the **Panorama** > **Managed Collectors** page.
4. (Optional) Enter the **Latitude** and **Longitude** to enable accurate placement of the M-Series appliance on the world map. The **App Scope** > **Traffic Maps** and **App Scope** > **Threat Maps** use these values.
5. Click **OK** to save your entries.

STEP 6 | Configure the DNS servers and Palo Alto Networks Update Server.

1. Select **Panorama** > **Setup** > **Services** and edit the settings.
2. Enter the IP address of the **Primary DNS Server** and (optionally) of the **Secondary DNS Server**.
3. Enter the URL or static address of the **Update Server** (default updates.paloaltonetworks.com).

   *Select **Verify Update Server Identity** if you want Panorama to verify that the Update Server from which it downloads software or content packages has an SSL certificate that a trusted authority signed. This option adds an additional level of security for communication between the Panorama management server and Update Server.*
4. Click **OK** to save your entries.

STEP 7 | Commit your configuration changes.

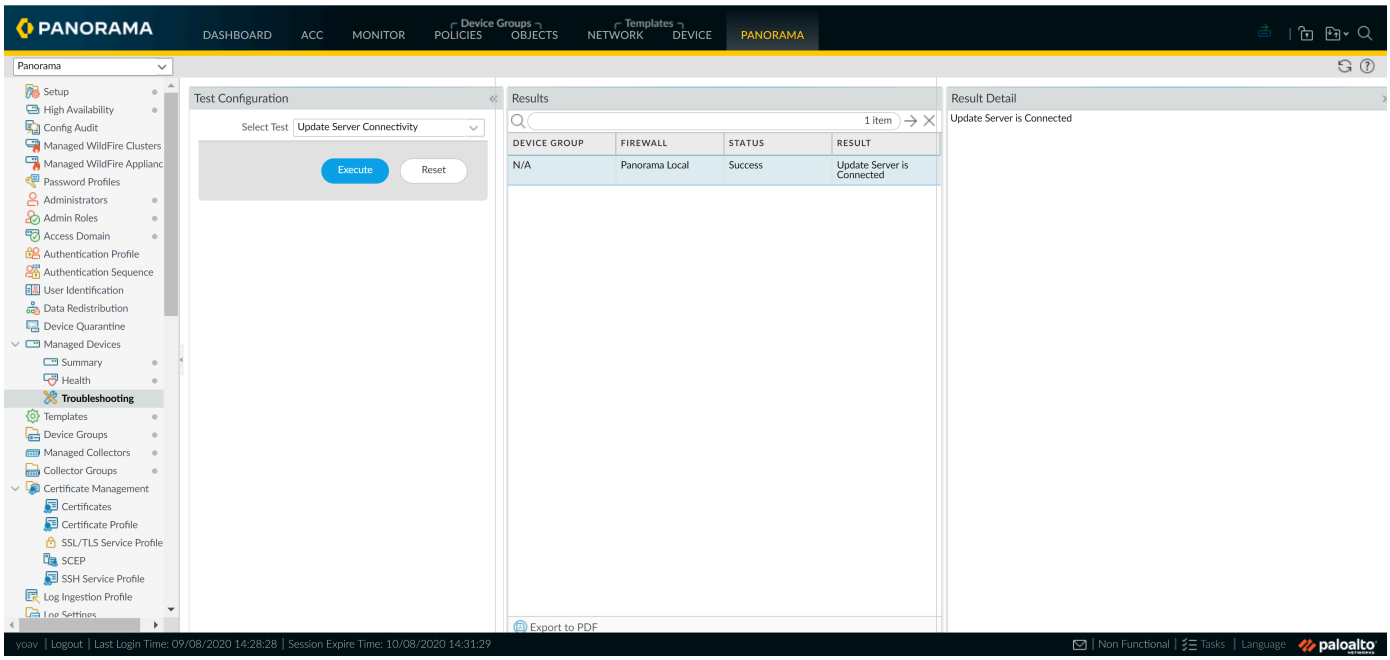
Select **Commit** > **Commit to Panorama** and **Commit** your changes.

*If you plan to use the M-Series appliance as a Panorama management server and you configured interfaces other than MGT, you must assign those interfaces to the **Device Log Collection** or **Collector Group Communication** functions when you Configure a Managed Collector. To make the interfaces operational, you must then Configure a Collector Group for the managed collector and perform a Collector Group commit.*

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Verify network access to external services required for Panorama management, such as the Palo Alto Networks Update Server.

1. Connect to the M-Series appliance in one of the following ways:
   - Attach a serial cable from your computer to the Console port on the M-Series appliance. Then use a terminal emulation software (9600-8-N-1) to connect.
   - Use terminal emulation software such as PuTTY to open an SSH session to the IP address that you specified for the MGT interface of the M-Series appliance during initial configuration.
2. Log in to the CLI when prompted. Use the default admin account and the password that you specified during initial configuration.
3. Use the Update Server Connectivity test to verify network connectivity to the Palo Alto Networks Update Server as shown in the following example.
   1. Select **Panorama** > **Managed Devices** > **Troubleshooting**, and select **Updates Server Connectivity** from the Select Test drop-down.
   2. **Execute** the update server connectivity test.



4. Use the following CLI command to retrieve information on the support entitlement for Panorama from the Update Server:

```
admin> request support check
```

If you have connectivity, the Update Server responds with the support status for Panorama. Because Panorama is not registered, the Update Server returns the following message:

```
Contact Us
https://www.paloaltonetworks.com/company/contact-us.html
Support Home
```

 ©2025 Palo Alto Networks, Inc.

```
https://www.paloaltonetworks.com/support/tabs/overview.html
Device not found on this update server
```

STEP 9 | Next steps...

1. Register Panorama and Install Licenses.
2. Install Content and Software Updates for Panorama.

As a best practice, replace the default certificate that Panorama uses to secure HTTPS traffic over the MGT interface.

Perform Initial Configuration of an Air Gapped M-Series Appliance

Perform the initial configuration for an air gapped M-Series appliance. By default, Panorama has an IP address of 192.168.1.1 and a username/password of admin/admin. For security reasons, you must change these settings before continuing with other configuration tasks. You must perform these initial configuration tasks either from the Management (MGT) interface or using a direct serial port connection to the console port on the M-700, M-600, M-500, M-300, or M-200 appliance.

The air gapped Panorama cannot connect to the Palo Alto Networks update server because an outbound internet connection is required. To activate licenses, upgrade the PAN-OS software version, and install dynamic content updates you must upload the relevant files to the air gapped Panorama manually.

If you are configuring an M-Series appliance in Log Collector mode with 10GB interfaces, you must complete this entire configuration procedure for the 10GB interfaces to display as Up.

STEP 1 | Gather the required information from your network administrator.

- Private IP address for the management (MGT) port
- Netmask
- Default gateway
- DNS server address
- NTP server address

STEP 2 | Install and power on M-Series appliance.

Review your M-Series appliance hardware reference guide for details and best practices.

STEP 3 | Connect to the M-Series appliance.

You must log in using the default **admin** username. You are immediately prompted to change the default `admin` password before you can continue. The new password must be a minimum of eight characters and include a minimum of one lowercase and one uppercase character, as well as one number or special character.

You can connect to the M-Series appliance in one of the following ways:

- Connect a serial cable from your computer to the Console port and connect to the M-Series appliance using terminal emulation software (9600-8-N-1). Wait a few minutes for the

 ©2025 Palo Alto Networks, Inc.

boot-up sequence to complete; when the M-Series appliance is ready, the prompt changes to the name of the M-Series, for example `M-500 login`.

- Log in to the Panorama CLI by connecting an RJ-45 Ethernet cable from your computer to the MGT interface on the M-Series appliance. From a browser, go to **`https://192.168.1.1`**.

  You may need to change the IP address on your computer to an address in the 192.168.1.0/24 network, such as 192.168.1.2, to access this URL.

STEP 4 | Configure the network settings for the air gapped M-Series appliance.

The following commands set the interface IP allocation to `static`, configures the IP address for the MGT interface, the Domain Name Server (DNS), and Network Time Protocol (NTP) server.

```
admin> configure
```

```
admin# set deviceconfig system type static
```

```
admin# set deviceconfig system ip-address <IP-Address> netmask
 <Netmask-IP> default-gateway <Gateway-IP>
```

```
admin# set deviceconfig system dns-settings servers primary <IP-
Address> secondary <IP-Address>
```

```
admin# set deviceconfig system ntp-servers primary-ntp-server ntp-
server-address <IP-Address>
```

```
admin# set deviceconfig system ntp-servers secondary-ntp-server
 ntp-server-address <IP-Address>
```

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Register the M-Series appliance with the Palo Alto Networks Customer Support Portal (CSP).

1. Log in to the Palo Alto Networks CSP.
2. Click **Register a Device**.
3. Select **Register device using Serial Number** and click **Next**.
4. Enter the required `Device Information`.
   - Enter the M-Series appliance **Serial Number**.
   - Check (enable) **Device will be used offline**.
   - Select the PAN-OS **OS Release** running on the M-Series appliance.
5. Enter the required `Location Information`.
   - Enter the **City** the M-Series appliance is located in,
   - Enter the **Postal Code** the M-Series appliance is located in,
   - Enter the **Country** the M-Series appliance is located in.
6. **Agree and Submit**.
7. **Skip this step** when prompted to generate the optional `Day 1 Configuration` config file.

STEP 6 | Download the Panorama license keys.

The license key files are required to activate your Panorama licenses when air gapped.

1. Log in to the Palo Alto Networks CSP.
2. Select **Product** > **Devices** and locate the M-Series appliance you added.
3. Download all license keys files from the download links available `License` column.

   You must download a license key file for each license you want to active on Panorama.

STEP 7 | Active the Panorama licenses.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Licenses** and **Manually upload license key**.

   Click **Choose File** to select the license key file you downloaded in the previous step and click **OK**.
3. Repeat this step to uploaded and activate all licenses.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | (Optional) Configure general Panorama settings as needed.

1. Select **Panorama** > **Setup** > **Management** and edit the General Settings.
2. Enter a **Hostname** for Panorama and enter your network **Domain** name. The domain name is just a label; it will not be used to join the domain.
3. Enter **Login Banner** text that informs users who are about to log in that they require authorization to access the Panorama management functions.

   As a best practice, avoid using welcoming verbiage. Additionally, you should ask your legal department to review the banner message to ensure it adequately warns that unauthorized access is prohibited.
4. Enter the **Latitude** and **Longitude** to enable accurate placement of the M-Series on the world map.
5. Click **OK**.
6. **Commit** and **Commit to Panorama**.

 ©2025 Palo Alto Networks, Inc.

STEP 9 | Upgrade the PAN-OS and dynamic content versions on Panorama.

Review the PAN-OS Upgrade Guide and PAN-OS Release Notes for detailed information about your target PAN-OS upgrade version.

1. Log in to the Palo Alto Networks CSP.
2. Download dynamic content updates.

   Alternatively, you can use a Secure Copy Protocol (SCP) server to automatically download dynamic content updates for Panorama, managed firewalls, Log Collectors, and WildFire appliances. An outbound internet connection is required for the SCP server to download dynamic content updates from the Palo Alto Networks Update Server.

   1. Select **Updates** > **Dynamic Updates**.
   2. Select the dynamic `Content type` you want to install.
   3. **Download** the dynamic content update to your local device.
   4. Repeat this step to download all required dynamic content updates.
3. Download a PAN-OS software update.
   1. Select **Updates** > **Software Updates**.
   2. For the `Content type`, select **Panorama M Base**. For the `Release type`, select **All**(default) or **Preferred**.
   3. In the `Download` column, click the PAN-OS version to download the software image to your local device.
4. Log in to the Panorama web interface.
5. Select **Panorama** > **Dynamic Updates** and **Upload** the dynamic content updates you downloaded.

   Repeat this step to **Browse** and select all the dynamic content release versions.
6. **Install** the dynamic content updates.
7. Select **Panorama** > **Software** and **Upload** the PAN-OS software image you download.
8. **Install** the PAN-OS software version.

   Panorama needs to restart to finish installing the PAN-OS software upgrade.

STEP 10 | Connect Panorama to your network.

1. Disconnect Panorama from your computer.
2. Connect the MGT port to a switch port on your management network using an RJ-45 Ethernet cable. Make sure that the switch port you cable on Panorama is configured for autonegotiation.

M-Series Setup Overview

Use the following procedures to set up an M-Series appliance:

- Set Up an M-Series Appliance in Management Only Mode
- Set Up an M-Series Appliance in Panorama Mode

 ©2025 Palo Alto Networks, Inc.

- Set Up an M-Series Appliance in Log Collector Mode

Set Up an M-Series Appliance in Management Only Mode

Set up the Panorama management server in Management Only mode to dedicate Panorama to managing firewalls and Dedicated Log Collectors. Panorama in Management Only mode have no log collection capabilities, except for config and system logs, and requires a Dedicated Log Collector to store logs.

*If you configured a local Log Collector, the local Log Collector still exists on Panorama when you change to Management Only mode despite having no log collection capabilities. Deleting the local Log Collector (**Panorama** > **Managed Collectors**) deletes the Eth1/1 interface configuration the local Log Collector uses by default. If you decide to delete the local Log Collector, you must reconfigure the Eth1/1 interface.*

STEP 1 | Rack mount the M-Series appliance. Refer to the M-Series Appliance Hardware Reference Guide for instructions.

STEP 2 | Perform Initial Configuration of the M-Series Appliance.

STEP 3 | Register Panorama and Install Licenses.

STEP 4 | Install content and software updates on Panorama.

STEP 5 | Change to Management Only mode.

1. Log in to the Panorama CLI.
2. Switch from Panorama mode to Management Only mode:
   ```
   request system system-mode management-only
   ```
3. Enter **Y** to confirm the mode change. The Panorama management server reboots. If the reboot process terminates your terminal emulation software session, reconnect to the Panorama management server to see the Panorama login prompt.

   If you see a `CMS Login` prompt, this means the Panorama management server has not finished rebooting. Press Enter at the prompt without typing a username or password.
4. Log back in to the CLI.
5. Verify that the switch to Management Only mode succeeded:
   ```
   show system info | match system-mode
   ```

   If the mode change succeeded, the output displays:

   ```
   system mode:management-only
   ```

STEP 6 | Set Up Administrative Access to Panorama

STEP 7 | Manage Firewalls

STEP 8 | Manage Log Collection

Set Up an M-Series Appliance in Panorama Mode

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Rack mount the M-Series appliance. Refer to the M-Series Appliance Hardware Reference Guide for instructions.

STEP 2 | Perform Initial Configuration of the M-Series Appliance.

STEP 3 | Register Panorama and Install Licenses.

STEP 4 | Install Content and Software Updates for Panorama.

STEP 5 | Configure each array. This task is required to make the RAID disks available for logging. Optionally, you can add disks to Increase Storage on the M-Series Appliance.

STEP 6 | Set Up Administrative Access to Panorama.

STEP 7 | Manage Firewalls.

STEP 8 | Manage Log Collection.

Set Up an M-Series Appliance in Log Collector Mode

STEP 1 | Rack mount the M-Series appliance. Refer to the M-Series Appliance Hardware Reference Guide for instructions.

STEP 2 | Perform Initial Configuration of the M-Series Appliance

STEP 3 | Register Panorama and Install Licenses

STEP 4 | Install Content and Software Updates for Panorama

STEP 5 | Configure each array. This task is required to make the RAID disks available for logging. Optionally, you can add disks to Increase Storage on the M-Series Appliance.

STEP 6 | Set Up the M-Series Appliance as a Log Collector

STEP 7 | Manage Log Collection

Set Up the M-Series Appliance as a Log Collector

If you want a dedicated appliance for log collection, configure an M-200, M-300, M-500, M-600, or M-700 appliance in Log Collector mode. To do this, you first perform the initial configuration of the appliance in Panorama mode, which includes licensing, installing software and content updates, and configuring the management (MGT) interface. You then switch the M-Series appliance to Log Collector mode and complete the Log Collector configuration. Additionally, if you want to use dedicated M-Series Appliance Interfaces (recommended) instead of the MGT interface for log collection and Collector Group communication, you must first configure the interfaces for the Panorama management server, then configure them for the Log Collector, and then perform a Panorama commit followed by a Collector Group commit.

Perform the following steps to set up a new M-Series appliance as a Log Collector or to convert an existing M-Series appliance that was previously deployed as a Panorama management server.

 ©2025 Palo Alto Networks, Inc.

If you are configuring an M-Series appliance in Log Collector mode with 10GB interfaces, you must complete this entire configuration procedure for the 10GB interfaces to display as Up.

Switching the M-Series appliance from Panorama mode to Log Collector mode reboots the appliance, deletes the local Log Collector, deletes any existing log data, and deletes all configurations except the management access settings. Switching the mode does not delete licenses, software updates, or content updates.

STEP 1 | Set up the Panorama management server that will manage the Log Collector if you have not already done so.

Perform one of the following tasks:

- Set Up the Panorama Virtual Appliance
- Set Up the M-Series Appliance

STEP 2 | Record the management IP addresses of the Panorama management server.

If you deployed Panorama in a high availability (HA) configuration, you need the IP address of each HA peer.

1. Log in to the web interface of the Panorama management server.
2. Record the **IP Address** of the solitary (non-HA) or active (HA) Panorama by selecting **Panorama** > **Setup** > **Management** and checking the Management Interface Settings.
3. For an HA deployment, record the **Peer HA IP Address** of the passive Panorama by selecting **Panorama** > **High Availability** and checking the Setup section.

STEP 3 | Set up the M-Series appliance that will serve as a Dedicated Log Collector.

If you previously deployed this appliance as a Panorama management server, you can skip this step because the MGT interface is already configured and the licenses and updates are already installed.

The M-Series appliance in Log Collector mode does not have a web interface for configuration tasks, only a CLI. Therefore, before changing the mode on the M-Series appliance, use the web interface in Panorama mode to:

1. Perform Initial Configuration of the M-Series Appliance.
2. Register Panorama and Install Licenses.
3. Install Content and Software Updates for Panorama.

STEP 4 | Access the CLI of the M-Series appliance.

1. Connect to the M-Series appliance in one of the following ways:
   - Attach a serial cable from your computer to the Console port on the M-Series appliance. Then use terminal emulation software (9600-8-N-1) to connect.
   - Use terminal emulation software such as PuTTY to open an SSH session to the IP address that you specified for the MGT interface of the M-Series appliance during initial configuration.
2. Log in to the CLI when prompted. Use the default admin account and the password that you specified during initial configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Switch from Panorama mode to Log Collector mode.

1. Switch to Log Collector mode by entering the following command:

   ```
   > request system system-mode logger
   ```

2. Enter **Y** to confirm the mode change. The M-Series appliance reboots. If the reboot process terminates your terminal emulation software session, reconnect to the M-Series appliance to see the Panorama login prompt.

   *If you see a **CMS Login** prompt, this means the Log Collector has not finished rebooting. Press Enter at the prompt without typing a username or password.*

3. Log back in to the CLI.
4. Verify that the switch to Log Collector mode succeeded:

   ```
   > show system info | match system-mode
   ```

   If the mode change succeeded, the output displays:

   ```
   system-mode: logger
   ```

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Configure the logging disks as RAID1 pairs.

If you previously deployed the appliance as a Panorama management server, you can skip this step because the disk pairs are already configured and available.

The time required to configure the drives varies from several minutes to a couple of hours, based on the amount of data on the drives.

1. Determine which disk pairs are present for configuring as RAID pairs on the M-Series appliance:

   ```
   > show system raid detail
   ```

   Perform the remaining steps to configure each disk pair that has `present` disks. This example uses disk pair A1/A2.
2. To add the first disk in the pair, enter the following command and enter **y** when prompted to confirm the request:

   ```
   > request system raid add A1
   ```

   Wait for the process to finish before adding the next disk in the pair. To monitor the progress of the RAID configuration, re-enter:

   ```
   > show system raid detail
   ```

   After the process finishes for the first disk, the output displays the disk pair status as `Available` but `degraded`.
3. Add the second disk in the pair:

   ```
   > request system raid add A2
   ```
4. Verify that the disk setup is complete:

   ```
   > show system raid detail
   ```

   After the process finishes for the second disk, the output displays the disk pair status as `Available` and `clean`:

   ```
   Disk Pair A      Available
   Status         clean
   ```

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Enable connectivity between the Log Collector and Panorama management server.

Enter the following commands at the Log Collector CLI, where <*IPaddress1*> is for the MGT interface of the solitary (non-HA) or active (HA) Panorama and <*IPaddress2*> is for the MGT interface of the passive (HA) Panorama, if applicable.

```
> configure
# set deviceconfig system panorama-server <IPaddress1> panorama-
server-2 <IPaddress2>
# commit
# exit
```

STEP 8 | Record the serial number of the Log Collector.

You need the serial number to add the Log Collector as a managed collector on the Panorama management server.

1. At the Log Collector CLI, enter the following command to display its serial number.

   ```
   > show system info | match serial
   ```

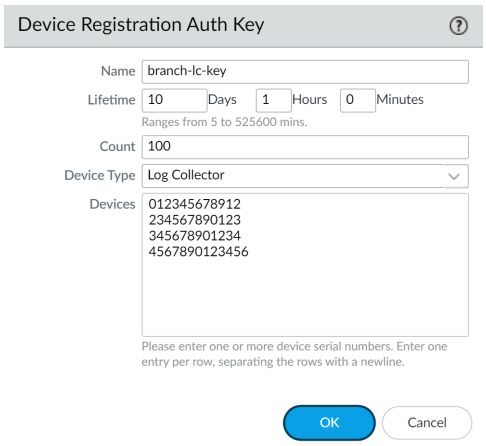
2. Record the serial number.

 ©2025 Palo Alto Networks, Inc.

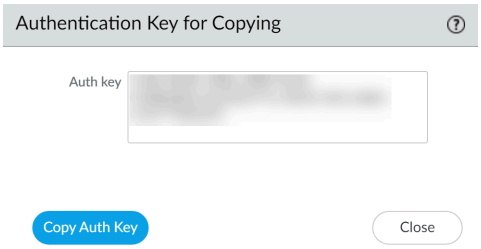
STEP 9 | Create a device registration authentication key.

1. Select **Panorama** > **Device Registration Auth Key** and **Add** a new authentication key.
2. Configure the authentication key.
   - **Name**—Add a descriptive name for the authentication key.
   - **Lifetime**—Specify the key lifetime for how long you can use the authentication key to onboard new Log Collectors.
   - **Count**—Specify how many times you can use the authentication key to onboard new Log Collectors.
   - **Device Type**—Specify that this authentication key is used to authenticate only a **Log Collector**.

     *You can select **Any** to use the device registration authentication key to onboard firewalls, Log Collectors, and WildFire appliances.*
   - (Optional) **Devices**—Enter one or more device serial numbers to specify for which Log Collectors the authentication key is valid.
3. Click **OK**.



4. **Copy Auth Key** and **Close**.



 ©2025 Palo Alto Networks, Inc.

STEP 10 | Add the Log Collector as a managed collector to the Panorama management server.

1. Select **Panorama** > **Managed Collectors** and **Add** a managed collector.
2. In the **General** settings, enter the serial number (**Collector S/N**) you recorded for the Log Collector.
3. In the **Panorama Server IP** field, enter the IP address or FQDN of the solitary (non-HA) or active (HA) Panorama. For HA deployments, enter the IP address or FQDN of the passive Panorama peer in the **Panorama Server IP 2** field.

   These IP addresses must specify a Panorama interface that has **Device Management and Device Log Collection** services enabled. By default, these services are enabled only on the MGT interface. However, you might have enabled the services on other interfaces when you Set Up the M-Series Appliance that is a Panorama management server.
4. Select **Interfaces**, click **Management**, and configure one or both of the following field sets for the MGT interface based on the IP protocols of your network.
   - IPv4—**IP Address**, **Netmask**, and **Default Gateway**
   - IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
5. Click **OK** twice to save your changes to the Log Collector.
6. Select **Commit** > **Commit to Panorama** and **Commit** your changes to the Panorama configuration.

   This step is required before you can enable logging disks.
7. Verify that **Panorama** > **Managed Collectors** lists the Log Collector you added. The Connected column displays a check mark to indicate that the Log Collector is connected to Panorama. You might have to wait a few minutes before the page displays the updated connection status.

   At this point, the Configuration Status column displays Out of Sync and the Run Time Status column displays disconnected. The status will change to In Sync and connected after you configure a Collector Group (Step Assign the Log Collector to a Collector Group.).

STEP 11 | Add the device registration authentication key to the Log Collector.

Add the device registration authentication key only to a Dedicated Log Collector. A Panorama in Panorama mode does not need to authenticate its own local Log Collector.

1. Log in to the Log Collector CLI.
2. Add the device registration authentication key.

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



 ©2025 Palo Alto Networks, Inc.

STEP 12 | Enable the logging disks.

1. Select **Panorama** > **Managed Collectors** and edit the Log Collector.
2. Select **Disks** and **Add** each RAID disk pair.
3. Click **OK** to save your changes.
4. Select **Commit** > **Commit to Panorama** and **Commit** your changes to the Panorama configuration.

STEP 13 | (Recommended) Configure the **Ethernet1**, **Ethernet2**, **Ethernet3**, **Ethernet4**, and **Ethernet5** interfaces if the Panorama management server and Log Collector will use them for **Device Log Collection** (receiving logs from firewalls) and **Collector Group Communication**.

If you previously deployed the Log Collector as a Panorama management server and configured these interfaces, you must reconfigure them because switching to Log Collector

 ©2025 Palo Alto Networks, Inc.

mode (Switch from Panorama mode to Log Collector mode.) would have deleted all configurations except the management access settings.

1. Configure each interface on the Panorama management server (other than the MGT interface) if you haven't already:
   1. Select **Panorama** > **Setup** > **Interfaces** and click the Interface Name.
   2. Select *<interface-name>* to enable the interface.
   3. Complete one or both of the following field sets based on the IP protocols of your network:

      IPv4—**IP Address**, **Netmask**, and **Default Gateway**

      IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
   4. Select the Device Management Services that the interface supports:

      Device Management and Device Log Collection—You can assign one or more interfaces.

      Collector Group Communication—You can assign only one interface.

      Device Deployment (software and content updates)—You can assign only one interface.
   5. Click **OK** to save your changes.
2. Configure each interface on the Log Collector (other than the MGT interface):
   1. Select **Panorama** > **Managed Collectors** and edit the Log Collector.
   2. Select **Interfaces** and click the name of the interface.
   3. Select *<interface-name>* to enable the interface.
   4. Complete one or both of the following field sets based on the IP protocols of your network:

      IPv4—**IP Address**, **Netmask**, and **Default Gateway**

      IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
   5. Select the Device Management Services that the interface supports:

      Device Log Collection—You can assign one or more interfaces.

      Collector Group Communication—You can assign only one interface.
   6. Click **OK** to save your changes to the interface.
3. Click **OK** to save your changes to the Log Collector.
4. Select **Commit** > **Commit to Panorama** and **Commit** your changes to the Panorama configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 14 | (Optional) If your deployment is using custom certificates for authentication between Panorama and managed devices, deploy the custom client device certificate. For more information, see Set Up Authentication Using Custom Certificates.

1. Select **Panorama** > **Certificate Management** > **Certificate Profile** and choose the certificate profile from the drop-down or click **New Certificate Profile** to create one.
2. Select **Panorama** > **Managed Collectors** > **Add** > **Communication** for a Log Collector.
3. Select the **Secure Client Communication** check box.
4. Select the type of device certificate the Type drop-down.
   - If you are using a local device certificate, select the **Certificate** and **Certificate Profile** from the respective drop-downs.
   - If you are using SCEP as the device certificate, select the **SCEP Profile** and **Certificate Profile** from the respective drop-downs.
5. Click **OK**.

 ©2025 Palo Alto Networks, Inc.

STEP 15 | (Optional) Configure Secure Server Communication on a Log Collector. For more information, see Set Up Authentication Using Custom Certificates.

1. Select **Panorama** > **Managed Collectors** > **Add** > **Communication**.
2. Verify that the **Custom Certificate Only** check box is not selected. This allows you to continue managing all devices while migrating to custom certificates.

   When the Custom Certificate Only check box is selected, the Log Collector does not authenticate and cannot receive logs from devices using predefined certificates.
3. Select the SSL/TLS service profile from the **SSL/TLS Service Profile** drop-down. This SSL/TLS service profile applies to all SSL connections between the Log Collector and devices sending it logs.
4. Select the certificate profile from the **Certificate Profile** drop-down.
5. Select **Authorize Client Based on Serial Number** to have the server check clients against the serial numbers of managed devices. The client certificate must have the special keyword $UDID set as the CN to authorize based on serial numbers.
6. In **Disconnect Wait Time (min)**, enter the number of minutes Panorama should wait before breaking and reestablishing the connection with its managed devices. This field is blank by default and the range is 0 to 44,640 minutes.

   The disconnect wait time does not begin counting down until you commit the new configuration.
7. (Optional) Configure an authorization list.
   1. Click **Add** under Authorization List.
   2. Select the **Subject** or **Subject Alt Name** as the Identifier type.
   3. Enter an identifier of the selected type.
   4. Click **OK**.
   5. Select **Check Authorization List** to enforce the authorization list.
8. Click **OK**.
9. Select **Commit** > **Commit to Panorama**.

STEP 16 | Assign the Log Collector to a Collector Group.

1. Configure a Collector Group. You must perform a Panorama commit and then a Collector Group commit to synchronize the Log Collector configuration with Panorama and to put

 ©2025 Palo Alto Networks, Inc.

the Eth1, Eth2, Eth3, Eth4, and Eth5 interfaces (if you configured them) in an operational state on the Log Collector.

In any single Collector Group, all the Log Collectors must run on the same Panorama model: all M-700 appliances, all M-600 appliances, all M-500 appliances, all M-300 appliances, all M-200 appliances, or all Panorama virtual appliances.

*As a best practice, **Enable log redundancy across collectors** if you add multiple Log Collectors to a single Collector group. This option requires each Log Collector to have the same number of logging disks.*

2. Select **Panorama** > **Managed Collectors** to verify that the Log Collector configuration is synchronized with Panorama.

   The Configuration Status column should display In Sync and the Run Time Status column should display connected.

3. Access the Log Collector CLI and enter the following command to verify that its interfaces are operational:

   ```
   > show interface all
   ```

   The output displays the `state` as `up` for each interface that is operational.

4. If the Collector Group has multiple Log Collectors, Troubleshoot Connectivity to Network Resources to verify they can communicate with each other by performing a Ping connectivity test for each interface that the Log Collectors use. For the `source` IP address, specify the interface of one of the Log Collectors. For the `host` IP address, specify the matching interface of another Log Collector in the same Collector Group.

STEP 17 | Next steps...

To enable the Log Collector to receive firewall logs:

1. Configure Log Forwarding to Panorama.
2. Verify Log Forwarding to Panorama.

Increase Storage on the M-Series Appliance

After you Perform Initial Configuration of the M-Series Appliance, you can increase log storage capacity of the appliance by upgrading the existing drive pairs to larger capacity drives or by installing additional drive pairs in empty drive bays. For example, you can choose to upgrade the existing 1TB drives to 2TB on an M-500 appliance, or you can add 2TB drives to the empty drive bays (B1 through D2).

The M-Series appliances leverage RAID 1 for data redundancy in the event of disk failure. Therefore, the pair of drives in a RAID 1 array need to be identical. However, you are free to mix drive capacities across different RAID 1 arrays. For example, the drives in the A1/A2 RAID 1 array can be 1TB drives, and the drives in the B1/B2 RAID 1 array can be 2TB drives.

 ©2025 Palo Alto Networks, Inc.

The following table lists the maximum number of drive bays (disks) and the available drive capacities supported on M-Series appliances.

Because each drive pair (A1/A2 for example) is in a RAID 1 array, the total storage capacity is half of the total drives installed. For example, if an M-500 appliance has 2TB drives installed in drive bays A1/A2 and B1/B2, the A1/A2 array provides 2TB total storage and the B1/B2 array provides another 2TB for a total of 4TB.

Appliance	Number of Supported Drive Bays (Disks)	Supported Drive Capacity
M-200 Appliance	4	8TB
M-300 Appliance	4	8TB
M-500 Appliance	24	1TB or 2TB
M-600 Appliance	12	8TB
M-700 Appliance	12	8TB

Before expanding log storage capacity, Determine Panorama Log Storage Requirements. If you need more log storage than a single M-Series appliance supports, you can add Dedicated Log Collectors (see Configure a Managed Collector) or you can Configure Log Forwarding from Panorama to External Destinations.

*You don't need to take the M-Series appliance offline to expand the storage when adding drives to an M-Series appliance that is already deployed. When the additional drives are configurable and available, the M-Series appliance redistributes the logs among all available drives. This log redistribution process happens in the background and does not impact uptime or the availability of the M-Series appliance. However, the process does diminish the maximum logging rate. The Redistribution State column (**Panorama** > **Collector Groups**) indicates the completion status of the process as a percentage.*

- Add Additional Drives to an M-Series Appliance
- Upgrade Drives on an M-Series Appliance

Add Additional Drives to an M-Series Appliance

STEP 1 | Install the new drives in the appropriate drive bays.

Make sure to add the drives sequentially in the next open drive bays. For example, add drives to B1 and B2 before adding drives to C1 and C2.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Access the command line interface (CLI) on the M-Series appliance.

Connect to the M-Series appliance in one of two ways:

- Connect a serial cable from your computer to the Console port and connect to the M-Series appliance using terminal emulation software (9600-8-N-1).
- Use terminal emulation software (such as PuTTY) to open a Secure Shell (SSH) session to the IP address of the M-Series appliance.

STEP 3 | When prompted, log in to the appliance.

Use the default administrator account and the assigned password.

STEP 4 | Configure each array.

The time required to mirror the data on the drive can take minutes, a few hours, or more than a day depending on the amount of data on the drive.

The following example uses the drives in bays B1 and B2.

1. Enter the following commands and confirm the request when prompted:

```
> request system raid add B1
> request system raid add B2
```

(RMA only) If you want to try and preserve the data on the disks, try the following commands instead:

```
> request system raid add B1 force no-format
> request system raid add B2 force no-format
```

2. To monitor the progress of the RAID configuration, enter the following command:

```
> show system raid detail
```

When the RAID set up is complete, the following response displays:

```
Disk Pair A       Available
   Status         clean
   Disk id A1     Present
       model        : ST91000640NS
       size         : 953869 MB
       status       : active sync
   Disk id A2     Present
       model        : ST91000640NS
       size         : 953869 MB
       status       : active sync
Disk Pair B       Available
   Status         clean
   Disk id B1     Present
       model        : ST91000640NS
       size         : 953869 MB
```

 ©2025 Palo Alto Networks, Inc.

```
        status       : active sync
    Disk id B2     Present
        model        : ST91000640NS
        size         : 953869 MB
        status       : active sync
```

STEP 5 | Make the array available for logging.

To enable the array for logging, you must first add the appliance as a managed collector on Panorama. If not already added, see Configure a Managed Collector.

1. Log in to the web interface of the Panorama management server that manages this Log Collector.
2. Select **Panorama** > **Managed Collectors** and edit the Log Collector.
3. Select **Disks** and **Add** each array.
4. Click **OK** to save your changes.
5. Select **Commit** > **Commit to Panorama** and **Commit** your changes.
6. Select **Commit** > **Push to Devices**, select the Collector Group, and **Push** your changes.

Upgrade Drives on an M-Series Appliance

STEP 1 | Access the command line interface (CLI) on the M-Series appliance.

Connect to the M-Series appliance in one of two ways:

- Connect a serial cable from your computer to the Console port and connect to the M-Series appliance using terminal emulation software (9600-8-N-1).
- Use terminal emulation software (such as PuTTY) to open a Secure Shell (SSH) session to the IP address of the M-Series appliance.

STEP 2 | When prompted, log in to the appliance.

Use the default administrator account and the assigned password.

STEP 3 | Verify that the RAID 1 status for the installed drives shows there are at least two functioning RAID 1 arrays. During the upgrade, you will upgrade one RAID 1 array at a time and there must be at least one other RAID 1 array that is available to the appliance. The appliance will show an abort error if you try to remove the only functioning array from the configuration.

Enter the following command to view RAID status:

```
> show system raid detail
```

For example, the following shows an output from an M-500 appliance with two available arrays (Disk Pair A and Disk Pair B). If there is only one available array, you must add a second array as described in Add Additional Drives to an M-Series Appliance before you upgrade the drives.

```
Disk Pair A                              Available
   Status                                    clean
   Disk id A1                              Present
```

 ©2025 Palo Alto Networks, Inc.

```
        model         : ST91000640NS
        size          : 953869 MB
        status        : active sync
    Disk id A2                          Present
        model         : ST91000640NS
        size          : 953869 MB
        status        : active sync
Disk Pair B                           Available
    Status                                clean
    Disk id B1                          Present
        model         : ST91000640NS
        size          : 953869 MB
        status        : active sync
    Disk id B2                          Present
        model         : ST91000640NS
        size          : 953869 MB
        status        : active sync
```

STEP 4 | Remove the first 1TB drive and replace it with a 2TB drive.

1. To remove the first drive from the RAID 1 array configuration (A1 in this example), enter the following command and enter **y** when prompted to confirm the request:

   ```
   > request system raid remove A1
   ```

2. Physically remove the first drive from the drive bay. Press the ejector button on the drive carrier in drive bay A1 to release the ejector handle. Then pull the handle toward you and slide the drive out of the appliance.
3. Remove a 2TB drive from its packaging and place the drive on a table next to the drive you just removed. Take note of how the drive is installed in the carrier because you will install the 2TB drive in this same carrier.
4. Remove the four screws holding the 1TB drive in the carrier and remove the drive from the carrier.
5. Attach the 2TB drive to the carrier using the same four screws you removed from the 1TB drive and then reinsert the carrier with the 2TB drive into drive bay A1.
6. Enter the following command to verify the 2TB drive is recognized:

   ```
   show system raid detail
   ```

   Verify that the A1 disk shows the correct model and size (about 2TB). If the model and size are not correct, run the above command again until the correct model and size are shown.

   If the wrong model and size are consistently shown, enter the following command:

   ```
   request system raid remove A1
   ```

   Wait for 30 seconds once you run the above command, then remove the disk and reinsert it and repeat the **show system raid detail** command to verify the size and model.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Copy the data from the remaining installed 1TB drive in the RAID 1 array to the newly installed 2TB drive in that array.

The time required to copy the data may vary from several minutes to a few hours, depending on the amount of data on the drive.

1. To copy the data from the 1TB drive in drive bay A2 to the newly installed 2TB drive in drive bay A1, enter the following command and enter **y** when prompted:

```
> request system raid copy from A2 to A1
```

2. To view the status of the copy process, run the following command:

```
> show system raid detail
```

Continue running this command to view the RAID detail output until you see that the array (A1/A2 in this example) shows `Available`.

At this point, drive A2 will show `not in use` because there is a drive size mismatch.

STEP 6 | Upgrade the second drive in the RAID 1 array to a 2TB drive.

1. Remove the second 1TB drive (from drive bay A2 in the current example) for the RAID 1 array configuration:

```
> request system raid remove A2
```

2. Insert the carrier with the newly installed 2TB drive into drive bay A2 and add it to the RAID 1 array configuration:

```
> request system raid add A2
```

The system will copy the data from A2 to A1 to mirror the drives.

3. To view the status of the copy process, run the following command:

```
> show system raid detail
```

Continue to view the RAID detail output until you see that the array (A1/A2 in this example) shows `Available` and both disks show `active sync`.

```
Disk Pair A       Available
   Status         clean
   Disk id A1     Present
       model        : ST2000NX0253
       size         : 1907138 MB
       status       : active sync
Disk id A2     Present
       model        : ST2000NX0253
```

 ©2025 Palo Alto Networks, Inc.

```
        size          : 1907138 MB
        status        : active sync
```

STEP 7 | Upgrade drives for additional RAID 1 arrays as needed.

To upgrade additional RAID 1 arrays to 2TB drives, repeat this procedure replacing the drive designators as applicable. For example, replace A1 with B1 and A2 with B2 to upgrade the drives in the B1/B2 RAID 1 array.

Configure Panorama to Use Multiple Interfaces

In a large-scale network, you can improve security and reduce congestion by implementing network segmentation, which involves segregating the subnetworks based on resource usage, user roles, and security requirements. Panorama supports network segmentation by enabling you to use multiple M-Series Appliance Interfaces for managing devices (firewalls, Log Collectors, and WildFire appliances and appliance clusters) and collecting logs; you can assign separate interfaces to the devices on separate subnetworks.

Using multiple interfaces to collect logs also provides the benefit of load balancing, which is particularly useful in environments where the firewalls forward logs at high rates to the Log Collectors. If you enable the **forward to all Log Collectors** setting in the Collector Group log forwarding preference list, logs are sent on all configured interfaces configured. Otherwise, logs are forwarded over a single interface, and if that interface goes down, log forwarding continues over the next configured interface. For example, you configure Eth1/1, Eth1/2, and Eth1/3 for log forwarding. In the event the Eth1/1 interface goes down, log forwarding continues over Eth1/2.

Because administrators access and manage Panorama over the MGT interface, securing that interface is especially important. One method for improving the security of the MGT interface is to offload Panorama services to other interfaces. In addition to device management and log collection, you can also offload Collector Group communication and deployment of software and content updates to firewalls, Log Collectors, and WildFire appliances and appliance clusters. By offloading these services, you can reserve the MGT interface for administrative traffic and assign it to a secure subnetwork that is segregated from the subnetworks where your firewalls, Log Collectors, and WildFire appliances and appliance clusters reside.

- Multiple Interfaces for Network Segmentation Example
- Configure Panorama for Network Segmentation

Multiple Interfaces for Network Segmentation Example

Figure 1 illustrates a deployment that uses multiple interfaces on M-500 appliances in Panorama mode and Log Collector mode. In this example, the interfaces support network segmentation as follows:

- **Panorama management network**—To protect the Panorama web interface, CLI, and XML API from unauthorized access, the MGT interface on Panorama connects to a subnetwork that only administrators can access.
- **Internet**—Panorama uses the MGT interface to communicate with external services such as the Palo Alto Networks Update Server.
- **Perimeter Gateway** and **Data Center**—Panorama uses a separate pair of interfaces to manage the firewalls and Log Collectors in each of these subnetworks. Managing firewalls typically generates less traffic than querying Log Collectors for report information. Therefore, Panorama

 ©2025 Palo Alto Networks, Inc.

uses 1Gbps interfaces (Eth1 and Eth2) for managing the firewalls and uses 10Gbps interfaces (Eth4 and Eth5) for querying and managing the Log Collectors. Each Log Collector uses its MGT interface to respond to the queries but uses its Eth4 and Eth5 interfaces for the heavier traffic associated with collecting logs from the firewalls.

- **Software and content updates**—The firewalls and Log Collectors in both subnetworks retrieve software and content updates over the Eth3 interface on Panorama.

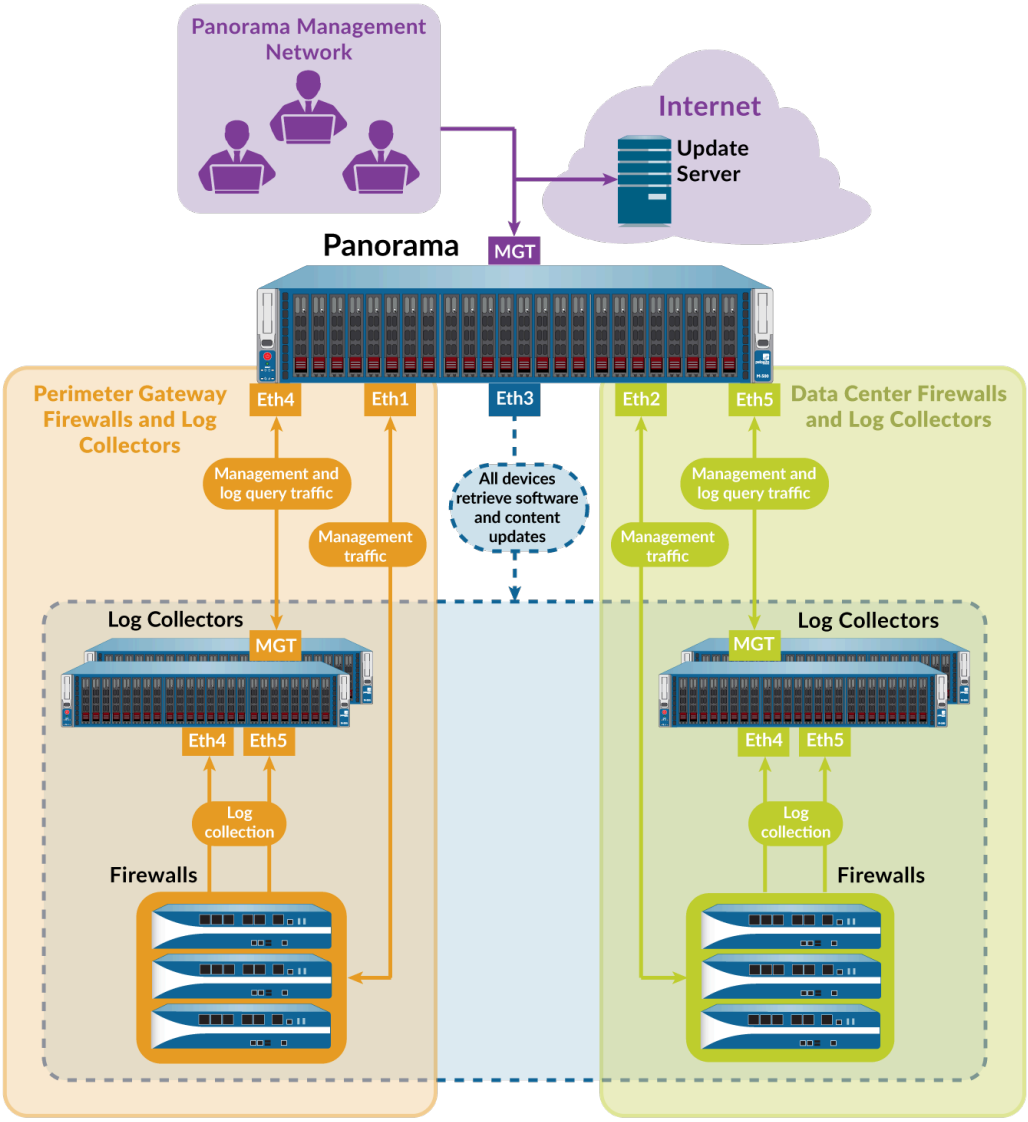


Figure 9: Multiple Panorama Interfaces

Configure Panorama for Network Segmentation

To offload Panorama services from the MGT interface to other interfaces, start by configuring the interfaces on the Panorama management server. If your network has heavy log traffic, remember that the Eth4 and Eth5 interfaces on the M-500, M-600, and M-700 appliances support higher throughput (10Gbps) than the other interfaces (1Gbps). Then, configure the Log Collectors in each subnetwork to connect with specific interfaces on Panorama. For each Log Collector, you also select an interface for Collector Group communication and one or more interfaces for collecting

 ©2025 Palo Alto Networks, Inc.

logs from firewalls. Finally, configure the firewalls in each subnetwork to connect with interfaces on Panorama.

If you are configuring an M-Series appliance in Log Collector mode with 10GB interfaces, you must complete this entire configuration procedure for the 10GB interfaces to display as Up.

Palo Alto Networks recommends that you specify the IP address, netmask (for IPv4) or prefix length (for IPv6), and default gateway for the MGT interface. If you omit one of these settings (such as the default gateway), you can access the M-Series appliance only through the console port for future configuration changes.

Perform the following steps to configure Panorama and Dedicated Log Collectors to use multiple interfaces:

STEP 1 | Verify that the Panorama appliances and firewalls support multiple interfaces, and have the prerequisite software versions and configurations.

- The M-Series appliances must run Panorama 8.0 or later to use a separate interface for deploying updates and to use multiple interfaces for device management and log collection. The M-200 and M-600 appliances must run Panorama 8.1 or later while the M-300 and M-700 must run Panorama 10.2 or later. Panorama appliances deployed on ESXi, vCloud, Air, Hyper-V and KVM must run Panorama 8.1 or later.
- If you deployed a Panorama or Log Collector as a virtual appliance, verify the Supported Interfaces for the Panorama Virtual Appliance.
- The M-Series appliances must run Panorama 6.1 or later to use separate interfaces for log collection or Collector Group communication.
- The initial configuration of each Panorama management server is complete. This includes configuration of the MGT interface.

  To configure an IPv6 IP address for the Panorama MGT interface, you must configure both an IPv4 and IPv6 to successfully configure Panorama using an IPv6 IP address. Panorama does not support configuring the MGT interface with only an IPv6 IP address.

- Log Collectors and Collector Groups are configured. This includes configuration of the MGT interface on the Log Collectors.

  To configure an IPv6 IP address for the MGT interface of a Log Collector, you must configure both an IPv4 and IPv6 to successfully configure Panorama using an IPv6 IP address. Panorama does not support configuring the MGT interface with only an IPv6 IP address.

- The initial configuration of the firewalls is complete, you have added the firewalls to Panorama as managed devices, and the firewalls in each subnetwork are assigned to a separate template.
- The initial configuration of WildFire appliances is complete and you have added WildFire appliances to Panorama as managed devices.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Configure the interfaces on the solitary (non-HA) or active (HA) Panorama management server.

Because the MGT interface was configured during initial Panorama configuration, you don't have to configure it again.

Perform these steps for each interface:

1. Log in to the Panorama Web Interface of the solitary (non-HA) or active (HA) Panorama management server.
2. Select **Panorama** > **Setup** > **Interfaces**.
3. Click an Interface Name to edit the interface.
4. Select <*interface-name*> to enable the interface.
5. Configure one or both of these field sets based on the IP protocols of your network:
   - IPv4—**IP Address**, **Netmask**, and **Default Gateway**
   - IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
6. Select the services that the interface supports:
   - **Device Management and Device Log Collection**—Manage firewalls, Log Collectors, and WildFire appliances and appliance clusters, collect logs that the Log Collectors generate, and query the Log Collectors for report information. To support a segmented network, you can enable these services on multiple interfaces.
   - **Collector Group Communication**—Communicate with the Collector Groups that Panorama manages across all subnetworks.
   - **Device Deployment**—Deploy software and content updates to managed firewalls, Log Collectors, and WildFire appliances and appliance clusters across all subnetworks.
7. Click **OK** to save your changes to the interface.
8. Click **Commit** > **Commit to Panorama** and **Commit** your changes.
9. Click **Commit** > **Push to Devices** and push the changes to the Collector Group that contain the Log Collectors you modified.

STEP 3 | (HA only) Configure the interfaces on the passive Panorama management server.

1. Log in to the Panorama Web Interface of the active Panorama management server.
2. Select **Panorama** > **Managed Collectors** and select the passive HA peer.
3. Select **Interfaces** and click an interface to edit.
4. Check the **Enable Interface** box to enable the interface.
5. Configure one or both of these field sets based on the IP protocols of your network:
   - IPv4—**IP Address**, **Netmask**, and **Default Gateway**
   - IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
6. Select the services that the interface supports:
   - **Device Management and Device Log Collection**—Manage firewalls, Log Collectors, and WildFire appliances and appliance clusters, collect logs that the Log Collectors

 ©2025 Palo Alto Networks, Inc.

generate, and query the Log Collectors for report information. To support a segmented network, you can enable these services on multiple interfaces.

- **Collector Group Communication**—Communicate with the Collector Groups that Panorama manages across all subnetworks.
- **Device Deployment**—Deploy software and content updates to managed firewalls, Log Collectors, and WildFire appliances and appliance clusters across all subnetworks.

7. Click **OK** to save your changes to the interface.
8. Select **Commit** > **Commit and Push** to commit your changes to Panorama and to push the changes to Collector Groups that contain the passive HA peer you modified.

STEP 4 | Configure each Log Collector to connect with a Panorama interface.

To support a segmented network, you can connect the Log Collectors in each subnetwork to separate Panorama interfaces. The interfaces must have **Device Management and Device Log Collection** enabled, as described in the previous step.

1. Log in to the Panorama Web Interface of the solitary (non-HA) or active (HA) Panorama management server.
2. Select **Panorama** > **Managed Collectors** and edit the Log Collector.
3. In the **Panorama Server IP** field, enter the IP address of an interface on the solitary (non-HA) or active (HA) Panorama.
4. (HA only) In the **Panorama Server IP 2** field, enter the IP address of an interface on the passive Panorama that will support **Device Management and Device Log Collection** if failover occurs on the active Panorama.
5. Click **OK** to save your changes.
6. Select **Commit** > **Commit and Push** to commit your changes to Panorama and to push the changes to Collector Groups that contain the Log Collector you modified.
7. Perform the following steps on each Dedicated Log Collector:
   1. Access the Log Collector CLI by using emulation software such as PuTTY to open a SSH session to the Log Collector using its MGT interface IP address. When prompted, log in using Panorama administrator credentials.
   2. Run the following commands, where *<IPaddress1>* is for the solitary (non-HA) or active (HA) Panorama and *<IPaddress2>* is for the passive Panorama (if applicable).

```
> configure
# set deviceconfig system panorama-server <IPaddress1>
 panorama-server-2 <IPaddress2>
# commit
```

 ©2025 Palo Alto Networks, Inc.

STEP 5 | (HA only) Configure an interface on the passive Panorama management server to deploy updates in case the active Panorama fails over.

1. Log in to the Panorama Web Interface of the passive Panorama management server.
2. Select **Panorama** > **Setup** > **Interfaces**.
3. Click an Interface Name to edit the interface.
4. Select *<interface-name>* to enable the interface.
5. Configure one or both of these field sets based on the IP protocols of your network:
   - IPv4—**IP Address**, **Netmask**, and **Default Gateway**
   - IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
6. Select **Device Deployment**.
7. Click **OK** to save your changes.
8. Click **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Configure the interfaces that the Log Collectors will use to collect logs from firewalls and communicate with other Log Collectors.

Because the MGT interface was configured during initial configuration of the Log Collectors, you don't have to configure it again.

1. Log in to the Panorama Web Interface of the solitary (non-HA) or active (HA) Panorama management server.
2. Select **Panorama** > **Managed Collectors** and edit the Log Collector.
3. Select **Interfaces** and perform the following steps for each interface:
   1. Click an interface name to edit that interface.
   2. Select *<interface-name>* to enable the interface.
   3. Configure one or both of the following field sets based on the IP protocols of your network.

      IPv4—**IP Address**, **Netmask**, and **Default Gateway**

      IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
   4. Select the functions that the interface supports:

      Device Log Collection—Collect logs from firewalls. You can load balance the logging traffic by enabling multiple interfaces to perform this function.

      Collector Group Communication—Communicate with other Log Collectors in the Collector Group.
   5. Click **OK** to save your changes to the interface.
4. Click **OK** to save your changes to the Log Collector.
5. Select **Commit** > **Commit and Push** to commit your changes to Panorama and to push the changes to Collector Groups that contain the Log Collectors you modified.
6. Select **Panorama** > **Managed Collectors** to verify that the Log Collectors are synchronized and connected with Panorama.

   The Configuration Status column should display `InSync` and the Run Time Status column should display `connected`.

STEP 7 | Configure the firewalls to connect with a Panorama interface.

To support a segmented network, you can connect the firewalls in each subnetwork to separate Panorama interfaces. The interfaces must have **Device Management and Device**

 ©2025 Palo Alto Networks, Inc.

Log Collection enabled. This step assumes that you use separate templates to configure the firewalls in separate subnetworks.

In this example deployment, Panorama uses these interfaces to manage the firewalls but not to collect firewall logs. You specify which Dedicated Log Collectors will collect firewall logs when you configure Collector Groups.

1. Log in to the Panorama Web Interface of the solitary (non-HA) or active (HA) Panorama management server.
2. On Panorama, select **Device** > **Setup** > **Management**, select a **Template** and edit the Panorama Settings.
3. In the first **Panorama Servers** field, enter the IP address of an interface on the solitary (non-HA) or active (HA) Panorama.
4. (HA only) In the second **Panorama Servers** field, enter the IP address of an interface on the passive Panorama that will support device management if failover occurs.
5. Click **OK** to save your changes.
6. Select **Commit** > **Commit and Push** to commit your changes to Panorama and push the template changes to firewalls.
7. Select **Panorama** > **Managed Devices** to verify that the firewalls are synchronized and connected with Panorama.

   The Device State column should display `Connected`. The Shared Policy and Template columns should display `InSync`.

 ©2025 Palo Alto Networks, Inc.

Register Panorama and Install Licenses

Before you can begin using Panorama for centralized management, logging, and reporting, you are required to register, activate, and retrieve the Panorama device management and support licenses. Every instance of Panorama requires valid licenses that entitle you to manage firewalls and obtain support. The firewall device management license enforces the maximum number of firewalls that Panorama can manage. This license is based on firewall serial numbers, not on the number of virtual systems on each firewall. The support license enables Panorama software updates and dynamic content updates (for the latest Applications and Threats signatures, as an example). Additionally, Panorama virtual appliances on AWS and Azure must be purchased from Palo Alto Networks, and cannot be purchased on the AWS or Azure marketplaces.

After upgrading your Panorama virtual appliance to PAN-OS 8.1, you are prompted if a capacity license has not been successfully installed or if the total number of firewalls being managed by Panorama exceeds the device management license. You have 180 days from the date of upgrade to install a valid device management license if no license has been installed. If the number of managed firewalls exceeds the device management license, you have 180 days to delete firewalls to meet the device management license requirements or upgrade your device management license. All commits fail if a valid device management license is not installed, or the existing device management license limit is not met, within 180 days of upgrade. To purchase a device management license, contact your Palo Alto Networks sales representative or authorized reseller.

If you want to use the cloud-based Strata Logging Service, you require a Strata Logging Service license, in addition to the firewall management license and premium support license. To purchase licenses, contact your Palo Alto Networks Systems Engineer or reseller.

If you are running an evaluation license for firewall management on your Panorama virtual appliance and want to apply a Panorama license that you purchased, perform the tasks Register Panorama *and* Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected.

- Register Panorama
- Activate a Panorama Support License
- Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected
- Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected
- Activate/Retrieve a Firewall Management License on the M-Series Appliance

Register Panorama

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Record the Panorama serial number or auth-code and record your Sales Order Number or Customer ID.

For the auth-code, Sales Order Number, or Customer ID, see the order fulfillment email that Palo Alto Networks Customer Service sent when you placed your order for Panorama.

For the serial number, the location depends on the model:

- M-Series appliance—Log in to the Panorama web interface and record the **Serial #** value in the **Dashboard** tab, General Information section.
- Panorama virtual appliance—See the order fulfillment email or refer to the serial number generated when provisioning Panorama using VM Flex licensing.



The Panorama virtual appliance is automatically registered when you allocate a serial number using VM Flex licensing.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Register Panorama in the Palo Alto Networks Customer Support Portal (CSP).

The steps depend on whether you already have a login for the Palo Alto Networks CSP.

- If this is the first Palo Alto Networks appliance you are registering and you do not yet have a CSP login:
  1. Go to the Palo Alto Networks CSP.
  2. Click **Create my account**.
  3. Enter **Your Email Address** and respond to the reCAPTCHA prompt.
  4. Click **Submit** after you successfully respond to the reCAPTCHA prompt.
  5. Select **Register device using Serial Number or Authorization Code** and click **Submit**
  6. Complete the fields in the **Create Contact Details** and **Create UserID and Password** sections.
  7. Enter the Panorama **Device Serial Number** or **Auth Code**.
  8. Enter your **Sales Order Number** or **Customer ID**.
  9. Respond to the reCAPTCHA prompt.
  10. Click **Submit** after you successfully respond to the reCAPTCHA prompt.
- If you already have a CSP login:
  1. Log in to the Palo Alto Networks CSP.
  2. Click **Assets** > **Devices** > **Register New Device**.

     *You can also **Register a Device** in the CSP **Support Home**.*
  3. Select **Register device using Serial Number** and click **Next**.
  4. Enter the Panorama **Serial Number**.
  5. Enter the **Device Name** to apply a name to search for and identify your Panorama.
  6. (Optional) Select a **Device Tag** to group Panorama with any other devices for which you have selected a device tag.

     The device tag must first be created at the account level (**Assets** > **Devices** > **Device Tag**) before it can be selected when you register Panorama.
  7. If the Panorama management server is not internet-connected, check **Device will be used offline** and select the **OS Release** version.
  8. Enter the required Location Information (as indicated by the asterisks) if you have purchased the 4 hour RMA.
  9. **Agree and Submit** the EULA.

     After you see the registration complete message, close the Device Registration dialog.

Activate a Panorama Support License

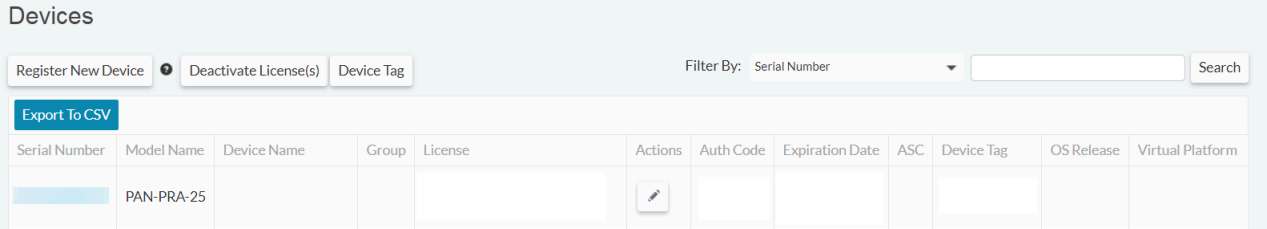
Before activating a Panorama support license on a Panorama M-Series appliance or Panorama virtual appliance, you must Register Panorama.

 ©2025 Palo Alto Networks, Inc.

If the support license expires, Panorama can still manage firewalls and collect logs, but software and content updates will be unavailable. The software and content versions on Panorama must be the same as or later than the versions on the managed firewalls, or else errors will occur. For details, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.

STEP 1 | Log in to the Palo Alto Networks customer support portal to activate the auth-code.

1. Select **Assets** > **Devices** and enter your Panorama serial number to Filter by the **Serial Number**.



2. Select the pencil icon in the Action column, select **Activate Auth-Code** and enter your support license **Authorization Code**, and click **Agree and Submit**.

STEP 2 | Log in to the Panorama web interface, and select **Panorama** > **Support** > **Activate feature using authorization code**.

STEP 3 | Enter the **Authorization Code** and click **OK**.

STEP 4 | Verify that the subscription is activated. Check the details (for example, the **Expiry Date**, support **Level**, and **Description**) in the Support section of the page.

Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected

In order to manage devices on Panorama, you need to activate a firewall management license generated by PAN-OS. The device management license you activate determines the number of devices Panorama can manage. Log Collectors and WildFire appliances are not treated as managed devices and do not count toward the number of devices allotted by the device management license.

Before activating and retrieving a firewall management license on the Panorama virtual appliance, you must Register Panorama. If you are running an evaluation license and want to apply a license that you purchased, you must still register and activate/retrieve the purchased license. Additionally, you must then change the serial number of the Panorama from the evaluation serial number to the production serial number.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Select **Panorama** > **Setup** > **Management** and edit the General Settings.

STEP 3 | Enter the Panorama **Serial Number** (included in the order fulfillment email) and click **OK**.

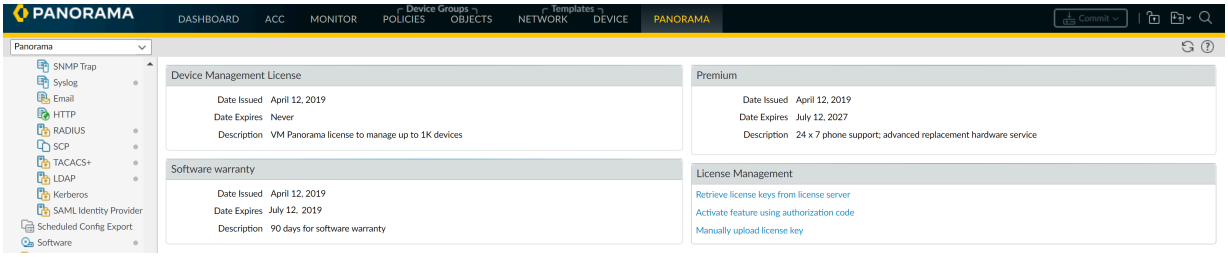
 ©2025 Palo Alto Networks, Inc.

STEP 4 | Select **Panorama** > **Licenses** to activate or retrieve the firewall management license:

- **Retrieve license keys from license server**—Panorama automatically retrieves and activates the firewall management license from the Panorama Update Server.
- **Activate feature using authorization code**—Enter the firewall management license authorization code and click **OK** to activate the license. The authorization code can be obtained from the order fulfillment email or by logging in to the Palo Alto Networks Customer Support web site by finding the Panorama management server.
- **Manually upload license key**—Log in to the Palo Alto Networks Customer Support web site, find your Panorama management server, and download the firewall management license key to your local device. After you download the license key, click **Choose File** to select the license key and click **OK**.

STEP 5 | Verify the firewall management license is activated.

The Device Management License section now appears displaying the date the license was issued, when the license expires, and a description of the firewall management license.



Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected

Before activating and retrieving a firewall management license on the Panorama virtual appliance, you must Register Panorama. In order to manage devices on Panorama, you will need to activate a device management license. The device management license you activate will determine the number of devices Panorama can manage. Log Collectors and WildFire appliances are not treated as managed devices and will not count toward the number of devices allotted by the device management license. If you are running an evaluation license and want to apply a license that you purchased, you must still register and activate/retrieve the purchased license.

After upgrading to PAN-OS 8.1, you will be prompted to retrieve a valid Panorama management license when you first log in to the Panorama web interface when Panorama has finished rebooting. To activate or retrieve the valid management license if the Panorama virtual appliance is offline or unable to reach the Palo Alto Networks update server, you must get the relevant appliance information for the Panorama virtual appliance and upload it to the Customer Support web site.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | (Initial Deployment only) Enter the Panorama **Serial Number**.

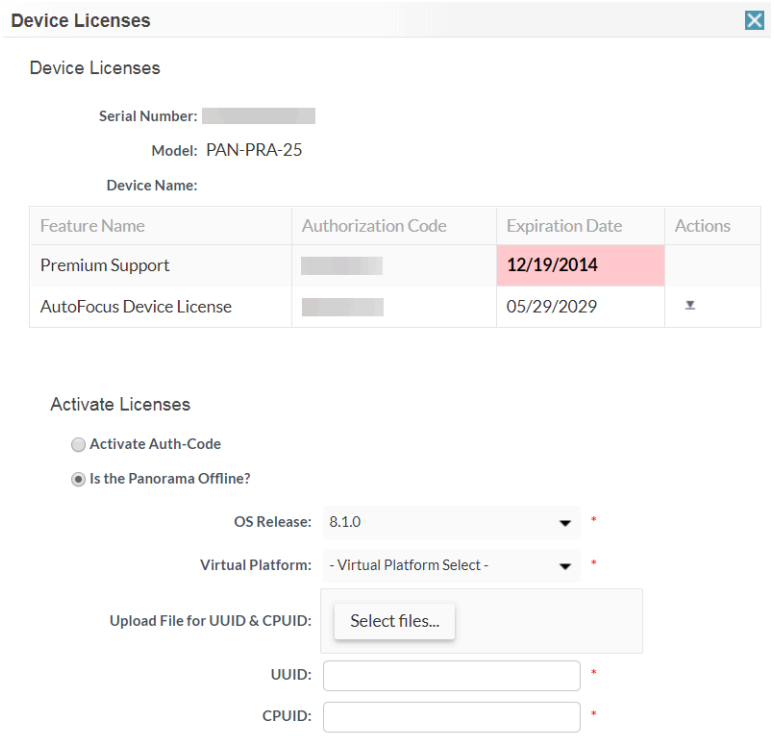
1. Select **Panorama** > **Setup** > **Management** and edit the General Settings.
2. Enter the Panorama **Serial Number** (included in the order fulfillment email) and click **OK**.
3. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Upload the Panorama virtual appliance information to the Customer Support website.

1. On the Retrieve Management License dialogue, click the **here** link to gather the UUID, CPUID, Panorama Version and Virtual Platform information. Click **Download Link** to download a XML file of the required Panorama information that can be uploaded to the Customer Support Portal.

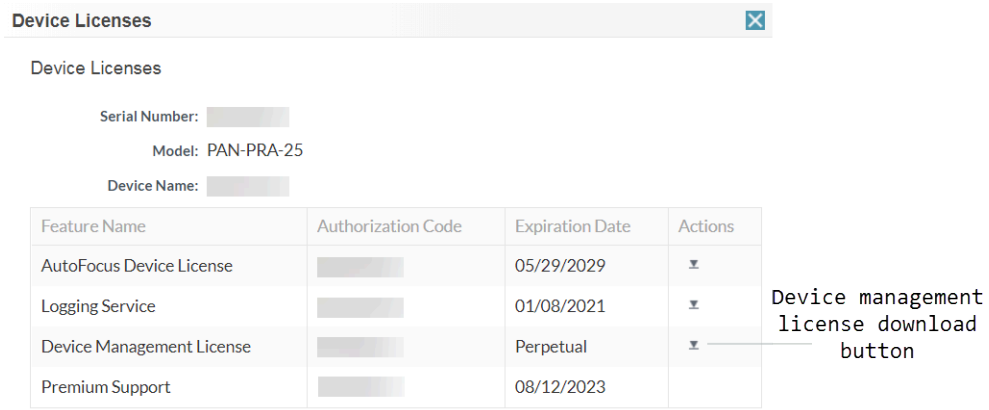
   On initial deployment, may need to log out and back in to the web interface to see the dialogue.
2. Log in to the Palo Alto Networks Customer Support web site.
3. Click **Get Support** in the upper right-hand corner.
4. Select **Assets** > **Devices**, find your Panorama virtual appliance and in the Action column, click the edit icon (✎).
5. Select **Is the Panorama Offline?** and enter the Panorama information gathered in Step 2, or click **Select files...** to upload the downloaded XML file.
6. **Agree and Submit** the EULA.



 ©2025 Palo Alto Networks, Inc.

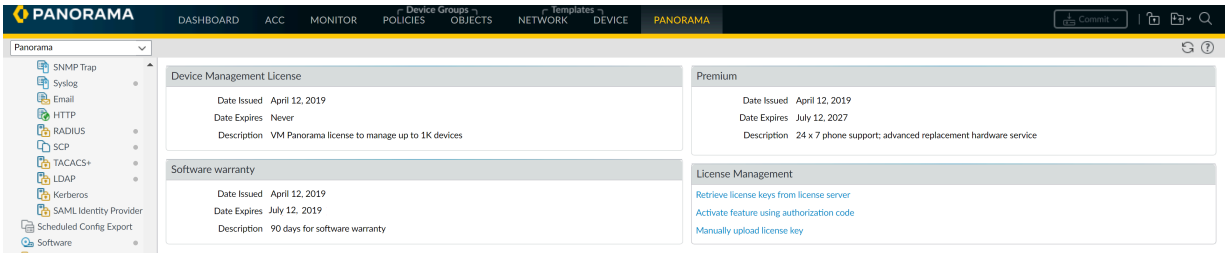
STEP 4 | Install the device management license.

1. In the Actions column, download the device management license.



2. In the Panorama web interface, click **Panorama** > **Licenses** and **Manually upload license key**.
3. Click **Choose file**, locate the downloaded device management license key and click **OK**.

STEP 5 | Confirm that the device management license was successfully uploaded by verify that the Device Management License displays with the license information.



Activate/Retrieve a Firewall Management License on the M-Series Appliance

In order to manage devices on Panorama, you need to activate a Capacity License. The Capacity License determines the number of devices Panorama can manage. Log Collectors and WildFire appliances are not treated as managed devices and do not count toward the number of devices allotted by the Capacity License.

Before activating and retrieving a Panorama firewall management license on the M-Series appliance:

- Register Panorama.
- Locate the auth-codes for the product/subscription you purchased. When you placed your order, Palo Alto Networks Customer Service sent you an email that listed the auth-code associated with the purchase. If you cannot locate this email, contact Palo Alto Networks Customer Support to obtain your codes before proceeding.

After you activate and retrieve the license, the **Panorama** > **Licenses** page displays the associated issuance date, expiration date, and the number of firewalls that the license enables Panorama to manage.

 ©2025 Palo Alto Networks, Inc.

To activate and retrieve the license, the options are:

- Use the web interface to activate and retrieve the license.

  Select this option if Panorama is ready to connect to the Palo Alto Networks update server (you completed the task Perform Initial Configuration of the M-Series Appliance) but you have not activated the license on the Palo Alto Networks Customer Support web site.

  1. Select **Panorama** > **Licenses** and click **Activate feature using authorization code**.
  2. Enter the **Authorization Code** and click **OK**. Panorama retrieves and activates the license.

- Retrieve the license key from the license server.

  If Panorama is not ready to connect to the update server (for example, you have not completed the initial M-Series appliance setup), you can activate the license on the Support website so that, when Panorama is ready to connect, you can then use the web interface to retrieve the activated license. The process of retrieving an activated license is faster than the process of both retrieving and activating.

  1. Activate the license on the Palo Alto Networks Customer Support web site.
     1. On a host with internet access, use a web browser to access the Palo Alto Networks Customer Support web site and log in.
     2. Select **Assets** > **Devices**, find your M-Series appliance and, in the Action column, click the edit icon ( ).
     3. Select **Activate Auth-Code**, enter the **Authorization Code** and click **Agree and Submit** to activate the license.
  2. Configure Panorama to connect to the update server: see Perform Initial Configuration of the M-Series Appliance.
  3. Select **Panorama** > **Licenses** and click **Retrieve license keys from the license server**. Panorama retrieves the activated license.

- Manually upload the license from a host to Panorama. Panorama must have access to that host.

  If Panorama is set up (you completed the task Perform Initial Configuration of the M-Series Appliance) but does not have a connection to the update server, activate the license on the

 ©2025 Palo Alto Networks, Inc.

Support website, download it to a host that has a connection to the update server, then upload it to Panorama.

1. Activate and download the license from the Palo Alto Networks Customer Support web site.
    1. On a host with internet access, use a web browser to access the Palo Alto Networks Customer Support web site and log in.
    2. Select **Assets** > **Devices**, find your M-Series appliance and, in the Action column, click the edit icon ( ).
    3. Select **Activate Auth-Code**, enter the **Authorization Code** and click **Agree and Submit** to activate the license.
    4. In the Action column, click the download icon and save the license key file to the host.
2. In the Panorama web interface, select **Panorama** > **Licenses**, click **Manually upload license key** and click **Browse**.
3. Select the key file you downloaded to the host and click **Open**.
4. Click **OK** to upload the activated license key.

 ©2025 Palo Alto Networks, Inc.

Install the Panorama Device Certificate

Where Can I Use This?	What Do I Need?
• NGFW (Managed by Panorama)	❑ Device management license ❑ Support license ❑ Outbound internet access ❑ Customer Support Portal (CSP) account with one of the following user roles: Super User, Standard User, Limited User, Threat Researcher, AutoFocus Trial Role, Group Super User, Group Standard User, Group Limited User, Group Threat Researcher, Authorized Support Center (ASC) User, and ASC Full Service User. ❑ Panorama superuser role

You must install the device certificate on the Panorama™ management server to use one or more cloud services. You only need to install a device certificate once. The device certificate has a 90-day lifetime. The firewall reinstalls the device certificate 15 days before the certificate expires. In the event Panorama is unable to reinstall the device certificate on its own, you may need to manually restore an expired device certificate.

To successfully install the device certificate, Panorama must have an outbound internet connection and the following Fully Qualified Domain Names (FQDN) and ports must be allowed on your network.

FQDN	Ports
• http://ocsp.paloaltonetworks.com • http://crl.paloaltonetworks.com • http://ocsp.godaddy.com	TCP 80
• https://api.paloaltonetworks.com • http://apitrusted.paloaltonetworks.com • https://certificatetrusted.paloaltonetworks.com • https://certificate.paloaltonetworks.com	TCP 443
• *.gpcloudservice.com	TCP 444 and TCP 443

 ©2025 Palo Alto Networks, Inc.

M-300 and M-700 appliances automatically install the device certificate when they first connect to the Palo Alto Networks CSP during the initial registration process. You do not need to manually install the device certificate for these M-Series appliances.

STEP 1 | Generate the One Time Password (OTP).

OTP lifetime is 60 minutes and expires if not used within the 60 minute lifetime.

Panorama may only attempt to retrieve the OTP from the CSP one time. If Panorama fails for any reason to fetch the OTP, the OTP expires and you must generate a new OTP.

1. Log in to the Customer Support Portal with a user role that has permission to generate an OTP.
2. Select **Products** > **Device Certificates** and **Generate OTP**.
3. For the **Device Type**, select **Generate OTP for Panorama** and click **Next**.
4. Select the **Panorama Device** serial number and **Generate OTP**.
5. **Generate OTP** and copy the OTP.

STEP 2 | Log in to the Panorama Web Interface as a Superuser.

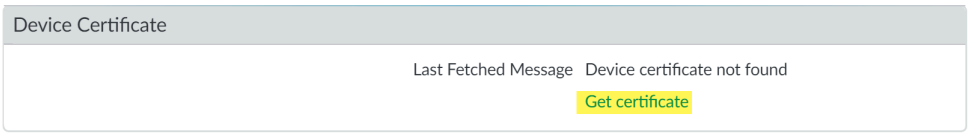
A Panorama admin with Superuser access privileges is required to required to apply the OTP used to install the device certificate on Panorama.

STEP 3 | Configure the Network Time Protocol (NTP) server.

An NTP server is required to validate the device certification expiration date, ensure the device certificate does not expire early or become invalid.

1. Select **Panorama** > **Setup** > **Services**.
2. Select **NTP** and enter the hostname or IP address of the **Primary NTP Server**.
3. (Optional) Enter a the hostname or IP address of the **Secondary NTP Server**.
4. (Optional) To authenticate time updates from the NTP server(s), for **Authentication Type**, select one of the following for each server.
   - **None** (default)—Disables NTP authentication.
   - **Symmetric Key**—Firewall uses symmetric key exchange (shared secrets) to authenticate time updates.
     - **Key ID**—Enter the Key ID (1-65534)
     - **Algorithm**—Select the algorithm to use in NTP authentication (**MDS** or **SHA1**)
5. Click **OK** to save your configuration changes.
6. Select **Commit** and **Commit to Panorama**.

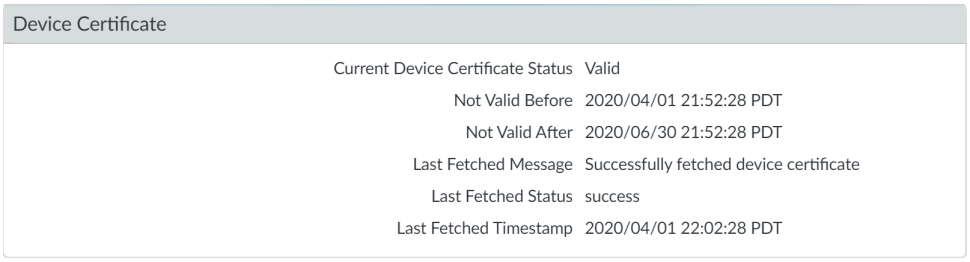
STEP 4 | Select **Panorama** > **Setup** > **Management** > **Device Certificate Settings** and **Get certificate**.



 ©2025 Palo Alto Networks, Inc.

STEP 5 | Enter the **One-time Password** you generated and click **OK**.

STEP 6 | Panorama successfully retrieves and installs the certificate.



 ©2025 Palo Alto Networks, Inc.

Install the Device Certificate for a Dedicated Log Collector

Where Can I Use This?	What Do I Need?
• NGFW (Managed by Panorama)	❑ Device management license ❑ Support license ❑ Outbound internet access ❑ Customer Support Portal (CSP) account with one of the following user roles: Super User, Standard User, Limited User, Threat Researcher, AutoFocus Trial Role, Group Super User, Group Standard User, Group Limited User, Group Threat Researcher, Authorized Support Center (ASC) User, and ASC Full Service User. ❑ Panorama superuser role

You must install the device certificate on the Dedicated Log Collector to use Device Telemetry. You only need to install a device certificate once. The device certificate has a 90-day lifetime. The Dedicated Log Collector reinstalls the device certificate 15 days before the certificate expires. In the event the Dedicated Log Collector is unable to reinstall the device certificate on its own, you may need to manually restore an expired device certificate.

To successfully install the device certificate, the Dedicated Log Collector must have an outbound internet connection and the following Fully Qualified Domain Names (FQDN) and ports must be allowed on your network.

You must manually install the device certificate on each Dedicated Log Collector individually. Installing the device certificate from the Panorama™ management server is not supported.

FQDN	Ports
• http://ocsp.paloaltonetworks.com • http://crl.paloaltonetworks.com • http://ocsp.godaddy.com	TCP 80
• https://api.paloaltonetworks.com • http://apitrusted.paloaltonetworks.com • https://certificatetrusted.paloaltonetworks.com • https://certificate.paloaltonetworks.com	TCP 443

 ©2025 Palo Alto Networks, Inc.

FQDN	Ports
• *.gpcloudservice.com	TCP 444 and TCP 443

M-300 and M-700 appliances automatically install the device certificate when they first connect to the Palo Alto Networks Customer Support Portal (CSP) during the initial registration process. You do not need to manually install the device certificate for these M-Series appliances.

STEP 1 | Log in to the Dedicated Log Collector CLI as a Superuser.

An admin with Superuser access privileges is required to required to apply the OTP used to install the device certificate on Panorama.

STEP 2 | View the current device certificate status on the Dedicated Log Collector.

```
admin>show device-certificate status
```

The Dedicated Log Collector displays one of the following responses:

- **Device certificate was never installed**—`No device certificate found`
- **Device certificate expired**—`Current device certificate status: Expired`

  The response also displays the lifetime of the previous device certificate and the date and time the last device certificate fetch was attempted.
- **Device certificate fetch failed**—Response displays the last time the device certificate fetch was attempted.

STEP 3 | Generate the One Time Password (OTP).

OTP lifetime is 60 minutes and expires if not used within the 60 minute lifetime.

The Dedicated Log Collector may only attempt to retrieve the OTP from the CSP one time. If the Dedicated Log Collector fails for any reason to fetch the OTP, the OTP expires and you must generate a new OTP.

1. Log in to the Customer Support Portal with a user role that has permission to generate an OTP.
2. Select **Products** > **Device Certificates** and **Generate OTP**.
3. For the **Device Type**, select **Generate OTP for Panorama** and click **Next**.
4. Select the **Panorama Device** serial number and **Generate OTP**.
5. **Generate OTP** and copy the OTP.

STEP 4 | Log in to the Dedicated Log Collector CLI.

An admin with Superuser access privileges is required to required to apply the OTP used to install the device certificate on Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Configure the Network Time Protocol (NTP) server.

An NTP server is required to validate the device certification expiration date, ensure the device certificate does not expire early or become invalid.

```
admin>configure
```

```
admin#set deviceconfig system ntp-servers primary-ntp-server ntp-server-address <ip_address>
```

```
admin#set deviceconfig system ntp-servers secondary-ntp-server ntp-server-address <ip_address>
```

```
admin>commit
```

```
admin>exit
```

STEP 6 | Install the device certificate.

```
admin>request certificate fetch otp <otp_value>
```

STEP 7 | Verify the device certificate successfully installed.

```
admin> show device-certificate status
```

A successful device certificate installation displays the following response:

```
Device Certificate information:
Current device certificate status: Valid
Not valid before: 2022/11/30 15:17:47 PST
Not valid after: 2023/02/28 15:17:47 PST
Last fetched timestamp: 2022/11/30 15:29:42 PST
Last fetched status: success
Last fetched info: Successfully fetched Device Certificate
```

 ©2025 Palo Alto Networks, Inc.

Transition to a Different Panorama Model

When your network requirements change (for example, the logging rate increases), you can migrate the Panorama management server and Dedicated Log Collectors to Panorama Models that better support those requirements.

Transitioning to a different Panorama models requires you import the Panorama configuration from the old Panorama to the new Panorama. Before you begin your transition, ensure that the old and new Panorama are in the same Panorama mode (Management Only or Panorama mode). See Panorama Models *for more information about the Panorama modes.*

This is required to successfully import the Panorama configuration to the new Panorama. For more information on changing the Panorama mode, see the M-Series Setup Overview *M-Series appliances and* Set Up the Panorama Virtual Appliance *for Panorama virtual appliances*

- Migrate from a Panorama Virtual Appliance to an M-Series Appliance
- Migrate a Panorama Virtual Appliance to a Different Hypervisor
- Migrate from an M-Series Appliance to a Panorama Virtual Appliance
- Migrate from an M-500 Appliance to an M-700 Appliance
- Migrate from an M-600 Appliance to an M-700 Appliance
- Migrate from an M-100 Appliance to an M-500 Appliance
- Migrate from an M-100 or M-500 Appliance to an M-200 or M-600 Appliance

Migrate from a Panorama Virtual Appliance to an M-Series Appliance

You can migrate the Panorama configuration from a Panorama virtual appliance to an M-Series appliance in Panorama mode. However, you cannot migrate the logs because the log format on the Panorama virtual appliance is incompatible with that on M-Series appliances. Therefore, if you want to maintain access to the old logs stored on the Panorama virtual appliance, you must continue running the Panorama virtual appliance after the migration. The M-Series appliance will collect the new logs that firewalls forward after the migration. After the pre-migration logs expire or become irrelevant due to aging, you can shut down the Panorama virtual appliance.

Legacy mode is no longer supported in PAN-OS 8.1 or later releases. If the old Panorama virtual appliance is in Legacy mode, you must change Panorama to Panorama mode before migrating to the new hypervisor in order to preserve the log settings and Log Collector forwarding configurations. Importing the configuration of the old Panorama in Legacy mode to a new Panorama in Panorama mode causes all log and log forwarding settings to be removed.

You cannot migrate logs between hypervisors. Therefore, if you want to maintain access to the logs stored on the old Panorama virtual appliance, you must continue running the old Panorama virtual appliance after the migration and add it as a managed Log Collector on the new Panorama virtual appliance. This allows the new Panorama virtual appliance to collect the new logs that firewalls forward after the migration, while maintaining access to the old log data. After the pre-

 ©2025 Palo Alto Networks, Inc.

migration logs expire or become irrelevant due to aging, you can shut down the Panorama virtual appliance.

If you store firewall logs on Dedicated Log Collectors (M-Series appliances in Log Collector mode) instead of on the Panorama virtual appliance, you can maintain access to the logs by migrating the Dedicated Log Collectors to the M-Series appliance in Panorama mode.

Policy rule usage data is not preserved when you transition to a different Panorama model. This means that all existing policy rule usage data from the old Panorama is no longer displayed after a successful migration to a new Panorama model. After a successful migration, Panorama begins tracking policy rule usage data based on the date the migration was completed. For example, the `Created` date displays the date the migration was completed.

STEP 1 | Plan the migration.

- Upgrade the software on the Panorama virtual appliance before the migration if the M-Series appliance requires a later release of the current software (the M-500 appliance requires Panorama 7.0 or a later release. The M-600 and M-200 appliances require Panorama 8.1 or later release. The M-700 and M-300 require Panorama 10.2 or later release.). For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
- Schedule a maintenance window for the migration. Although firewalls can buffer logs after the Panorama virtual appliance goes offline and then forward the logs after the M-Series appliance comes online, completing the migration during a maintenance window minimizes the risk that logs will exceed the buffer capacities and be lost during the transition between Panorama models.
- Consider whether to maintain access to the Panorama virtual appliance after the migration to access existing logs. The most efficient approach is to assign a new IP address to the Panorama virtual appliance and reuse its old IP address for the M-Series appliance. This ensures that the Panorama virtual appliance remains accessible and that firewalls can point to the M-Series appliance without you reconfiguring the Panorama IP address on each firewall.

STEP 2 | Purchase the new M-Series appliance, and migrate your subscriptions to the new appliance.

1. Purchase the new M-Series appliance.
2. Purchase the new support license and migration license.
3. At the time you purchase the new M-Series appliance, provide your sales representative the serial number and device management auth-code of the Panorama virtual appliance you are phasing out, as well as a license migration date of your choosing. On receipt of your M-Series appliance, register the appliance and activate the device management and support licenses using the migration and support auth-codes provided by Palo Alto Networks. On the migration date, the device management license on the Panorama virtual appliance is decommissioned, and you can no longer manage devices or collect logs using the Panorama virtual appliance. However, the support license is preserved and the Panorama appliance remains under support. You can complete the migration after the effective date, but you are unable to commit any configuration changes on the now decommissioned Panorama virtual appliance.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | (Legacy mode only) On the old Panorama virtual appliance, change to Panorama mode .

> *This step is required to preserve the log data, settings and log forwarding configuration of the Panorama virtual appliance. If you export the Panorama configuration while in Legacy mode, these settings are lost. You must complete Step 9 if you do not change Panorama to Panorama mode before continuing.*
>
> *Continue to the next step if the Panorama virtual appliance is already in Panorama or Management Only mode.*

STEP 4 | Export the Panorama configuration from the Panorama virtual appliance.

1. Log in to the Panorama virtual appliance and select **Panorama** > **Setup** > **Operations**.
2. Click **Save named Panorama configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**.
3. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**. Panorama exports the configuration to your client system as an XML file.

STEP 5 | Power off the Panorama virtual appliance if you won't need to access to it after the migration or assign a new IP address to its management (MGT) interface if you will need access to it.

To power off the Panorama virtual appliance, see the documentation for your VMware product.

To change the IP address on the Panorama virtual appliance:

1. Select **Panorama** > **Setup** > **Management**, and edit the Management Interface Settings.
2. Enter the new **IP Address** and click **OK**.
3. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 6 | Perform the initial setup of the M-Series appliance.

1. Rack mount the M-Series appliance. Refer to the M-Series Appliance Hardware Reference Guide for instructions.
2. Perform Initial Configuration of the M-Series Appliance to define the network connections required to activate licenses and install updates.
3. Register Panorama.
4. Activate a Panorama Support License.
5. For Panorama-managed Prisma Access, ensure that you transfer the licenses to the new panorama appliance.
6. Activate/Retrieve a Firewall Management License on the M-Series Appliance. Use the auth-code associated with the migration license.
7. Install Content and Software Updates for Panorama. Install the same versions as those on the Panorama virtual appliance.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Load the Panorama configuration snapshot that you exported from the Panorama virtual appliance into the M-Series appliance.

*The Panorama **Policy** rule **Creation** and **Modified** dates are updated to reflect the date you commit the imported Panorama configuration on the new Panorama. The universially unique identifier (UUID) for each policy rule persists when you migrate the Panorama configuration.*

*The **Creation** and **Modified** for managed firewalls are not impacted when you monitor policy rule usage for a managed firewall because this data is stored locally on the managed firewall and not on Panorama.*

1. On the M-Series appliance, select **Panorama** > **Setup** > **Operations**.
2. Click **Import named Panorama configuration snapshot**, **Browse** to the Panorama configuration file you exported from the Panorama virtual appliance, and click **OK**.
3. Click **Load named Panorama configuration snapshot**, select the **Name** of the configuration you just imported, select a **Decryption Key** (the master key for Panorama), and click **OK**. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file.
4. If errors occurred, save them to a local file. Resolve each error to ensure the migrated configuration is valid.

STEP 8 | Modify the configuration on the M-Series appliance.

Required if the M-Series appliance will use different values than the Panorama virtual appliance. If you will maintain access to the Panorama virtual appliance to access its logs, use a different hostname and IP address for the M-Series appliance.

1. Select **Panorama** > **Setup** > **Management**.
2. Edit the General Settings, modify the **Hostname**, and click **OK**.
3. Edit the Management Interface Settings, modify the values as necessary, and click **OK**.

STEP 9 | Add the default managed collector and Collector Group back to the M-Series appliance.

Loading the configuration from the Panorama virtual appliance (Step 7) removes the default managed collector and Collector Group that are predefined on each M-Series appliance.

1. Configure a Managed Collector that is local to the M-Series appliance.
2. Configure a Collector Group for the default managed collector.
3. Select **Commit** > **Commit to Panorama** and **Commit** your changes to the Panorama configuration.

STEP 10 | Recover Managed Device Connectivity to Panorama for managed firewalls and Dedicated Log Collectors added using the device registration authentication key.

This is required when transitioning from one Panorama model to another.

 ©2025 Palo Alto Networks, Inc.

STEP 11 | Synchronize the M-Series appliance with the firewalls to resume firewall management.

Complete this step during a maintenance window to minimize network disruption.

1. On the M-Series appliance, select **Panorama** > **Managed Devices** and verify that the Device State column displays **Connected** for the firewalls.

   At this point, the Shared Policy (device groups) and Template columns display **Out of sync** for the firewalls.
2. Push your changes to device groups and templates:
   1. Select **Commit** > **Push to Devices** and **Edit Selections**.
   2. Select **Device Groups**, select every device group, **Include Device and Network Templates**, and click **OK**.
   3. **Push** your changes.
3. In the **Panorama** > **Managed Devices** page, verify that the Shared Policy and Template columns display **In sync** for the firewalls.

After you migrate to a different Panorama model, if there are connectivity issues between Panorama and the managed firewalls, recover the connectivity of the managed devices to Panorama *to resolve the issues.*

Migrate a Panorama Virtual Appliance to a Different Hypervisor

Migrate the Panorama configuration of a Panorama virtual appliance from one supported hypervisor to another supported hypervisor in Management Only mode or Panorama mode. Before migrating to the Panorama virtual appliance to a new hypervisor, review the Panorama Models to ensure that the new hypervisor you are migrating to is supported. Additionally, if your Panorama configuration has multiple interfaces configuration for device management includes multiple interfaces for device management, log collection, Collector Group communication, licensing and software updates, review Setup Prerequisites for the Panorama Virtual Appliance to verify that the hypervisor you are migrating to supports multiple interfaces.

Legacy mode is no longer supported in PAN-OS 8.1 or later releases. If the old Panorama virtual appliance is in Legacy mode, you must change Panorama to Panorama mode before migrating to the new hypervisor in order to preserve the log settings and Log Collector forwarding configurations. Importing the configuration of the old Panorama in Legacy mode to a new Panorama in Panorama mode causes all log and log forwarding settings to be removed.

You cannot migrate logs from Panorama virtual appliance. Therefore, if you want to maintain access to the logs stored on the old Panorama virtual appliance, you must continue running the old Panorama virtual appliance in Log Collector mode after the migration and add it as a managed Log Collector on the new Panorama virtual appliance. This allows the new Panorama virtual appliance to collect the new logs that firewalls forward after the migration, while maintaining access to the old log data. After the pre-migration logs expire or become irrelevant due to aging, you can shut down the Panorama virtual appliance.

 ©2025 Palo Alto Networks, Inc.

If you store firewall logs on Dedicated Log Collectors (Panorama virtual appliance in Log Collector mode) instead of on the Panorama virtual appliance, you can maintain access to the logs by migrating the Dedicated Log Collectors to the new Panorama virtual appliance in Panorama mode.

Policy rule usage data *is not preserved when you transition to a different Panorama model. This means that all existing policy rule usage data from the old Panorama is no longer displayed after a successful migration to a new Panorama model. After a successful migration, Panorama begins tracking policy rule usage data based on the date the migration was completed. For example, the* `Created` *date displays the date the migration was completed.*

STEP 1 | Plan the migration.

- Upgrade the software on the Panorama virtual appliance before the migration if the new Panorama virtual appliance requires a later release of the current software. For the minimum PAN-OS version for each hypervisor, see Panorama Hypervisor Support. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
- Schedule a maintenance window for the migration. Although firewalls can buffer logs after the Panorama virtual appliance goes offline and then forward the logs after the new Panorama virtual appliance comes online, completing the migration during a maintenance window minimizes the risk that logs will exceed the buffer capacities and be lost during the transition between hypervisors.
- Consider whether to maintain access to the old Panorama virtual appliance after the migration to access existing logs. The most efficient approach is to assign a new IP address to the old Panorama virtual appliance and reuse its old IP address for the Panorama virtual appliance. This ensures that the old Panorama virtual appliance remains accessible and that firewalls can point to the new Panorama virtual appliance without you reconfiguring the Panorama IP address on each firewall.

  If you to intend to maintain access to the old Panorama virtual appliance, you must purchase a new device management license and support license for the new Panorama virtual appliance before you can complete the migration successfully.

STEP 2 | (Legacy mode only) On the old Panorama virtual appliance, Set up a Panorama Virtual Appliance in Panorama Mode.

*This step is required to preserve the log settings (**Panorama** > **Log Settings**) on the old Panorama virtual appliance. If you export the Panorama configuration while in Legacy mode, these settings are lost.*

Continue to the next step if the Panorama virtual appliance is already in Panorama or Management Only mode.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Export the Panorama configuration from the old Panorama virtual appliance.

1. Log in to the Panorama Web Interface.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Export named Panorama configuration snapshot**, select `running-config.xml` and click **OK**. Panorama exports the configuration to your client system as an XML file.
4. Locate the `running-config.xml` file you exported and rename the XML file. This is required to import the configuration as Panorama does not support importing an XML file with the name `running-config.xml`.

STEP 4 | Install the Panorama virtual appliance.

STEP 5 | Migrate the serial number of the old Panorama virtual appliance to the new Panorama virtual appliance.

This step is required to migrate all subscriptions and the device management license tied to the Panorama serial number and only if you intend to shut down the old Panorama virtual appliance. If you do intend on maintaining access to the old Panorama virtual appliance, continue to the next step.

You have up to 90 days to shut down the old Panorama virtual appliance. Running multiple Panorama virtual appliances with the same serial number violates the EULA.

1. Log in to the Panorama web interface of the old Panorama virtual appliance.
2. In the **Dashboard**, copy the `Serial #` of the old Panorama virtual appliance located in the General Information widget.
3. Log in to the Panorama web interface of the new Panorama virtual appliance.
4. Add the serial number of the old Panorama virtual appliance to the new Panorama virtual appliance.
   1. Select **Panorama** > **Setup** > **Management** and edit the General Settings.
   2. Enter (paste) the **Serial Number** and click **OK**.
   3. Select **Commit** and **Commit to Panorama**.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Perform the initial setup of the new Panorama virtual appliance.

1. Perform Initial Configuration of the Panorama Virtual Appliance to define the network connections required to activate licenses and install updates.
2. (For maintaining access to the old Panorama virtual appliance only) Register Panorama.
3. (For maintaining access to the old Panorama virtual appliance only) Activate a Panorama Support License.
4. For Panorama-managed Prisma Access, ensure that you transfer the licenses to the new panorama appliance.
5. (For maintaining access to the old Panorama virtual appliance only) Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected. Use the auth-code associated with the migration license.
6. Install Content and Software Updates for Panorama. Install the same versions as those on the old Panorama virtual appliance.

   This step is required before loading configuration from the old Panorama virtual appliance. Ensure that all required content updates are installed to avoid security outages.
7. Select **Panorama** > **Plugins** and install all plugins that were installed on the old Panorama virtual appliance.

STEP 7 | Power off the old Panorama virtual appliance if you won't need to access to it after the migration or assign a new IP address to its management (MGT) interface if you will need access to it.

To power off the Panorama virtual appliance, see the supported documentation for the hypervisor on which the old Panorama virtual appliance has been deployed.

To change the IP address on the Panorama virtual appliance:

1. On the web interface of the old Panorama virtual appliance, select **Panorama** > **Setup** > **Management**, and edit the Management Interface Settings.
2. Enter the new **IP Address** and click **OK**.
3. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 8 | (Prisma Access) Transfer the Prisma Access license from the old Panorama virtual appliance to the new Panorama virtual appliance.

 ©2025 Palo Alto Networks, Inc.

STEP 9 | Load the Panorama configuration snapshot that you exported from the old Panorama virtual appliance into the new Panorama virtual appliance.

> *The Panorama **Policy** rule **Creation** and **Modified** dates are updated to reflect the date you commit the imported Panorama configuration on the new Panorama. The universially unique identifier (UUID) for each policy rule persists when you migrate the Panorama configuration.*
>
> *The **Creation** and **Modified** for managed firewalls are not impacted when you monitor policy rule usage for a managed firewall because this data is stored locally on the managed firewall and not on Panorama.*

1. Log in to the Panorama Web Interface of the new Panorama virtual appliance.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Import named Panorama configuration snapshot**, **Browse** to the Panorama configuration file you exported from the Panorama virtual appliance, and click **OK**.
4. Click **Load named Panorama configuration snapshot**, select the **Name** of the configuration you just imported, leave the **Decryption Key** blank (empty), and click **OK**. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file.
5. If errors occurred, save them to a local file. Resolve each error to ensure the migrated configuration is valid.

STEP 10 | Modify the configuration on the new Panorama virtual appliance.

Required if the new Panorama virtual appliance will use different values than the old Panorama virtual appliance. If you will maintain access to the old Panorama virtual appliance to access its logs, use a different hostname and IP address for the new Panorama virtual appliance.

1. Select **Panorama** > **Setup** > **Management**.
2. Edit the General Settings, modify the **Hostname**, and click **OK**.
3. Edit the Management Interface Settings, modify the values as necessary, and click **OK**.
4. Select **Commit** > **Commit to Panorama** and **Commit** your changes to the Panorama configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 11 | Add the default managed collector and Collector Group to the new Panorama virtual appliance.

Loading the configuration from the old Panorama virtual appliance (Step 7) removes the default managed collector and Collector Group that are predefined on each Panorama virtual appliance in Panorama mode.

1. To maintain access to logs stored on the old Panorama virtual appliance, change to Log Collector mode and add the Dedicated Log Collector to the new Panorama virtual appliance.
   1. Set Up The Panorama Virtual Appliance as a Log Collector.
   2. Configure a Managed Collector.
2. Configure a Managed Collector that is local to the Panorama virtual appliance.
3. Configure a Collector Group for the default managed collector.
4. Select **Commit** > **Commit to Panorama** and **Commit** your changes to the Panorama configuration.

STEP 12 | Recover Managed Device Connectivity to Panorama for managed firewalls and Dedicated Log Collectors added using the device registration authentication key.

This is required when transitioning from one Panorama model to another.

STEP 13 | Synchronize the new Panorama virtual appliance with the firewalls to resume firewall management.

Complete this step during a maintenance window to minimize network disruption.

1. On the new Panorama virtual appliance, select **Panorama** > **Managed Devices** and verify that the Device State column displays **Connected** for the firewalls.

   At this point, the Shared Policy (device groups) and Template columns display **Out of sync** for the firewalls.
2. Push your changes to device groups and templates:
   1. Select **Commit** > **Push to Devices** and **Edit Selections**.
   2. Select **Device Groups**, select every device group, **Include Device and Network Templates**, and click **OK**.
   3. **Push** your changes.
3. In the **Panorama** > **Managed Devices** page, verify that the Shared Policy and Template columns display **In sync** for the firewalls.

After you migrate to a different Panorama model, if there are connectivity issues between Panorama and the managed firewalls, recover the connectivity of the managed devices to Panorama *to resolve the issues.*

 ©2025 Palo Alto Networks, Inc.

Migrate from an M-Series Appliance to a Panorama Virtual Appliance

You can migrate the Panorama configuration from an M-100, M-200, M-300, M-500, M-600, M-700 appliance to a Panorama virtual appliance in Panorama mode. However, you cannot migrate the logs because the log format on the M-Series appliances is incompatible with that on the Panorama virtual appliances. Therefore, if you want to maintain access to the old logs stored on the M-Series appliance, you must continue running the M-Series appliance as a Dedicated Log Collector after the migration and add it to the Panorama virtual appliance as a managed collector.

If your Panorama management server is part of a high availability configuration, you must deploy a second Panorama virtual appliance of the same hypervisor or cloud environment, and purchase the required device management and support licenses. See Panorama HA Prerequisites for a full list of HA requirements.

Policy rule usage data *is not preserved when you transition to a different Panorama model. This means that all existing policy rule usage data from the old Panorama is no longer displayed after a successful migration to a new Panorama model. After a successful migration, Panorama begins tracking policy rule usage data based on the date the migration was completed. For example, the* `Created` *date displays the date the migration was completed.*

STEP 1 | Plan the migration.

- Upgrade the M-Series appliance to PAN-OS 10.2 or later release before the migrating to the Panorama virtual appliance. To upgrade Panorama, see Install Content and Software Updates for Panorama. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
- Schedule a maintenance window for the migration. Although firewalls can buffer logs after the M-Series appliance goes offline and then forward the logs after the Panorama virtual appliance comes online, completing the migration during a maintenance window minimizes the risk that logs will exceed the buffer capacities during the transition to a different Panorama model.

STEP 2 | Purchase management and support licenses for the new Panorama virtual appliance.

1. Contact your sales representative to purchase the new device management and support licenses.
2. Provide your sales representative the serial number of the M-Series appliance you to plan phase out, the serial number and support auth code you received when you purchased the new Panorama virtual appliance, and the date when you expect your migration from the old device to the new virtual appliance to be completed. Before the migration date, register the serial number and activate support auth code on the new virtual appliance so that you can begin your migration. The capacity auth code on the old M-Series appliance is automatically removed on the expected migration completion date you provided.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Perform the initial setup of the Panorama virtual appliance.

1. Set Up the Panorama Virtual Appliance.
2. Perform Initial Configuration of the Panorama Virtual Appliance to define the network connections required to activate licenses and install updates.
3. Register Panorama.
4. Activate a Panorama Support License.
5. For Panorama-managed Prisma Access, ensure that you transfer the licenses to the new panorama appliance.
6. Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected
7. Install Content and Software Updates for Panorama. Install the same versions as those on the M-Series appliance.

STEP 4 | Edit the M-Series appliance Panorama interface configuration to only use the management interface.

The Panorama virtual appliance supports only the management interface for device management and log collection.

1. Log in to the Panorama Web Interface of the M-Series appliance.
2. Select **Panorama** > **Setup** > **Management**.
3. Edit the General Settings, modify the **Hostname**, and click **OK**.
4. Select **Interfaces** and edit the **Management** interface to enable the required services.
5. Disable services for the remaining interfaces.
6. Select **Commit** > **Commit to Panorama**.

STEP 5 | Add the IP address of the new Panorama virtual appliance.

On the M-Series appliance, add the Public IP address of the Panorama virtual appliance as the second Panorama Server to manage devices from the new Panorama management server. If the Panorama virtual appliance is deployed on Alibaba Cloud, AWS, Azure, GCP, or OCI, use the public IP address.

1. Select **Device** > **Setup**.
2. In the Template context drop-down, select the template or template stack containing the Panorama server configuration.
3. Edit the Panorama Settings.
4. Enter the Panorama virtual appliance public IP address and click **OK**.
5. Select **Commit** > **Commit and Push**.

STEP 6 | Export the configuration from the M-Series appliance.

1. Select **Panorama** > **Setup** > **Operations**.
2. Click **Save named Panorama configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**.
3. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**. Panorama exports the configuration to

 ©2025 Palo Alto Networks, Inc.

your client system as an XML file. Save the configuration to a location external to the Panorama appliance.

STEP 7 | Power off the M-Series appliance or assign a new IP address to the management (MGT) interface.

If the M-Series appliance is in Panorama mode and has logs stored on the local Log Collector that you need access on the new Panorama virtual appliance, you must change the IP address on the M-Series appliance in order to add it to the Panorama virtual appliance as a managed Log Collector.

- **To Power off the M-Series appliance:**

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Setup** > **Operations**, and under Device Operations, **Shutdown Panorama**. Click **Yes** to confirm the shutdown.

- **To change the IP address on the M-Series appliance:**

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Setup** > **Management**, and edit the Management Interface Settings.
3. Enter the new **IP Address** and click **OK**.
4. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 8 | Load the Panorama configuration snapshot that you exported from the M-Series appliance into the Panorama virtual appliance.

*The Panorama **Policy** rule **Creation** and **Modified** dates are updated to reflect the date you commit the imported Panorama configuration on the new Panorama. The universially unique identifier (UUID) for each policy rule persists when you migrate the Panorama configuration.*

*The **Creation** and **Modified** for managed firewalls are not impacted when you monitor policy rule usage for a managed firewall because this data is stored locally on the managed firewall and not on Panorama.*

1. Log in to the Panorama web interface of the Panorama virtual appliance, and select **Panorama** > **Setup** > **Operations**.
2. Click **Import named Panorama configuration snapshot**, **Browse** to the Panorama configuration file you exported from the M-Series appliance, and click **OK**.
3. Click **Load named Panorama configuration snapshot**, select the **Name** of the configuration you just imported, select a **Decryption Key** (the master key for Panorama), and click **OK**. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file.

If errors occurred, save them to a local file. Resolve each error to ensure the migrated configuration is valid. The configuration has been loaded once the commit is successful.

 ©2025 Palo Alto Networks, Inc.

STEP 9 | Change the M-Series appliance to Log Collector mode to preserve existing log data.

Logging data is erased if you change to Log Collector mode while the logging disks are still inserted in the M-Series appliance. Logging disks must be removed before changing mode to avoid log data loss.

Generating the metadata for each disk pair rebuilds the indexes. Therefore, depending on the data size, this process can take a long time to complete. To expedite the process, you can launch multiple CLI sessions and run the metadata regeneration command in each session to complete the process simultaneously for every pair. For details, see Regenerate Metadata for M-Series Appliance RAID Pairs.

1. Remove the RAID disks from the old M-Series appliance.
   1. Power off the M-Series appliance by pressing the Power button until the system shuts down.
   2. Remove the disk pairs. For details, refer to the disk replacement procedure in the M-Series Appliance Hardware Reference Guides.
2. Power on the M-Series appliance by pressing the Power button.
3. Configure an `admin` superuser administrator account.

   If an `admin` administrator account already is already created, continue to the next step.

   An `admin` account with superuser privileges must be created before you switch to Log Collector mode or you lose access to the M-Series appliance after switching modes.
4. Log in to the Panorama CLI on the old M-Series appliance.
5. Switch from Panorama mode to Log Collector mode.
   - Switch to Log Collector mode by entering the following command:

     ```
     > request system system-mode logger
     ```
   - Enter **Y** to confirm the mode change. The M-Series appliance reboots. If the reboot process terminates your terminal emulation software session, reconnect to the M-Series appliance to see the Panorama login prompt.

     *If you see a **`CMS Login`** prompt, this means the Log Collector has not finished rebooting. Press Enter at the prompt without typing a username or password.*
   - Log back in to the CLI.
   - Verify that the switch to Log Collector mode succeeded:

     ```
     > show system info | match system-mode
     ```

     If the mode change succeeded, the output displays:

 ©2025 Palo Alto Networks, Inc.

```
> system-mode: logger
```

6. Insert the disks back into the old M-Series appliance. For details, refer to the disk replacement procedure in the M-Series Appliance Hardware Reference Guides.

   You must maintain the disk pair association. Although you can place a disk pair from slot A1/A2 on the into slot B1/B2, you must keep the disks together in the same slot; otherwise, Panorama might not restore the data successfully.
7. Enable the disk pairs by running the following CLI command for each pair:

   ```
   > request system raid add <slot> force no-format
   ```

   For example:

   ```
   > request system raid add A1 force no-format
   > request system raid add A2 force no-format
   ```

   The **force** and **no-format** arguments are required. The **force** argument associates the disk pair with the new appliance. The **no-format** argument prevents reformatting of the drives and retains the logs stored on the disks.
8. Generate the metadata for each disk pair.

   ```
   > request metadata-regenerate slot <slot_number>
   ```

   For example:

   ```
   > request metadata-regenerate slot 1
   ```
9. Enable connectivity between the Log Collector and Panorama management server.

   Enter the following commands at the Log Collector CLI, where *<IPaddress1>* is for the MGT interface of the solitary (non-HA) or active (HA) Panorama and *<IPaddress2>* is for the MGT interface of the passive (HA) Panorama, if applicable.

   ```
   > configure
   # set deviceconfig system panorama-server <IPaddress1>
    panorama-server-2 <IPaddress2>
   # commit
   # exit
   ```

STEP 10 | Recover Managed Device Connectivity to Panorama for managed firewalls and Dedicated Log Collectors added using the device registration authentication key.

This is required when transitioning from one Panorama model to another.

 ©2025 Palo Alto Networks, Inc.

STEP 11 | Synchronize the Panorama virtual appliance with the firewalls to resume firewall management.

Complete this step during a maintenance window to minimize network disruption.

1. On the Panorama virtual appliance, select **Panorama** > **Managed Devices** and verify that the Device State column displays the firewalls as **Connected**.

   At this point, the Shared Policy (device groups) and Template columns display **Out of sync** for the firewalls.
2. Push your changes to device groups and templates:
   1. Select **Commit** > **Push to Devices** and **Edit Selections**.
   2. Select **Device Groups**, select every device group, and **Include Device and Network Templates**.
   3. Select **Collector Groups**, select every collector group, and click **OK**.
   4. **Push** your changes.
3. In the **Panorama** > **Managed Devices** page, verify that the Shared Policy and Template columns display **In sync** for the firewalls.

STEP 12 | (HA only) Set up the Panorama HA peer.

If the Panorama management servers are in a high availability configuration, perform the steps below on the HA peer.

1. Perform the initial setup of the Panorama virtual appliance.
2. Edit the M-Series appliance Panorama interface configuration to only use the management interface.
3. Add the IP address of the new Panorama virtual appliance.
4. Power off the M-Series appliance or assign a new IP address to the management (MGT) interface.
5. Change the M-Series appliance to Log Collector mode to preserve existing log data.

STEP 13 | (HA only) Modify the Panorama virtual appliance HA peer configuration.

1. On an HA peer, Log in to the Panorama Web Interface, select **Panorama** > **High Availability** and edit the **Setup**.
2. In the **Peer HA IP Address** field, enter the new IP address of the HA peer and click **OK**.
3. Select **Commit** > **Commit to Panorama** and **Commit** your change
4. Repeat these steps on the other peer in the HA peer.

 ©2025 Palo Alto Networks, Inc.

STEP 14 | (HA only) Synchronize the Panorama peers.

1. Access the **Dashboard** on one of the HA peers and select **Widgets** > **System** > **High Availability** to display the HA widget.
2. **Sync to peer**, click **Yes**, and wait for the **Running Config** to display **Synchronized**.
3. Access the **Dashboard** on the remaining HA peer and select **Widgets** > **System** > **High Availability** to display the HA widget.
4. Verify that the **Running Config** displays **Synchronized**.

Migrate from an M-500 Appliance to an M-700 Appliance

You can migrate the Panorama configurations, managed firewalls, and log collectors from an M-500 appliance to an M-700 appliance. You can migrate Panorama configurations between the appliances when both the appliances are running the same PAN-OS version. However, the M-500 appliance supports up to PAN-OS version 10.1, while the M-700 appliance requires at least PAN-OS version 10.2.

To migrate the Panorama configurations across appliances with different PAN-OS versions, you must use an intermediate virtual appliance that supports both versions, and perform the migration in the following two phases:

- First, migrate the configurations from the M-500 appliance to the intermediate Panorama virtual appliance. For more information about migrating an M-Series appliance to a Panorama virtual appliance see, Migrate from an M-Series Appliance to a Panorama Virtual Appliance.
- Next, upgrade the intermediate Panorama virtual appliance to a preferred PAN-OS version, and migrate the configurations from the intermediate Panorama virtual appliance to the M-700 appliance running the same preferred PAN-OS version. For more information about migrating a Panorama virtual appliance to an M-Series appliance, see Migrate from a Panorama Virtual Appliance to an M-Series Appliance.

Ensure that all the Log Collectors in the Collector Group are the same Panorama model. For example, if you want to add the local Log Collector on the new M-700 appliance to a Collector Group, the target Collector Group must contain only M-700 appliances. The same is true for the local Log Collector for an M-700 appliance.



This procedure assumes you are no longer using the M-500 appliance for device management or log collection. If you intend to continue using the M-500 appliance as a log collector, you must get a device management license for the M-500 appliance. Without a device management license, you cannot use the M-500 appliance as a log collector.

If you do not plan to use the M-500 appliance as a log collector, but the M-500 appliance contains log data that you must access at a later date, use the Panorama web interface to query and generate reports using the existing log data. Palo Alto Networks recommends reviewing the log retention policy before decommissioning the M-500 appliance.

 ©2025 Palo Alto Networks, Inc.

Policy rule usage data *is not preserved when you migrate to a different Panorama model. This indicates that all the existing policy rule usage data from the old Panorama model is no longer displayed after you migrate to a new Panorama model. After a successful migration, Panorama begins tracking policy rule usage data based on the date the migration was completed. For example, the* `Created` *date displays the date the migration was completed.*

STEP 1 | Plan the migration.

- Ensure that both the M-500 appliance and the intermediate Panorama virtual appliance are running the same PAN-OS version. Upgrade the M-700 appliance to a recommended supported PAN-OS version.

  In the second phase of the migration, before migrating the configurations from the Panorama Virtual appliance to the M-700 appliance, you must upgrade the Panorama virtual appliance to the same PAN-OS version that is running on the M-700 appliance. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
- Ensure that the M-500 appliance, the intermediate Panorama virtual appliance, and the M-700 appliance are on the same system mode.
- Schedule a maintenance window for the migration. Firewalls can buffer logs after the M-500 appliance goes offline and then forward the logs after the M-700 appliance comes online. However, completing the migration during a maintenance window ensures that the logs do not exceed the buffer capacities and are not lost during the transition between the Panorama models.

STEP 2 | Purchase the new M-700 appliance, and migrate your subscriptions to the new appliance.

1. Purchase the new M-700 appliance.
2. Purchase the new support license and migration license.
3. When purchasing the new M-700 appliance, provide your sales representative with the serial number and device management auth-code of the M-500 appliance that you are phasing out, and the date when you expect your migration. After you receive the M-700 appliance, register it and activate the device management and support licenses by using the migration and support auth-codes from Palo Alto Networks. On the migration date, the device management license on the M-500 will be decommissioned, preventing you from managing devices or collecting logs using the M-500 appliance. However, the support license is preserved and the Panorama appliance remains under support. You can complete the migration after the effective date, but you will not be able to commit any configuration changes on the decommissioned M-500 appliance. Palo Alto Networks allows up to a 90 day migration grace period when migrating between M-Series appliances. Contact your Palo Alto Networks sales representative for more information about your migration.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Obtain and apply an evaluation or temporary license on the intermediate Panorama virtual appliance.

1. Log in to the Palo Alto Networks Customer Support Portal.
2. Select **Assets** > **Devices** > **Register New Device**.
3. In the **Device Type** window, select **Register device using Serial Number or Authorization Code**, and click **Next**.
4. To activate the Panorama software, enter the serial number you received in the **Request for Software Evaluation Approved** email.
5. If you plan to use the Panorama software offline, select **Device will be used Offline**, and enter the required information.
6. Review the EULA and Support Agreement.
7. If you agree, click **Agree** and **Submit**.
8. After successful registration, the Assets screen displays the newly registered and activated Eval Panorama.

STEP 4 | Perform the initial setup of the intermediate Panorama virtual appliance. For details, see Perform the initial setup of the Panorama virtual appliance.

STEP 5 | Edit the M-500 interface configuration to use only the management interface.

The Panorama virtual appliance supports only the management interface for device management and log collection.

1. Log in to the Panorama web interface of the M-Series appliance.
2. Select **Panorama** > **Setup** > **Management**.
3. Edit the **General Settings**, modify the **Hostname**, and click **OK**.
4. Select **Panorama** > **Setup** > **Interfaces** > **Management interface**, and enable the required services.
5. Disable the services for the other interfaces.
6. Select **Commit** > **Commit to Panorama**.

STEP 6 | Export the Panorama configuration from the M-500 appliance.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Save named Panorama configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**.
4. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**.

   Panorama exports the configuration to your client system as an XML file.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Load the Panorama configuration snapshot that you exported from the M-500 appliance into the Panorama virtual appliance.

> *The Panorama **Policy** rule **Creation** and **Modified** dates are updated to reflect the date you commit the imported Panorama configuration on the new Panorama. The universally unique identifier (UUID) for each policy rule persists when you migrate the Panorama configuration.*
>
> *The **Creation** and **Modified** for managed firewalls are not impacted when you monitor policy rule usage for a managed firewall because this data is stored locally on the managed firewall and not on Panorama.*

1. Log in to the Panorama web interface of the Panorama virtual appliance.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Import named Panorama configuration snapshot**.
4. **Browse** for the configuration file you exported from the M-500 appliance, and click **OK**.
5. Click **Load named Panorama configuration snapshot**, and select the **Name** of the configuration you just imported.
6. Select a **Decryption Key** (the master key for Panorama) and click **OK**.
7. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file. If errors occur, save the errors to a local file. Resolve each error to ensure the migrated configuration is valid.

STEP 8 | Log in to the Panorama web interface of the M-700 appliance, select **Panorama** > **Setup** > **Interfaces**, and verify that the IP address on the management interface is different from the IP address of the M-500 appliance.

This is to ensure that the connectivity to the Panorama virtual appliance is not disrupted post commit.

STEP 9 | Select **Commit** > **Commit to Panorama** > **Validate Commit** to review and resolve any configuration issues. Commit the Panorama configuration.

STEP 10 | Upgrade the intermediate Panorama virtual appliance to the same version installed on the M-700 appliance.

STEP 11 | Export the Panorama configuration from the Panorama virtual appliance.

1. Log in to the Panorama web interface of the Panorama virtual appliance.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Save named Panorama configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**.
4. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**.

   Panorama exports the configuration to your client system as an XML file.

STEP 12 | Perform the initial setup of the M-700 appliance. For details, see Perform the initial setup of the M-Series appliance.

 ©2025 Palo Alto Networks, Inc.

STEP 13 | Load the Panorama configuration snapshot that you exported from the Panorama virtual appliance to the M-700 appliance.

> *The Panorama **Policy** rule **Creation** and **Modified** dates are updated to reflect the date you commit the imported Panorama configuration on the new Panorama. The universally unique identifier (UUID) for each policy rule persists when you migrate the Panorama configuration.*
>
> *The **Creation** and **Modified** for managed firewalls are not impacted when you monitor policy rule usage for a managed firewall because this data is stored locally on the managed firewall and not on Panorama.*

1. Log in to the Panorama web interface of the Panorama virtual appliance.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Import named Panorama configuration snapshot**.
4. **Browse** for the configuration file you exported from the Panorama virtual appliance, and click **OK**.
5. Click **Load named Panorama configuration snapshot**, and select the **Name** of the configuration you just imported.
6. Select a **Decryption Key** (the master key for Panorama) and click **OK**.
7. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file. If errors occur, save the errors to a local file. Resolve each error to ensure the migrated configuration is valid.

STEP 14 | Review the network configuration on the M-700 appliance.

1. (Optional) Log in to the Panorama web interface of the M-500 appliance, select **Panorama** > **Setup** > **Operations**, and click **Shutdown Panorama**.

   Shut down the M-500 appliance if you plan to have the same IP address on both the M-500 and M-700 appliances.
2. Log in to the Panorama web interface of the M-700 appliance, select **Panorama** > **Setup** > **Interfaces**, and verify the network configuration on the Management interface to ensure that the connectivity to the M-700 appliance is not disrupted post commit.
3. Ensure that all the interface configurations are set up based on your requirements for the M-700 appliance.

STEP 15 | Select **Commit** > **Commit to Panorama** > **Validate Commit** to review and resolve any configuration issues. Commit the Panorama configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 16 | Generate a new device registration authentication key for managed device connectivity.

1. In the Panorama web interface of the M-700 appliance, 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 the duration of the validity of the authentication key.
   - **Count**—Enter the number of devices that will use the authentication key for connecting to Panorama.
   - **Device Type**—Specify whether the authentication key may be used for **Firewalls**, **Log Collectors**, or **Any** device.
3. Click **OK**.
4. **Copy Auth Key** and **Close**.

STEP 17 | After you complete the migration, connectivity to the managed firewalls is lost. Recover connectivity to the managed firewalls.

1. Log in to the firewall CLI.
2. Reset the secure connection state.

   This command resets the managed device connection to Panorama and is irreversible.

```
admin>
request sc3 reset
```

3. Restart the management server on the managed device.

```
admin>
  debug software restart process management-server
```

4. (Optional) If the IP address of the M-700 appliance is different from the M-500 appliance, update the panorama-server IP address.

```
admin> configure
admin# set deviceconfig system panorama local-panorama
 panorama-server <panorama-ip> [panorama-server-2 <panorama-
ha-peer-ip>]
admin# commit
admin# exit
```

5. Add the device registration authentication key you copied in Step 16.

```
admin>
      request authkey set <auth_key>
```

6. Verify the managed device connectivity to Panorama.

```
admin>
```

 ©2025 Palo Alto Networks, Inc.

```
show panorama-status
```

Verify that the IP address of the M-700 appliance appears and the Panorama server **Connected** status displays **yes**.

STEP 18 | After you complete the migration, connectivity to the managed log collectors is lost. Recover connectivity to the managed log collectors.

1. Log in to the Dedicated Log Collector CLI.
2. Reset the secure connection state.

   This command resets the managed device connection to Panorama and is irreversible.

   ```
   admin>
   request sc3 reset
   ```

3. Restart the management server on the managed device.

   ```
   admin>
     debug software restart process management-server
   ```

4. (Optional) If the IP address of the M-700 appliance is different from the M-500 appliance, update the panorama-server IP address.

   ```
   admin> configure
   admin # set deviceconfig system panorama-server <panorama-ip>
    [panorama-server-2 <panorama-ha-peer-ip>]
   admin# Commit
   admin# exit
   ```

5. Add the device registration authentication key you copied in Step 16.

   ```
   admin>
         request authkey set <auth_key>
   ```

6. Verify the managed device connectivity to Panorama.

   ```
   admin>
   show panorama-status
   ```

   Verify that the IP address of the M-700 appliance appears and the Panorama server **Connected** status displays **yes**.

STEP 19 | Select **Commit** > **Commit to Panorama** > **Validate Commit** to review and resolve any configuration issues. Commit the Panorama configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 20 | Synchronize the M-700 appliance with the managed devices.

1. Select **Commit** > **Push to Devices** and **Edit Selections**.
2. Select all the devices under **Device Groups**, **Templates**, and **Collector Groups**, and click **OK**.
3. Push your changes.
4. Select **Panorama** > **Managed Devices** > **Summary**, and verify that all the firewalls are connected. Also, verify that the shared policy and template configurations of the firewalls are **In sync** with Panorama.
5. Select **Panorama** > **Managed Collectors**, and verify that the configuration status is **In Sync** with Panorama, and the health status is **Green** for all the log collectors.

Migrate from an M-600 Appliance to an M-700 Appliance

This procedure describes the Panorama configuration migration from an M-600 appliance to an M-700 appliance in Panorama mode. Log migration is not supported. The M-600 appliance logging disk form factor is not supported on the M-600 and M-700 appliances.

This procedure assumes you are no longer using the M-600 appliance for device management or log collection. If you intend to continue using the M-600 appliance as log collector, you must get a device management license for the M-600 appliance. Without a device management license, you cannot use the M-600 appliance as a log collector.

You can continue to access the existing log data at a later date if you do not plan to use the M-600 appliance as a log collector. After you successfully migrate to the new M-700 appliance, power on the M-600 appliance to query and generate reports from the Panorama web interface of the decommissioned M-series appliance. Palo Alto Networks recommends reviewing the log retention policy before decommissioning the M-600 appliance.

Policy rule usage data *is not preserved when you migrate to a different Panorama model. This indicates that all the existing policy rule usage data from the old Panorama model is no longer displayed after you migrate to a new Panorama model. After a successful migration, Panorama begins tracking policy rule usage data based on the date the migration was completed. For example, the* `Created` *date displays the date the migration was completed.*

STEP 1 | Plan the migration.

- Upgrade the M-600 appliance to the latest software and content version.
- Install the latest software and content versions on the M-700 appliance. Both the M-600 and M-700 appliances must run on the same PAN-OS version. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
- Ensure that both the M-600 and M-700 appliances are on the same system mode.
- Schedule a maintenance window for the migration. Firewalls can buffer logs after the M-600 appliance goes offline and then forward the logs after the M-700 appliance comes online. However, completing the migration during a maintenance window ensures that the

 ©2025 Palo Alto Networks, Inc.

logs do not exceed the buffer capacities and are not lost during the transition between Panorama models.

STEP 2 | Purchase the new M-700 appliance, and migrate your subscriptions to the new appliance.

1. Purchase the new M-700 appliance.
2. Purchase the new support license and migration license.
3. When purchasing the new M-700 appliance, provide your sales representative with the serial number and device management auth-code of the M-700 appliance that you are phasing out, and the date when you expect your migration. After you receive the M-700 appliance, register it and activate the device management and support licenses by using the migration and support auth-codes from Palo Alto Networks. On the migration date, the device management license on the M-600 will be decommissioned, preventing you from managing devices or collecting logs using the M-600 appliance. However, the support license is preserved and the Panorama appliance remains under support. You can complete the migration after the effective date, but you will not be able to commit any configuration changes on the decommissioned M-600 appliance. Palo Alto Networks allows up to a 90 day migration grace period when migrating between M-Series appliances. Contact your Palo Alto Networks sales representative for more information about your migration.

STEP 3 | Export the Panorama configuration from the M-600 appliance.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Save named Panorama configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**.
4. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**.

   Panorama exports the configuration to your client system as an XML file.

STEP 4 | Perform the initial setup of the M-700 appliance.

1. Rack mount the M-700 appliances. For more information, refer to the M-300 and M-700 Appliance Hardware Reference guide.
2. Perform Initial Configuration of the M-Series Appliance to define the network connections required to activate licenses and install updates.
3. Register Panorama.
4. Activate a Panorama Support License.
5. For Panorama Managed Prisma Access, ensure that you transfer the licenses to the new panorama appliance.
6. (FIPS-CC only) Retrieve the license key from the license server when migrating from normal mode to FIPS-CC mode.
7. Activate a firewall management license. Use the auth-code associated with the migration license.
8. Install Content and Software Updates for Panorama. Install the same versions as those on the M-600 appliance.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Load the Panorama configuration snapshot that you exported from the M-600 appliance into the M-700 Panorama mode.

*The Panorama **Policy** rule **Creation** and **Modified** dates are updated to reflect the date you commit the imported Panorama configuration on the new Panorama. The universally unique identifier (UUID) for each policy rule persists when you migrate the Panorama configuration.*

*The **Creation** and **Modified** for managed firewalls are not impacted when you monitor policy rule usage for a managed firewall because this data is stored locally on the managed firewall and not on Panorama.*

1. Log in to the Panorama web interface of the M-700 appliance.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Import named Panorama configuration snapshot**.
4. **Browse** for the configuration file you exported from the M-600 appliance, and click **OK**.
5. Click **Load named Panorama configuration snapshot**, and select the **Name** of the configuration you just imported.
6. Select a **Decryption Key** (the master key for Panorama) and click **OK**.
7. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file. If errors occur, save the errors to a local file. Resolve each error to ensure the migrated configuration is valid.

STEP 6 | Review the network configuration on the M-700 appliance.

1. (Optional) Log in to the Panorama web interface of the M-600 appliance, select **Panorama** > **Setup** > **Operations**, and click **Shutdown Panorama**.

   Shut down the M-600 appliance if you plan to have the same IP address on both the M-600 and M-700 appliances.
2. Log in to the Panorama web interface of the M-700 appliance, select **Panorama** > **Setup** > **Interfaces**, and verify the network configuration on the Management interface to ensure that the connectivity to the M-700 appliance is not disrupted post commit.
3. Ensure that all the interface configurations are set up based on your requirements for the M-700 appliance.

STEP 7 | Select **Commit** > **Commit to Panorama** > **Validate Commit** to review and resolve any configuration issues. Commit the Panorama configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Generate a new device registration authentication key for managed device connectivity.

1. In the Panorama web interface of the M-700 appliance, 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 the duration of the validity of the authentication key.
   - **Count**—Enter the number of devices that will use the authentication key for connecting to Panorama.
   - **Device Type**—Specify whether the authentication key may be used for **Firewalls**, **Log Collectors**, or **Any** device.
3. Click **OK**.
4. **Copy Auth Key** and **Close**.

STEP 9 | After you complete the migration, connectivity to the managed firewalls is lost. Recover connectivity to the managed firewalls.

1. Log in to the firewall CLI.
2. (Optional) If the IP address of the M-700 appliance is different from the M-600 appliance, update the panorama-server IP address.

```
admin> configure
admin# set deviceconfig system panorama local-panorama
 panorama-server <panorama-ip> [panorama-server-2 <panorama-
ha-peer-ip>]
admin# commit
admin# exit
```

3. Reset the secure connection state.

   This command resets the managed device connection to Panorama and is irreversible.

```
admin>
request sc3 reset
```

4. Restart the management server on the managed device.

```
admin>
  debug software restart process management-server
```

5. Add the device registration authentication key you copied in Step 8.

```
admin>
      request authkey set <auth_key>
```

6. Verify the managed device connectivity to Panorama.

```
admin>
```

 ©2025 Palo Alto Networks, Inc.

```
show panorama-status
```

Verify that the IP address of the M-700 appliance appears and the Panorama server **Connected** status displays **yes**.

STEP 10 | After you complete the migration, connectivity to the managed log collectors is lost. Recover connectivity to the managed log collectors.

1. Log in to the Dedicated Log Collector CLI.
2. (Optional) If the IP address of the M-700 appliance is different from the M-600 appliance, update the panorama-server IP address.

```
admin> configure
admin # set deviceconfig system panorama-server <panorama-ip>
 [panorama-server-2 <panorama-ha-peer-ip>]
admin# Commit
admin# exit
```

3. Reset the secure connection state.

   This command resets the managed device connection to Panorama and is irreversible.

```
admin>
request sc3 reset
```

4. Restart the management server on the managed device.

```
admin>
  debug software restart process management-server
```

5. Add the device registration authentication key you copied in Step 8.

```
admin>
      request authkey set <auth_key>
```

6. Verify the managed device connectivity to Panorama.

```
admin>
show panorama-status
```

Verify that the IP address of the M-700 appliance appears and the Panorama server **Connected** status displays **yes**.

STEP 11 | Select **Commit** > **Commit to Panorama** > **Validate Commit** to review and resolve any configuration issues. Commit the Panorama configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 12 | Synchronize the M-700 appliance with the managed devices.

1. Select **Commit** > **Push to Devices** and **Edit Selections**.
2. Select all the devices under **Device Groups**, **Templates**, and **Collector Groups**, and click **OK**.
3. Push your changes.
4. Select **Panorama** > **Managed Devices** > **Summary**, and verify that all the firewalls are connected. Also, verify that the shared policy and template configurations of the firewalls are **In sync** with Panorama.
5. Select **Panorama** > **Managed Collectors**, and verify that the configuration status is **In Sync** with Panorama, and the health status is **Green** for all the log collectors.

Migrate from an M-100 Appliance to an M-500 Appliance

You can migrate the Panorama configuration and firewall logs from an M-100 appliance to an M-500 appliance in Panorama mode (Panorama management server). You can also migrate the firewall logs from an M-100 appliance to an M-500 appliance in Log Collector mode (Dedicated Log Collector). Because all the Log Collectors in a Collector Group must be the same Panorama model, you must migrate all or none of the M-100 appliances in any Collector Group.

In the following procedure, the Panorama management server is deployed in an active/passive high availability (HA) configuration, you will migrate both the configuration and logs, and the M-500 appliances will reuse the IP addresses from the M-100 appliances.

This procedure assumes you are no longer using the M-100 for device management or log collection. If you plan on using the decommissioned M-100 appliance as a Dedicated Log Collector, a device management license is required on the M-100. Without a device management license, you are unable to use the M-100 as a Dedicated Log Collector.

If you do not plan on using the M-100 appliance as a Dedicated Log Collector, but the M-100 appliance contains log data that you must access at a later date, you may still query and generate reports using the existing log data. Palo Alto Networks recommends reviewing the log retention policy before decommissioning the M-100 appliance.

If you will migrate only the logs and not the Panorama configuration, perform the task Migrate Logs to a New M-Series Appliance in Log Collector Mode *or* Migrate Logs to a New M-Series Appliance in Panorama Mode.

If you will migrate to a new Panorama management server that is not deployed in an HA configuration and the new Panorama must access logs on existing Dedicated Log Collectors, perform the task Migrate Log Collectors after Failure/RMA of Non-HA Panorama.

Policy rule usage data *is not preserved when you transition to a different Panorama model. This means that all existing policy rule usage data from the old Panorama is no longer displayed after a successful migration to a new Panorama model. After a successful migration, Panorama begins tracking policy rule usage data based on the date the migration was completed. For example, the* `Created` *date displays the date the migration was completed.*

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Plan the migration.

- Upgrade the software on the M-100 appliance if its current release is earlier than 7.0; the M-500 appliance requires Panorama 7.0 or a later release. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
- Forward the System and Config logs that Panorama and Log Collectors generate to an external destination before the migration if you want to preserve those logs. The M-Series appliance in Panorama mode stores these log types on its SSD, which you cannot move between models. You can move only the RAID drives, which store firewall logs.
- Schedule a maintenance window for the migration. Although firewalls can buffer logs after the M-100 appliance goes offline and then forward the logs after the M-500 appliance comes online, completing the migration during a maintenance window minimizes the risk that logs will exceed the buffer capacities and be lost during the transition between Panorama models.

STEP 2 | Purchase the new M-500 appliance, and migrate your subscriptions to the new appliance.

1. Purchase the new M-500 appliance.
2. Purchase the new support license and migration license.
3. At the time you purchase the new M-500 appliance, provide your sales representative the serial number and device management auth-code of the M-100 appliance you are phasing out, as well as a license migration date of your choosing. On receipt of your M-500 appliance, register the appliance and activate the device management and support licenses using the migration and support auth-codes provided by Palo Alto Networks. On the migration date, the device management license on the M-100 is decommissioned, and you can no longer manage devices or collect logs using the M-100 appliance. However, the support license is preserved and the Panorama appliance remains under support. You can complete the migration after the effective date, but you are unable to commit any configuration changes on the now decommissioned M-100 appliance.

STEP 3 | Export the Panorama configuration from each M-100 appliance in Panorama mode.

Perform this task on each M-100 appliance HA peer:

1. Log in to the M-100 appliance and select **Panorama** > **Setup** > **Operations**.
2. Click **Save named Panorama configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**.
3. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**. Panorama exports the configuration to your client system as an XML file.

STEP 4 | Power off each M-100 appliance in Panorama mode.

1. Log in to the M-100 appliance HA peer that you will power off.
2. Select **Panorama** > **Setup** > **Operations**, and click **Shutdown Panorama**.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Perform the initial setup of each M-500 appliance.

1. Rack mount the M-500 appliances. Refer to the M-500 Appliance Hardware Reference Guide for instructions.
2. Perform Initial Configuration of the M-Series Appliance to define the network connections required to activate licenses and install updates.
3. Register Panorama.
4. Activate a Panorama Support License.
5. For Panorama Managed Prisma Access, ensure that you transfer the licenses to the new panorama appliance.
6. Activate a firewall management license. Use the auth-code associated with the migration license.
7. Install Content and Software Updates for Panorama. Install the same versions as those on the M-100 appliance.
8. (Dedicated Log Collector only) Set Up the M-Series Appliance as a Log Collector.

STEP 6 | Load the Panorama configuration snapshot that you exported from each M-100 appliance into each M-500 appliance in Panorama mode (both HA peers).



*The Panorama **Policy** rule **Creation** and **Modified** dates are updated to reflect the date you commit the imported Panorama configuration on the new Panorama. The universially unique identifier (UUID) for each policy rule persists when you migrate the Panorama configuration.*

*The **Creation** and **Modified** for managed firewalls are not impacted when you monitor policy rule usage for a managed firewall because this data is stored locally on the managed firewall and not on Panorama.*

Perform this task on each M-500 appliance HA peer:

1. Log in to the M-500 appliance and select **Panorama** > **Setup** > **Operations**.
2. Click **Import named Panorama configuration snapshot**, **Browse** to the configuration file you exported from the M-100 appliance that has the same HA priority (primary or secondary) as the M-500 appliance will have, and click **OK**.
3. Click **Load named Panorama configuration snapshot**, select the **Name** of the configuration you just imported, select a **Decryption Key** (the master key for Panorama), and click **OK**. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file. If errors occurred, save them to a local file. Resolve each error to ensure the migrated configuration is valid.
4. Select **Commit** > **Commit to Panorama** and **Validate Commit**. Resolve any errors before proceeding.
5. **Commit** your changes to the Panorama configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Synchronize the configuration between the M-500 appliance HA peers in Panorama mode.

1. On the active M-500 appliance, select the **Dashboard** tab and, in the High Availability widget, click **Sync to peer**.
2. In the High Availability widget, verify that the **Local** (primary M-500 appliance) is **active**, the **Peer** is passive, and the **Running Config** is **synchronized**.

STEP 8 | Move the RAID drives from each M-100 appliance to its replacement M-500 appliance to migrate the logs collected from firewalls.

In the following tasks, skip any steps that you already completed on the M-500 appliance.

- Migrate Logs to a New M-Series Appliance in Panorama Mode. Migrate logs from the M-100 appliance only if it uses a default managed collector for log collection.
- Migrate Logs to a New M-Series Appliance in Log Collector Mode.

STEP 9 | Recover Managed Device Connectivity to Panorama for managed firewalls and Dedicated Log Collectors added using the device registration authentication key.

This is required when transitioning from one Panorama model to another.

STEP 10 | Synchronize the active M-500 appliance in Panorama mode with the firewalls to resume firewall management.

Complete this step during a maintenance window to minimize network disruption.

1. In the active M-500 appliance, select **Panorama** > **Managed Devices**, and verify that the Device State column displays **Connected** for the firewalls.

   At this point, the Shared Policy (device groups) and Template columns display **Out of sync** for the firewalls.
2. Push your changes to device groups and templates:
   1. Select **Commit** > **Push to Devices** and **Edit Selections**.
   2. Select **Device Groups**, select every device group, **Include Device and Network Templates**, and click **OK**.
   3. **Push** your changes.
3. In the **Panorama** > **Managed Devices** page, verify that the Shared Policy and Template columns display **In sync** for the firewalls.

After you migrate to a different Panorama model, if there are connectivity issues between Panorama and the managed firewalls, recover the connectivity of the managed devices to Panorama to resolve the issues.

Migrate from an M-100 or M-500 Appliance to an M-200 or M-600 Appliance

This procedure describes the Panorama configuration migration for the following M-Series appliances in Panorama mode (Panorama management server):

 ©2025 Palo Alto Networks, Inc.

- M-100 appliance to an M-200 or M-600 appliance.

  Log migration is not supported. The M-100 appliance logging disk form factor is not supported on the M-200 and M-600 appliances.
- M-500 appliance to an M-200 or M-600 appliance.

  Log migration is not supported. The M-500 appliance logging disk form factor is not supported on the M-200 and M-600 appliances.

Additionally, all the Log Collectors in a Collector Group must be the same Panorama model. For example, if you want to add the local Log Collector on the new M-200 appliance to a Collector Group, the target Collector Group must contain only M-200 appliances. The same is true for the local Log Collector for an M-600 appliance.

This procedure assumes you are no longer using the M-100 or M-500 appliance for device management or log collection. If you plan on using the decommissioned M-100 or M-500 appliance as a Dedicated Log Collector, a device management license is required on the M-100 or M-500 appliance. Without a device management license, you are unable to use the M-100 or M-500 as a Dedicated Log Collector.

You may still access existing log data at a later date if you do not plan on using the M-100 or M-500 appliance as a Dedicated Log Collector. After you have successfully migrate to the new M-Series appliance, power on the M-100 or M-500 appliance to query and generate reports from the Panorama web interface of the decommissioned M-Series appliance. Palo Alto Networks recommends reviewing the log retention policy before decommissioning the M-100 or M-500 appliance.

Policy rule usage data *is not preserved when you transition to a different Panorama model. This means that all existing policy rule usage data from the old Panorama is no longer displayed after a successful migration to a new Panorama model. After a successful migration, Panorama begins tracking policy rule usage data based on the date the migration was completed. For example, the* `Created` *date displays the date the migration was completed.*

STEP 1 | Plan the migration.

- Upgrade the software on the M-100 or M-500 appliance to a supported PAN-OS release. Review the Palo Alto Network Compatibility Matrix for the minimum supported PAN-OS version.

  See the Palo Alto Networks End-of-Life Summary for a list of currently supported PAN-OS versions. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
- Schedule a maintenance window for the migration. Although firewalls can buffer logs after the M-100 or M-500 appliance goes offline and then forward the logs after the M-200 or M-600 appliance comes online, completing the migration during a maintenance window minimizes the risk that logs will exceed the buffer capacities and be lost during the transition between Panorama models.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Purchase the new M-200 or M-600 appliance, and migrate your subscriptions to the new appliance.

1. Purchase the new M-200 or M-600 appliance.
2. Purchase the new support license and migration license.
3. At the time you purchase the new M-200 or M-600 appliance, provide your sales representative the serial number and device management auth-code of the M-100 or M-500 appliance you are phasing out, as well as a license migration date of your choosing. On receipt of your M-200 or M-600 appliance, register the appliance and activate the device management and support licenses using the migration and support auth-codes provided by Palo Alto Networks. On the migration date, the device management license on the M-100 or M-500 is decommissioned, and you can no longer manage devices or collect logs using the M-100 or M-500 appliance. However, the support license is preserved and the Panorama appliance remains under support. You can complete the migration after the effective date, but you are unable to commit any configuration changes on the now decommissioned M-100 or M-500 appliance.

   Palo Alto Networks allows up to a 90 day migration grace period when migrating between M-Series appliances. Please contact your Palo Alto Networks sales representative for more information regarding your migration.

STEP 3 | Export the Panorama configuration from each M-100 or M-500 appliance in Panorama mode.

(HA configuration) Perform this step on each M-100 or M-500 appliance HA peer. Keep track of the HA priority (Primary or Secondary) of the M-100 or M-500 appliance.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Save named Panorama configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**.
4. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**. Panorama exports the configuration to your client system as an XML file.

STEP 4 | In the Panorama web interface of the M-100 or M-500 appliance HA peer that you will power off, select **Panorama** > **Setup** > **Operations** and **Shutdown Panorama**.

(HA configuration) Repeat this step for both M-100 or M-500 appliance HA peers.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Perform the initial setup of the M-200 or M-600 appliance.

(HA configuration) Repeat this step for both M-200 or M-600 appliance HA peers.

1. Rack mount the M-500 appliances. Refer to the M-200 and M-600 Appliance Hardware Reference Guide for instructions.
2. Perform Initial Configuration of the M-Series Appliance to define the network connections required to activate licenses and install updates.
3. Register Panorama.
4. Activate a Panorama Support License.
5. For Panorama Managed Prisma Access, ensure that you transfer the licenses to the new panorama appliance.
6. (FIPS-CC only) Retrieve the license key from the license server when migrating from normal mode to FIPS-CC mode.
7. Activate a firewall management license. Use the auth-code associated with the migration license.
8. Install Content and Software Updates for Panorama. Install the same versions as those on the M-100 or M-500 appliance.
9. (Dedicated Log Collector only) Set Up the M-Series Appliance as a Log Collector.

STEP 6 | Load the Panorama configuration snapshot that you exported from each M-100 or M-500 appliance into each M-200 or M-600 appliance in Panorama mode.

(HA configuration) Repeat this step for both M-200 or M-600 appliances HA peers.

*The Panorama **Policy** rule **Creation** and **Modified** dates are updated to reflect the date you commit the imported Panorama configuration on the new Panorama. The universally unique identifier (UUID) for each policy rule persists when you migrate the Panorama configuration.*

*The **Creation** and **Modified** for managed firewalls are not impacted when you monitor policy rule usage for a managed firewall because this data is stored locally on the managed firewall and not on Panorama.*

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Setup** > **Operations**.
3. Click **Import named Panorama configuration snapshot**.
4. **Browse** for the configuration file you exported from the M-100 or M-500 appliance that has the same HA priority (Primary or Secondary) as the M-200 or M-600 appliance will have and click **OK**.
5. **Load named Panorama configuration snapshot** and select the **Name** of the configuration you just imported.
6. Select a **Decryption Key** (the master key for Panorama) and click **OK**.
7. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file. If

 ©2025 Palo Alto Networks, Inc.

errors occurred, save them to a local file. Resolve each error to ensure the migrated configuration is valid.

8. Select **Commit** > **Commit to Panorama** and **Validate Commit**. Resolve any errors before proceeding.
9. **Commit** your changes to the Panorama configuration.

STEP 7 | Synchronize the configuration between the M-200 or M-600 appliance HA peers in Panorama mode.

1. In the Panorama web interface of the active M-200 or M-600 appliance, select the **Dashboard**.
2. In the High Availability widget, click **Sync to peer**.
3. In the High Availability widget, verify that the **Local** (Primary M-200 appliance) is **active**, the **Peer** is passive, and the **Running Config** is `synchronized`.

STEP 8 | Recover Managed Device Connectivity to Panorama for managed firewalls and Dedicated Log Collectors added using the device registration authentication key.

This is required when transitioning from one Panorama model to another.

STEP 9 | Synchronize the active M-200 or M-600 appliance in Panorama mode with the firewalls to resume firewall management.

Complete this step during a maintenance window to minimize network disruption.

1. In the active M-200 or M-600 appliance, select **Panorama** > **Managed Devices**, and verify that the Device State column displays `Connected` for the firewalls.

   At this point, the Shared Policy (device groups) and Template columns display `Out of Sync` for the firewalls.
2. Push your changes to device groups and templates:
   1. Select **Commit** > **Push to Devices** and **Edit Selections**.
   2. Select **Device Groups**, select every device group, **Include Device and Network Templates**, and click **OK**.
   3. **Push** your changes.
3. In the **Panorama** > **Managed Devices** page, verify that the Shared Policy and Template columns display **In sync** for the firewalls.

After you migrate to a different Panorama model, if there are connectivity issues between Panorama and the managed firewalls, recover the connectivity of the managed devices to Panorama to resolve the issues.

 ©2025 Palo Alto Networks, Inc.

Access and Navigate Panorama Management Interfaces

Panorama provides three management interfaces:

- **Web interface**—The Panorama web interface has a look and feel similar to the firewall web interface. If you are familiar with the latter, you can easily navigate, complete administrative tasks, and generate reports from the Panorama web interface. This graphical interface enables you to access Panorama using HTTPS and it is the best way to perform administrative tasks. See Log in to the Panorama Web Interface and Navigate the Panorama Web Interface. If you need to enable HTTP access to Panorama, edit the Management Interface Settings on the **Panorama** > **Setup** > **Management** tab.
- **Command line interface (CLI)**—The CLI is a no-frills interface that allows you to type commands in rapid succession to complete a series of tasks. The CLI supports two command modes—operational and configuration—and each has its own hierarchy of commands and statements. When you become familiar with the nesting structure and the syntax for the commands, the CLI enables quick response times and administrative efficiency. See Log in to the Panorama CLI.
- **XML API**—The XML-based API is provided as a web service that is implemented using HTTP/HTTPS requests and responses. It enables you to streamline your operations and integrate with existing, internally developed applications and repositories. For details on using the Panorama API, refer to the PAN-OS and Panorama XML API Usage Guide.

Log in to the Panorama Web Interface

The Panorama web interface for the M-600 appliance in Panorama or Management Only mode is inaccessible if you have 6 or more concurrent API calls. A `504 Gateway Timeout` error is displayed when you try to log in to the Panorama web interface of an M-600 appliance with 6 or more concurrent API calls.

STEP 1 | Launch an internet browser and enter the Panorama IP address using a secure connection (https://<IP address>).

STEP 2 | Log in to Panorama according to the type of authentication used for your account. If logging in to Panorama for the first time, use the default value **`admin`** for your username and password.

- **SAML**—Click **Use Single Sign-On** (SSO). If Panorama performs authorization (role assignment) for administrators, enter your **Username** and **Continue**. If the SAML identity provider (IdP) performs authorization, **Continue** without entering a **Username**. In both cases, Panorama redirects you to the IdP, which prompts you to enter a username and password. After you authenticate to the IdP, the Panorama web interface displays.
- **Any other type of authentication**—Enter your user **Name** and **Password**. Read the login banner and select **I Accept and Acknowledge the Statement Below** if the login page has the banner and check box. Then click **Login**.

STEP 3 | Read and **Close** any messages of the day.

 ©2025 Palo Alto Networks, Inc.

Navigate the Panorama Web Interface

Use the Panorama web interface to configure Panorama, manage and monitor firewalls, Log Collectors, and WildFire appliances and appliance clusters, and access the web interface of each firewall through the **Context** drop-down. Refer to the Panorama online help for details on the options and fields in each web interface tab. The following is an overview of the tabs:

Tab	Description
Dashboard	View general information about the Panorama model and network access settings. This tab includes widgets that display information about applications, logs, system resources, and system settings.
ACC	View the overall risk and threat level on the network, based on information that Panorama gathered from the managed firewalls.
Monitor	View and manage logs and reports.
Device Groups > **Policies**	Create centralized policy rules and apply them to multiple firewalls/device groups. You must Add a Device Group for this tab to display.
Device Groups > **Objects**	Define policy objects that policy rules can reference and that managed firewalls/device groups can share. You must Add a Device Group for this tab to display.
Templates > **Network**	Configure network setting, such as network profiles, and apply them to multiple firewalls. You must Add a Template for this tab to display.
Templates > **Device**	Configure device settings, such as server profiles and admin roles, and apply them to multiple firewalls. You must Add a Template for this tab to display.
Panorama	Configure Panorama, manage licenses, set up high availability, access software updates and security alerts, manage administrative access, and manage the deployed firewalls, Log Collectors, and WildFire appliances and appliance clusters.

Log in to the Panorama CLI

You can log in to the Panorama CLI using a serial port connection or remotely using a Secure Shell (SSH) client.

 ©2025 Palo Alto Networks, Inc.

- Use SSH to log in to the Panorama CLI.

  The same instructions apply to an M-Series appliance in Log Collector mode.

  Optionally, you can Configure an Administrator with SSH Key-Based Authentication for the CLI.

  1. Ensure the following prerequisites are met:
     - You have a computer with network access to Panorama.
     - You know the Panorama IP address.
     - The Management interface supports SSH, which is the default setting. If an administrator disabled SSH and you want to re-enable it: select **Panorama** > **Setup** > **Interfaces**, click **Management**, select **SSH**, click **OK**, select **Commit** > **Commit to Panorama**, and **Commit** your changes to the Panorama configuration.
  2. To access the CLI using SSH:
     1. Enter the Panorama IP address in the SSH client and use port 22.
     2. Enter your administrative access credentials when prompted. After you log in, the message of the day displays, followed by the CLI prompt in Operational mode. For example:

        ```
        admin@ABC_Sydney>
        ```

- Use a serial port connection to log in to the Panorama CLI.
  1. Make sure that you have the following:
     - A null-modem serial cable that connects Panorama to a computer with a DB-9 serial port
     - A terminal emulation program running on the computer
  2. Use the following settings in the terminal emulation software to connect: 9600 baud; 8 data bits; 1 stop bit; No parity; No hardware flow control.
  3. Enter your administrative access credentials when prompted. After you log in, the message of the day displays, followed by the CLI prompt in Operational mode.

- Change to Configuration mode.

  To switch to Configuration mode, enter the following command at the prompt:

  ```
  admin@ABC_Sydney> configure
  ```

  The prompt changes to `admin@ABC_Sydney#`.

 ©2025 Palo Alto Networks, Inc.

Set Up Administrative Access to Panorama

Panorama implements Role-Based Access Control (RBAC) to enable you to specify the privileges and responsibilities of administrators. The following topics describe how to create administrator roles, access domains, and accounts for accessing the Panorama web interface and command line interface (CLI):

- Configure an Admin Role Profile
- Configure an Admin Role Profile for Selective Push to Managed Firewalls
- Configure an Access Domain
- Configure Administrative Accounts and Authentication
- Configure Tracking of Administrator Activity

Configure an Admin Role Profile

Admin Role profiles are custom Administrative Roles that enable you to define granular administrative access privileges to ensure protection for sensitive company information and privacy for end users. As a best practice, create Admin Role profiles that allow administrators to access only the areas of the management interfaces required to perform their jobs.

STEP 1 | Select **Device** > **Admin Roles** and select the **Template** in which to configure a firewall admin role profile.

You must create an Admin Role profile on the firewall and assign it to the Panorama management server Admin Role profile to allow administrators to context switch between Panorama and managed firewall web interfaces.

STEP 2 | Select **Panorama** > **Admin Roles** and click **Add**.

STEP 3 | Enter a **Name** for the profile and select the **Role** type: **Panorama** or **Device Group and Template**.

STEP 4 | Configure access privileges to each functional area of Panorama (**Web UI**) by toggling the icons to the desired setting: Enable (read-write), Read Only, or Disable.

*If administrators with custom roles will commit device group or template changes to managed firewalls, you must give those roles read-write access to **Panorama** > **Device Groups** and **Panorama** > **Templates**. If you upgrade from an earlier Panorama version, the upgrade process provides read-only access to those nodes.*

STEP 5 | If the **Role** type is **Panorama**, configure access to the **XML API** by toggling the Enabled/Disabled icon for each functional area.

STEP 6 | If the **Role** type is **Panorama**, select an access level for the **Command Line** interface: **None** (default), **superuser**, **superreader**, or **panorama-admin**.

STEP 7 | (Optional) To allow **Panorama** administrators to **Context Switch** between the Panorama and firewall web interface, enter the name of **Device Admin Role** you configured in Step 1.

STEP 8 | Click **OK** to save the profile.

 ©2025 Palo Alto Networks, Inc.

Configure an Admin Role Profile for Selective Push to Managed Firewalls

To allow for greater control of configuration changes of managed firewalls, create an admin role profile to enable a Panorama administrator to push configuration for one or more Panorama administrators from the Panorama™ management server to managed firewalls. After you commit selective configuration changes to Panorama, you can select specific Panorama admin changes to review the configuration changes and then push only those changes made by the selected admins to your managed firewalls. Leveraging selective pushes to managed firewalls also reduces the risk of pushing incomplete device group and template configurations to managed firewalls by allowing you to explicitly exclude incomplete configuration changes when you push to managed firewalls. This helps mitigate and avoid potential outages and configuration related issues that could cause network disruptions,.

Administrators with Superuser or Panorama admin role privileges can push and review object level changes of other administrators by default. However, you can modify the Panorama administrator admin roles to modify the object level configuration privileges as needed.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | (Optional) Select **Device** > **Admin Roles** and select the **Template** in which to configure a firewall admin role profile.

You must create an Admin Role profile on the firewall and assign it to the Panorama management server Admin Role profile to allow administrators to context switch between Panorama and managed firewall web interfaces.

STEP 3 | Select **Panorama** > **Admin Roles** and **Add** a new admin role.

STEP 4 | Enter a descriptive **Name** for the admin role.

STEP 5 | Select the **Panorama** admin role.

STEP 6 | Select **Web UI** and navigate to the Commit privileges.

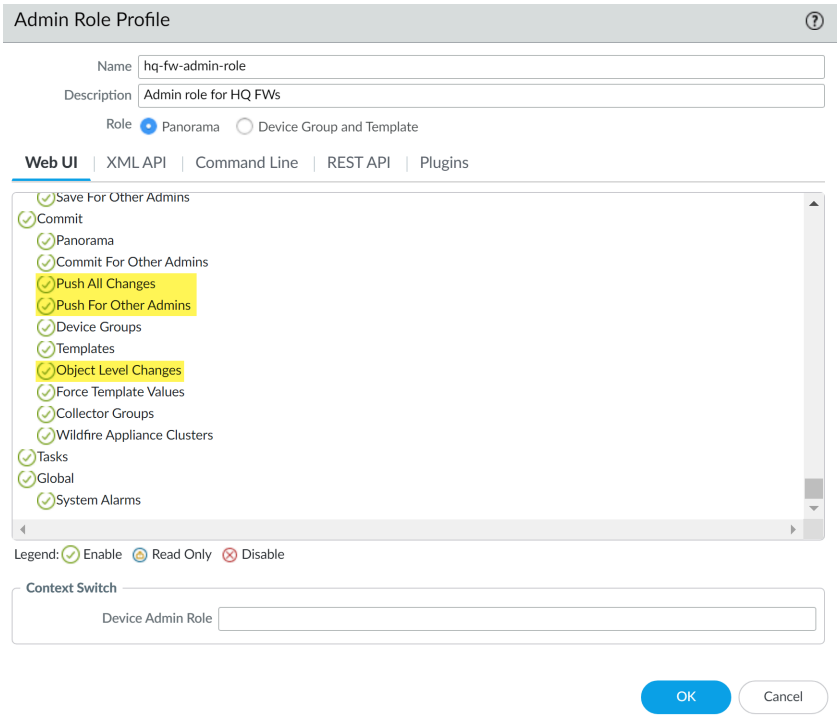
 ©2025 Palo Alto Networks, Inc.

STEP 7 | Configure the object level configuration privileges as needed.

All object level configuration privileges are enabled by default.

The default Superuser or Panorama admin role privileges support full object level configuration privileges.

- **Push All Changes**—Allow the administrator to push all changes made by all admins.
- **Push For Other Admins**—Allows the administrator select and push configuration changes made by other administrators.
- **Object Level Changes**—Allows the administrator to view individual configuration objects to push. If disabled, the list of configuration objects is not displayed in the Push Scope.



STEP 8 | (Optional) To allow **Panorama** administrators to **Context Switch** between the Panorama and firewall web interface, enter the name of **Device Admin Role** you configured in Step 1.

STEP 9 | Click **OK**.

STEP 10 | Configure a custom Panorama administrator and select the **Admin Role** you created.

STEP 11 | **Commit** and **Commit to Panorama**.

Configure an Access Domain

Use Access Domains to define access for Device Group and Template administrators for specific device groups and templates, and also to control the ability of those administrators to switch context to the web interface of managed firewalls. Panorama supports up to 4,000 access domains.

STEP 1 | Select **Panorama** > **Access Domain** and click **Add**.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Enter a **Name** to identify the access domain.

STEP 3 | Select an access privilege for **Shared Objects**:

- **write**—Administrators can perform all operations on Shared objects. This is the default value.
- **read**—Administrators can display and clone but cannot perform other operations on Shared objects. When adding non-Shared objects or cloning Shared objects, the destination must be a device group within the access domain, not the Shared location.
- **shared-only**—Administrators can add objects only to the Shared location. 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 this option is for an organization that requires all objects to be in a single, global repository.

STEP 4 | Toggle the icons in the **Device Groups** tab to enable read-write or read-only access for device groups in the access domain.

*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.*

STEP 5 | Select the **Templates** tab and **Add** each template you want to assign to the access domain.

STEP 6 | Select the **Device Context** tab, select firewalls to assign to the access domain, and click **OK**. Administrators can access the web interface of these firewalls by using the **Context** drop-down in Panorama.

Configure Administrative Accounts and Authentication

If you have already configured an authentication profile or you don't require one to authenticate administrators, you are ready to Configure a Panorama Administrator Account. Otherwise, perform one of the other procedures listed below to configure administrative accounts for specific types of authentication.

- Configure a Panorama Administrator Account
- Configure Local or External Authentication for Panorama Administrators
- Configure a Panorama Administrator with Certificate-Based Authentication for the Web Interface
- Configure an Administrator with SSH Key-Based Authentication for the CLI
- Configure RADIUS Authentication for Panorama Administrators
- Configure TACACS+ Authentication for Panorama Administrators
- Configure SAML Authentication for Panorama Administrators

Configure a Panorama Administrator Account

Administrative accounts specify Administrative Roles and authentication for Panorama administrators. The service that you use to assign roles and perform authentication determines

 ©2025 Palo Alto Networks, Inc.

whether you add the accounts on Panorama, on an external server, or both (see Administrative Authentication). For an external authentication service, you must configure an authentication profile before adding an administrative account (see Configure Administrative Accounts and Authentication). If you already configured the authentication profile or you will use the authentication mechanism that is local to Panorama, perform the following steps to add an administrative account on Panorama.

STEP 1 | Modify the number of supported administrator accounts.

Configure the total number of supported concurrent administrative accounts sessions for Panorama in the normal operational mode or in FIPS-CC mode. You can allow up to four concurrent administrative account sessions or configure Panorama to support an unlimited number of concurrent administrative account sessions.

1. Select **Panorama** > **Setup** > **Management** and edit the Authentication Settings.
2. Edit the **Max Session Count** to specify the number of supported concurrent sessions (range is **0** to **4**) allowed for all administrator and user accounts.

   Enter **0** to configure Panorama to support an unlimited number of administrative accounts.
3. Edit the **Max Session Time** in minutes for an administrative account. Default is 720 minutes.
4. Click **OK**.
5. **Commit** and **Commit to Panorama**.

You can also configure the total number of supported concurrent sessions by logging in to the Panorama CLI.

```
admin> configure
```

```
admin# set deviceconfig setting management admin-session
 max-session-count <0-4>
```

```
admin# set deviceconfig setting management admin-session
 max-session-time <0, 60-1499>
```

```
admin# commit
```

STEP 2 | Select **Panorama** > **Administrators** and **Add** an account.

STEP 3 | Enter a user **Name** for the administrator.

STEP 4 | Select an **Authentication Profile** or sequence if you configured either for the administrator.

This is required if Panorama will use Kerberos SSO or an external service for authentication.

If Panorama will use local authentication, set the **Authentication Profile** to **None** and enter a **Password** and then **Confirm Password**.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Select the **Administrator Type**:

- **Dynamic**—Select a predefined administrator role.
- **Custom Panorama Admin**—Select the Admin Role **Profile** you created for this administrator (see Configure an Admin Role Profile).
- **Device Group and Template Admin**—Map access domains to administrative roles as described in the next step.

STEP 6 | (Device Group and Template Admin only) In the Access Domain to Administrator Role section, click **Add**, select an Access Domain from the drop-down (see Configure an Access Domain), click the adjacent Admin Role cell, and select an Admin Role profile.

STEP 7 | Click **OK** to save your changes.

STEP 8 | Select **Commit** > **Commit to Panorama** and **Commit** your changes.

Configure Local or External Authentication for Panorama Administrators

You can use an external authentication service or the service that is local to Panorama to authenticate administrators who access Panorama. These authentication methods prompt administrators to respond to one or more authentication challenges, such as a login page for entering a username and password.

If you use an external service to manage both authentication and authorization (role and access domain assignments), see:

- Configure RADIUS Authentication for Panorama Administrators
- Configure TACACS+ Authentication for Panorama Administrators
- Configure SAML Authentication for Panorama Administrators

To authenticate administrators without a challenge-response mechanism, you can Configure a Panorama Administrator with Certificate-Based Authentication for the Web Interface *and* Configure an Administrator with SSH Key-Based Authentication for the CLI.

STEP 1 | (External authentication only) Enable Panorama to connect to an external server for authenticating administrators.

1. Select **Panorama** > **Server Profiles**, select the service type (**RADIUS**, **TACACS+**, **SAML**, **LDAP**, or **Kerberos**), and configure a server profile:
   - Configure RADIUS Authentication for Panorama Administrators.

     You can use a RADIUS server to support RADIUS authentication services or multi-factor authentication(MFA) services.

   - Configure TACACS+ Authentication for Panorama Administrators.
   - Add a SAML IdP server profile. You cannot combine Kerberos single sign-on (SSO) with SAML SSO; you can use only one type of SSO service.
   - Add a Kerberos server profile.
   - Add a LDAP Server Profile.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | (Optional) Define password complexity and expiration settings if Panorama uses local authentication.

These settings help protect Panorama against unauthorized access by making it harder for attackers to guess passwords.

1. Define global password complexity and expiration settings for all local administrators.
   1. Select **Panorama** > **Setup** > **Management** and edit the Minimum Password Complexity settings.
   2. Select **Enabled**.
   3. Define the password settings and click **OK**.
2. Define a Password Profile.

   You assign the profile to administrator accounts for which you want to override the global password expiration settings.
   1. Select **Panorama** > **Password Profiles** and **Add** a profile.
   2. Enter a **Name** to identify the profile.
   3. Define the password expiration settings and click **OK**.

STEP 3 | (Kerberos SSO only) Create a Kerberos keytab.

A keytab is a file that contains Kerberos account information for Panorama. To support Kerberos SSO, your network must have a Kerberos infrastructure.

STEP 4 | Configure an authentication profile.



If your administrative accounts are stored across multiple types of servers, you can create an authentication profile for each type and add all the profiles to an authentication sequence.

In the authentication profile, specify the **Type** of authentication service and related settings:

- **External service**—Select the **Type** of external service and select the **Server Profile** you created for it.
- **Local authentication**—Set the **Type** to **None**.
- **Kerberos SSO**—Specify the **Kerberos Realm** and **Import** the **Kerberos Keytab** you created.

STEP 5 | (Device group and template administrators only) Configure an Access Domain.

Configure one or more access domains.

STEP 6 | (Custom roles only) Configure an Admin Role Profile.

Configure one or more Admin Role profiles.

For custom Panorama administrators, the profile defines access privileges for the account. For device group and template administrators, the profile defines access privileges for one or more access domains associated with the account.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Configure an administrator.

1. Configure a Panorama Administrator Account.
   - Assign the **Authentication Profile** or sequence that you configured.
   - (Device Group and Template Admin only) Map the access domains to Admin Role profiles.
   - (Local authentication only) Select a **Password Profile** if you configured one.
2. Select **Commit** > **Commit to Panorama** and **Commit** your changes.
3. (Optional) Test authentication server connectivity to verify that Panorama can use the authentication profile to authenticate administrators.

Configure a Panorama Administrator with Certificate-Based Authentication for the Web Interface

As a more secure alternative to password-based authentication to the Panorama web interface, you can configure certificate-based authentication for administrator accounts that are local to Panorama. Certificate-based authentication involves the exchange and verification of a digital signature instead of a password.

Configuring certificate-based authentication for any administrator disables the username/password logins for all administrators on Panorama and all administrators thereafter require the certificate to log in.

STEP 1 | Generate a certificate authority (CA) certificate on Panorama.

You will use this CA certificate to sign the client certificate of each administrator.

Create a self-signed root CA certificate.

Alternatively, you can import a certificate from your enterprise CA.

STEP 2 | Configure a certificate profile for securing access to the web interface.

1. Select **Panorama** > **Certificate Management** > **Certificate Profile** and click **Add**.
2. Enter a **Name** for the certificate profile and set the **Username Field** to **Subject**.
3. Select **Add** in the CA Certificates section and select the **CA Certificate** you just created.
4. Click **OK** to save the profile.

STEP 3 | Configure Panorama to use the certificate profile for authenticating administrators.

1. Select the **Panorama** > **Setup** > **Management** and edit the Authentication Settings.
2. Select the **Certificate Profile** you just created and click **OK**.

STEP 4 | Configure the administrator accounts to use client certificate authentication.

Configure a Panorama Administrator Account for each administrator who will access the Panorama web interface. Select the **Use only client certificate authentication (Web)** check box.

If you have already deployed client certificates that your enterprise CA generated, skip to Step 8. Otherwise, continue with Step 5.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Generate a client certificate for each administrator.

Generate a certificate on Panorama. In the **Signed By** drop-down, select the CA certificate you created.

STEP 6 | Export the client certificates.

1. Export the certificates.
2. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

   Panorama restarts and terminates your login session. Thereafter, administrators can access the web interface only from client systems that have the client certificate you generated.

STEP 7 | Import the client certificate into the client system of each administrator who will access the web interface.

Refer to your web browser documentation as needed to complete this step.

STEP 8 | Verify that administrators can access the web interface.

1. Open the Panorama IP address in a browser on the computer that has the client certificate.
2. When prompted, select the certificate you imported and click **OK**. The browser displays a certificate warning.
3. Add the certificate to the browser exception list.
4. Click **Login**. The web interface should appear without prompting you for a username or password.

Configure an Administrator with SSH Key-Based Authentication for the CLI

For administrators who use Secure Shell (SSH) to access the Panorama CLI, SSH keys provide a more secure authentication method than passwords. SSH keys almost eliminate the risk of brute-force attacks, provide the option for two-factor authentication (private key and passphrase), and don't send passwords over the network. SSH keys also enable automated scripts to access the CLI.

STEP 1 | Use an SSH key generation tool to create an asymmetric key pair on the client system of the administrator.

The supported key formats are IETF SECSH and Open SSH. The supported algorithms are DSA (1024 bits) and RSA (768-4096 bits).

For the commands to generate the key pair, refer to your SSH client documentation.

The public key and private key are separate files. Save both to a location that Panorama can access. For added security, enter a passphrase to encrypt the private key. Panorama prompts the administrator for this passphrase during login.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Configure the administrator account to use public key authentication.

1. Configure a Panorama Administrator Account.
   - Configure one of two authentication methods to use as a fallback if SSH key authentication fails:

     External authentication service—Select an **Authentication Profile**.

     Local authentication—Set the **Authentication Profile** to **None** and enter a **Password** and **Confirm Password**.
   - Select the **Use Public Key Authentication (SSH)** check box, click **Import Key**, **Browse** to the public key you just generated, and click **OK**.
2. Click **OK** to save the administrative account.
3. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 3 | Configure the SSH client to use the private key to authenticate to Panorama.

Perform this task on the client system of the administrator. Refer to your SSH client documentation as needed to complete this step.

STEP 4 | Verify that the administrator can access the Panorama CLI using SSH key authentication.

1. Use a browser on the client system of the administrator to go to the Panorama IP address.
2. Log in to the Panorama CLI as the administrator. After entering a username, you will see the following output (the key value is an example):

   ```
   Authenticating with public key “dsa-key-20130415”
   ```
3. If prompted, enter the passphrase you defined when creating the keys.

Configure RADIUS Authentication for Panorama Administrators

You can use a RADIUS server to authenticate administrative access to the Panorama web interface. You can also define Vendor-Specific Attributes (VSAs) on the RADIUS server to manage administrator authorization. Using VSAs enables you to quickly change the roles, access domains, and user groups of administrators through your directory service, which is often easier than reconfiguring settings on Panorama.



You can use a RADIUS server to authenticate administrative access to the Panorama web interface. You can also define Vendor-Specific Attributes (VSAs) on the RADIUS server to manage administrator authorization. Using VSAs enables you to quickly change the roles, access domains, and user groups of administrators through your directory service, which is often easier than reconfiguring settings on Panorama.

You can Import the Palo Alto Networks RADIUS dictionary into RADIUS server to define the authentication attributes needed for communication between Panorama and the RADIUS server.

You can also use a RADIUS server to implement multi-factor authentication (MFA) for administrators.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Add a RADIUS server profile.

The profile defines how Panorama connects to the RADIUS server.

1. Select **Panorama** > **Server Profiles** > **RADIUS** and **Add** a profile.
2. Enter a **Profile Name** to identify the server profile.
3. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–20).

   If you use the server profile to integrate Panorama with an MFA service, enter an interval that gives administrators enough time to respond to the authentication challenge. For example, if the MFA service prompts for a one-time password (OTP), administrators need time to see the OTP on their endpoint device and then enter the OTP in the MFA login page.
4. Select the **Authentication Protocol** (default is **CHAP**) that Panorama uses to authenticate to the RADIUS server.

   *Select **CHAP** if the RADIUS server supports that protocol; it is more secure than **PAP**.*
5. **Add** each RADIUS server and enter the following:
   - **Name** to identify the server
   - **RADIUS Server** IP address or FQDN
   - **Secret**/**Confirm Secret** (a key to encrypt usernames and passwords)
   - Server **Port** for authentication requests (default is 1812)
6. Click **OK** to save the server profile.

STEP 2 | Assign the RADIUS server profile to an authentication profile.

The authentication profile defines authentication settings that are common to a set of administrators.

1. Select **Panorama** > **Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the authentication profile.
3. Set the **Type** to **RADIUS**.
4. Select the **Server Profile** you configured.
5. Select **Retrieve user group from RADIUS** to collect user group information from VSAs defined on the RADIUS server.

   Panorama matches the group information against the groups you specify in the Allow List of the authentication profile.
6. Select **Advanced** and, in the Allow List, **Add** the administrators that are allowed to authenticate with this authentication profile.
7. Click **OK** to save the authentication profile.

STEP 3 | Configure Panorama to use the authentication profile for all administrators.

1. Select **Panorama** > **Setup** > **Management** and edit the Authentication Settings.
2. Select the **Authentication Profile** you configured and click **OK**.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the roles and access domains that define authorization settings for administrators.

1. Configure an Admin Role Profile if the administrator uses a custom role instead of a predefined (dynamic) role.
2. Configure an Access Domain if the administrator uses a Device Group and Template role.

STEP 5 | Commit your changes.

Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 6 | Configure the RADIUS server.

Refer to your RADIUS server documentation for the specific instructions to perform these steps:

1. Add the Panorama IP address or hostname as the RADIUS client.
2. Add the administrator accounts.

   *If the RADIUS server profile specifies **CHAP** as the **Authentication Protocol**, you must define accounts with reversibly encrypted passwords. Otherwise, CHAP authentication will fail.*
3. Define the vendor code for Panorama (25461) and define the RADIUS VSAs for the role, access domain, and user group of each administrator.

   When you predefine dynamic administrator roles for users, use lower-case to specify the role (for example, enter **`superuser`**, not **`SuperUser`**).

STEP 7 | Verify that the RADIUS server performs authentication and authorization for administrators.

1. Log in the Panorama web interface using an administrator account that you added to the RADIUS server.
2. Verify that you can access only the web interface pages that are allowed for the role you associated with the administrator.
3. In the **Monitor**, **Policies**, and **Objects** tabs, verify that you can access only the device groups that are allowed for the access domain you associated with the administrator.

Configure TACACS+ Authentication for Panorama Administrators

You can use a TACACS+ server to authenticate administrative access to the Panorama web interface. You can also define Vendor-Specific Attributes (VSAs) on the TACACS+ server to manage administrator authorization. Using VSAs enables you to quickly change the roles, access domains, and user groups of administrators through your directory service, which is often easier than reconfiguring settings on Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Add a TACACS+ server profile.

The profile defines how Panorama connects to the TACACS+ server.

1. Select **Panorama** > **Server Profiles** > **TACACS+** and **Add** a profile.
2. Enter a **Profile Name** to identify the server profile.
3. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–20).
4. Select the **Authentication Protocol** (default is **CHAP**) that Panorama uses to authenticate to the TACACS+ server.

   *Select **CHAP** if the TACACS+ server supports that protocol; it is more secure than **PAP**.*
5. **Add** each TACACS+ server and enter the following:
   - **Name** to identify the server
   - **TACACS+ Server** IP address or FQDN
   - **Secret**/**Confirm Secret** (a key to encrypt usernames and passwords)
   - Server **Port** for authentication requests (default is 49)
6. Click **OK** to save the server profile.

STEP 2 | Assign the TACACS+ server profile to an authentication profile.

The authentication profile defines authentication settings that are common to a set of administrators.

1. Select **Panorama** > **Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the profile.
3. Set the **Type** to **TACACS+**.
4. Select the **Server Profile** you configured.
5. Select **Retrieve user group from TACACS+** to collect user group information from VSAs defined on the TACACS+ server.

   Panorama matches the group information against the groups you specify in the Allow List of the authentication profile.
6. Select **Advanced** and, in the Allow List, **Add** the administrators that are allowed to authenticate with this authentication profile.
7. Click **OK** to save the authentication profile.

STEP 3 | Configure Panorama to use the authentication profile for all administrators.

1. Select **Panorama** > **Setup** > **Management** and edit the Authentication Settings.
2. Select the **Authentication Profile** you configured and click **OK**.

STEP 4 | Configure the roles and access domains that define authorization settings for administrators.

1. Configure an Admin Role Profile if the administrator will use a custom role instead of a predefined (dynamic) role.
2. Configure an Access Domain if the administrator uses a Device Group and Template role.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Commit your changes.

Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 6 | Configure the TACACS+ server to authenticate and authorize administrators.

Refer to your TACACS+ server documentation for the specific instructions to perform these steps:

1. Add the Panorama IP address or hostname as the TACACS+ client.
2. Add the administrator accounts.

   *If you selected **CHAP** as the **Authentication Protocol**, you must define accounts with reversibly encrypted passwords. Otherwise, CHAP authentication will fail.*
3. Define TACACS+ VSAs for the role, access domain, and user group of each administrator.

   *When you predefine dynamic administrator roles for users, use lower-case to specify the role (for example, enter **`superuser`**, not **`SuperUser`**).*

STEP 7 | Verify that the TACACS+ server performs authentication and authorization for administrators.

1. Log in the Panorama web interface using an administrator account that you added to the TACACS+ server.
2. Verify that you can access only the web interface pages that are allowed for the role you associated with the administrator.
3. In the **Monitor**, **Policies**, and **Objects** tabs, verify that you can access only the virtual systems that are allowed for the access domain you associated with the administrator.

Configure SAML Authentication for Panorama Administrators

You can use Security Assertion Markup Language (SAML) 2.0 for administrative access to the Panorama web interface (but not the CLI). You can also use SAML attributes to manage administrator authorization. SAML attributes enable you to quickly change the roles, access domains, and user groups of administrators through your directory service instead of reconfiguring settings on Panorama.

To configure SAML single sign-on (SSO) and single logout (SLO), you must register Panorama and the identity provider (IdP) with each other to enable communication between them. If the IdP provides a metadata file containing registration information, you can import it onto Panorama to register the IdP and to create an IdP server profile. The server profile defines how to connect to the IdP and specifies the certificate that the IdP uses to sign SAML messages. You can also use a certificate for Panorama to sign SAML messages. Using certificates is optional but recommended to secure communications between Panorama and the IdP.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | (Recommended) Obtain the certificates that the IdP and Panorama will use to sign SAML messages.

If the certificates don't specify key usage attributes, all usages are allowed by default, including signing messages. In this case, you can obtain certificates by any method.

If the certificates do specify key usage attributes, one of the attributes must be Digital Signature, which is not available on certificates that you generate on Panorama. In this case, you must import the certificates:

- **Certificate Panorama uses to sign SAML messages**—Import the certificate from your enterprise certificate authority (CA) or a third-party CA.
- **Certificate the IdP uses to sign SAML messages**—Import a metadata file containing the certificate from the IdP (see the next step). The IdP certificate is limited to the following algorithms:
  - **Public key algorithms**—RSA (1,024 bits or larger) and ECDSA (all sizes).
  - **Signature algorithms**—SHA1, SHA256, SHA384, and SHA512.

STEP 2 | Add a SAML IdP server profile.

The server profile registers the IdP with Panorama and defines how they connect.

In this example, you import a SAML metadata file from the IdP so that Panorama can automatically create a server profile and populate the connection, registration, and IdP certificate information.

*If the IdP doesn't provide a metadata file, select **Panorama** > **Server Profiles** > **SAML Identity Provider**, **Add** the server profile, and manually enter the information (consult your IdP administrator for the values).*

1. Export the SAML metadata file from the IdP to a client system that Panorama can access.

   The certificate specified in the file must meet the requirements listed in the preceding step. Refer to your IdP documentation for instructions on exporting the file.
2. Select **Panorama** > **Server Profiles** > **SAML Identity Provider** and **Import** the metadata file onto Panorama.
3. Enter a **Profile Name** to identify the server profile.
4. **Browse** to the **Identity Provider Metadata** file.
5. (Recommended) Select **Validate Identity Provider Certificate** (default) to have Panorama validate the **Identity Provider Certificate**.

   Validation occurs only after you assign the server profile to an authentication profile and **Commit**. Panorama uses the **Certificate Profile** in the authentication profile to validate the certificate.

   Validating the certificate is a best practice for improved security.
6. Enter the **Maximum Clock Skew**, which is the allowed difference in seconds between the system times of the IdP and Panorama at the moment when Panorama validates

 ©2025 Palo Alto Networks, Inc.

IdP messages (default is 60; range is 1 to 900). If the difference exceeds this value, authentication fails.

7. Click **OK** to save the server profile.
8. Click the server profile Name to display the profile settings. Verify that the imported information is correct and edit it if necessary.

STEP 3 | Configure an authentication profile.

The authentication profile specifies a SAML IdP server profile and defines options for the authentication process, such as SLO.

1. Select **Panorama** > **Authentication Profile** and **Add** a profile.
2. Enter a **Name** to identify the profile.
3. Set the **Type** to **SAML**.
4. Select the **IdP Server Profile** you configured.
5. Select the **Certificate for Signing Requests**.

   Panorama uses this certificate to sign messages it sends to the IdP.
6. (Optional) **Enable Single Logout** (disabled by default).
7. Select the **Certificate Profile** that Panorama will use to validate the **Identity Provider Certificate**.
8. Enter the **Username Attribute** that IdP messages use to identify users (default **`username`**).

   *When you predefine dynamic administrator roles for users, use lower-case to specify the role (for example, enter **`superuser`**, not **`SuperUser`**). If you manage administrator authorization through the IdP identity store, specify the **Admin Role Attribute** and **Access Domain Attribute** also.*
9. Select **Advanced** and **Add** the administrators who are allowed to authenticate with this authentication profile.
10. Click **OK** to save the authentication profile.

STEP 4 | Configure Panorama to use the authentication profile for all administrators.

1. Select **Panorama** > **Setup** > **Management**, edit the Authentication Settings, and select the **Authentication Profile** you configured.
2. Select **Commit** > **Commit to Panorama** to activate your changes on Panorama and to validate the **Identity Provider Certificate** that you assigned to the SAML IdP server profile.

STEP 5 | Create a SAML metadata file to register Panorama on the IdP.

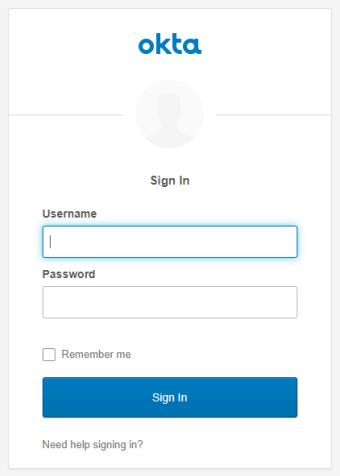
1. Select **Panorama** > **Authentication Profile** and, in the Authentication column for the authentication profile you configured, click **Metadata**.
2. Set the **Management Choice** to **Interface** (default is selected) and select the management (MGT) interface.
3. Click **OK** and save the metadata file to your client system.
4. Import the metadata file into the IdP server to register Panorama. Refer to your IdP documentation for instructions.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Verify that administrators can authenticate using SAML SSO.

1. Go to the URL of the Panorama web interface.
2. Click **Use Single Sign-On**.
3. Click **Continue**.

   Panorama redirects you to authenticate to the IdP, which displays a login page. For example:



4. Log in using your SSO username and password.

   After you successfully authenticate on the IdP, it redirects you back to Panorama, which displays the web interface.
5. Use your Panorama administrator account to request access to another SSO application.

   Successful access indicates SAML SSO authentication succeeded.

Configure Tracking of Administrator Activity

Track administrator activity on the web interface and CLI of your Panorama™ management server, managed firewalls, and Log Collectors to 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.

When an event occurs, an audit log is generated and forwarded to the specified syslog server each time an administrator navigates through the web interface or when an operational command is executed in the CLI. An audit log is generated for each navigation or commend executed. Take for example if you want to create a new address object. An audit log is generated when you click on **Objects**, and a second audit log is generated when you then click on Addresses.

Audit logs are only visible as syslogs forwarded to your syslog server and cannot be viewed in the Panorama or managed firewall web interface. Audit logs can only be forwarded to a syslog server, cannot be forwarded to Strata Logging Service, and are not stored locally on the firewall, Panorama, or Log Collector.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Configure a syslog server profile to forward audit logs of administrator activity for Panorama, managed firewalls, and Log Collectors.

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

1. Select **Panorama** > **Server Profiles** > **Syslog** and **Add** a new syslog server profile.
2. Configure a syslog server profile.

STEP 2 | Configure administrator activity tracking for your managed firewalls.

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

1. Select **Device** > **Setup** > **Management** and edit the Logging and Reporting Settings.
2. Configure Tracking of Administrator Activity.
3. Select **Commit** and **Commit and Push**.

STEP 3 | Configure administrator activity tracking for Panorama.

1. Select **Panorama** > **Setup** > **Management** and edit the Logging and Reporting Settings.
2. Select **Log Export and Reporting**.
3. In the Log Admin Activity section, configure what administrator activity to track.
   - **Operational Commands**—Generate an audit log when an administrator executes an operational or debug command in the CLI or an operational command triggered from

 ©2025 Palo Alto Networks, Inc.

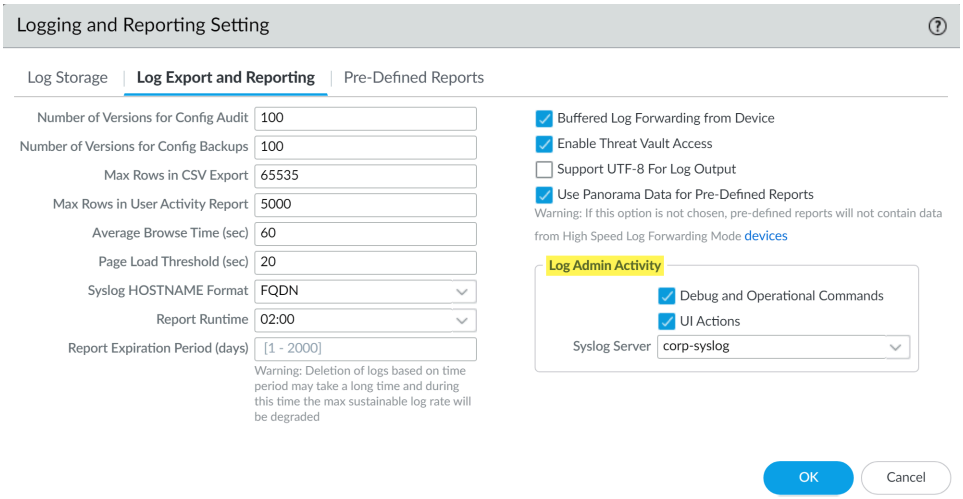
the web interface. See the CLI Operational Command Hierarchy for a full list of PAN-OS operational and debug commands.

- **UI Actions**—Generate an audit log when an administrator navigates throughout the web interface. This includes navigation between configuration tabs, as well as individual objects within a tab.

  For example, an audit log is generated when an administrator navigates from the **ACC** to the **Policies** tab. Additionally, an audit log is generated when an administrator navigates from **Objects** > **Addresses** to **Objects** > **Tags**.

- **Syslog Server**—Select a target syslog server profile to forward audit logs.

4. Click **OK**



5. Select **Commit** and **Commit to Panorama**.

STEP 4 | Configure administrator activity tracking for Log Collectors in a Collector Group.

1. Select **Panorama** > **Collector Groups** and click a Collector Group.
2. Select **Audit**.
3. In the Log Admin Activity section, configure audit tracking for CLI activity.

   You can only track CLI activity for Log Collectors because Log Collectors you can only access Log Collectors through the CLI.

   - **Operational Commands**—Generate an audit log when an administrator executes an operational or debug command in the CLI. See the CLI Operational Command Hierarchy for a full list of PAN-OS operational and debug commands.
   - **Syslog Server**—Select a target syslog server profile to forward audit logs.
4. Click **OK**.
5. Select **Commit** and **Commit to Panorama**.

 ©2025 Palo Alto Networks, Inc.

Set Up Authentication Using Custom Certificates

By default, Palo Alto Networks devices use predefined certificates for mutual authentication to establish the SSL connections used for management access and inter-device communication. However, you can configure authentication using custom certificates instead. Additionally, you can use custom certificates to secure the High Availability (HA) connections between Panorama HA peers. Custom certificates allow you to establish a unique chain of trust to ensure mutual authentication between Panorama and the managed firewalls and log collectors. See Certificate Management for detailed information about the certificates and how to deploy them on Panorama, Log Collectors, and firewalls.

The following topics describe how to configure and manage custom certificates using Panorama.

- How Are SSL/TLS Connections Mutually Authenticated?
- Configure Authentication Using Custom Certificates on Panorama
- Configure Authentication Using Custom Certificates on Managed Devices
- Add New Client Devices
- Change Certificates

How Are SSL/TLS Connections Mutually Authenticated?

In a regular SSL connection, only the server needs to identify itself to the client by presenting its certificate. However, in mutual SSL authentication, the client presents its certificate to the server as well. Panorama, the primary Panorama HA peer, Log Collectors, WildFire appliances, and PAN-DB appliances can act as the server. Firewalls, Log Collectors, WildFire appliances, and the secondary Panorama HA peer can act as the client. The role that a device takes on depends the deployment. For example, in the diagram below, Panorama manages a number of firewalls and a collector group and acts as the server for the firewalls and Log Collectors. The Log Collector acts as the server to the firewalls that send logs to it.

To deploy custom certificates for mutual authentication in your deployment, you need:

- **SSL/TLS Service Profile**—An SSL/TLS service profile defines the security of the connections by referencing your custom certificate and establishing the SSL/TLS protocol versions used by the server device to communicate with client devices.
- **Server Certificate and Profile**—Devices in the server role require a certificate and certificate profile to identify themselves to the client devices. You can deploy this certificate from your enterprise public key infrastructure (PKI), purchase one from a trusted third-party CA, or generate a self-signed certificate locally. The server certificate must include the IP address or FQDN of the device's management interface in the certificate common name (CN) or Subject Alt Name. The client firewall or Log Collector matches the CN or Subject Alt Name in the certificate the server presents against the server's IP address or FQDN to verify the server's identity.

  Additionally, use the certificate profile to define certificate revocation status (OCSP/CRL) and the actions taken based on the revocation status.
- **Client Certificates and Profile**—Each managed device requires a client certificate and certificate profile. The client device uses its certificate to identify itself to the server device. You can deploy certificates from your enterprise PKI, using Simple Certificate Enrollment

 ©2025 Palo Alto Networks, Inc.

Protocol (SCEP), purchase one from a trusted third-party CA, or generate a self-signed certificate locally.

Custom certificates can be unique to each client device or common across all devices. The unique device certificates uses a hash of the serial number of the managed device and CN. The server matches the CN or the subject alt name against the configured serial numbers of the client devices. For client certificate validation based on the CN to occur, the username must be set to Subject common-name. The client certificate behavior also applies to Panorama HA peer connections.

You can configure the client certificate and certificate profile on each client device or push the configuration from Panorama to each device as part of a template.

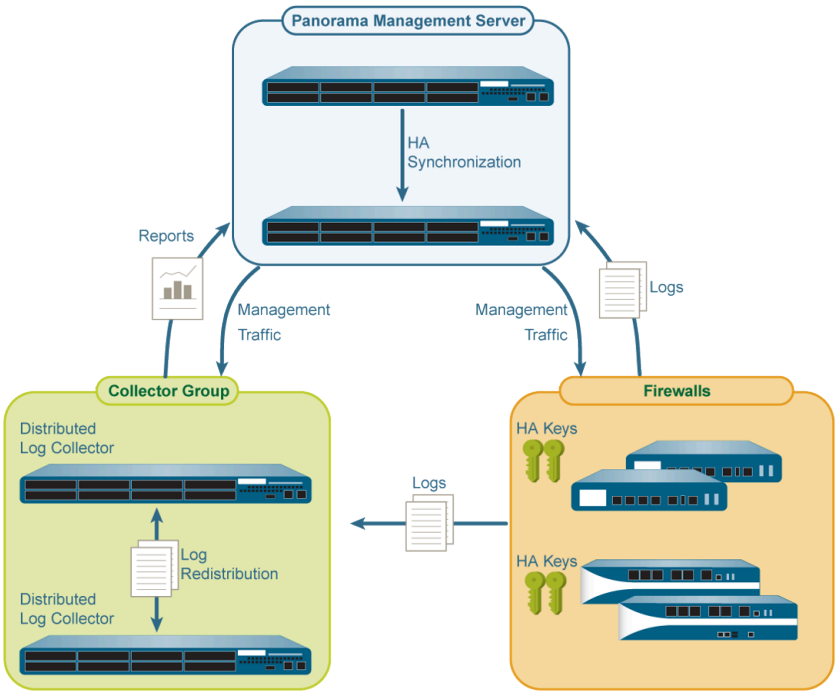


Figure 10: SSL/TLS Authentication

Configure Authentication Using Custom Certificates on Panorama

Complete the following procedure to configure the server side (Panorama) to use custom certificates instead of predefined certificates for mutual authentication with managed devices in your deployment. See Set Up Authentication Using Custom Certificates Between HA Peers to configure custom certificates on a Panorama HA pair.

STEP 1 | Deploy the server certificate.

You can deploy certificates on Panorama or a server Log Collector by generating a self-signed certificate on Panorama or obtaining a certificate from your enterprise certificate authority (CA) or a trusted third-party CA.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | On Panorama, configure a certificate profile This certificate profile defines what certificate to use and what certificate field to look for the IP address or FQDN in.

1. Select **Panorama** > **Certificate Management** > **Certificate Profile**.
2. Configure a certificate profile.

   If you configure an intermediate CA as part of the certificate profile, you must include the root CA as well.

STEP 3 | Configure an SSL/TLS service profile.

1. Select **Panorama** > **Certificate Management** > **SSL/TLS Service Profile**.
2. Configure an SSL/TLS profile to define the certificate and protocol that Panorama and its managed devices use for SSL/TLS services.

STEP 4 | Configure Secure Server Communication on Panorama or a Log Collector in the server role.

1. Select one of the following navigation paths:
   - For Panorama: **Panorama** > **Setup** > **Management** and **Edit** the Secure Communications Settings
   - For a Log Collector: **Panorama** > **Managed Collectors** > **Add** > **Communication**
2. Select the **Customize Secure Server Communication** option.
3. Verify that the **Allow Custom Certificate Only** check box is not selected. This allows you to continue managing all devices while migrating to custom certificates.

   When the Custom Certificate Only check box is selected, Panorama does not authenticate and cannot manage devices using predefined certificates.
4. Select the **SSL/TLS Service Profile**. This SSL/TLS service profile applies to all SSL connections between Panorama, firewalls, Log Collectors, and Panorama HA peers.
5. Select the **Certificate Profile** that identifies the certificate to use to establish secure communication with clients such as firewalls.
6. (Optional) Configure an authorization list. The authorization list adds an additional layer of security beyond certificate authentication. The authorization list checks the client certificate Subject or Subject Alt Name. If the Subject or Subject Alt Name

 ©2025 Palo Alto Networks, Inc.

presented with the client certificate does not match an identifier on the authorization list, authentication is denied.

You can also authorize client devices based on their serial number.

1. **Add** an Authorization List.
2. Select the **Subject** or **Subject Alt Name** configured in the certificate profile as the Identifier type.
3. Enter the Common Name if the identifier is Subject or and IP address, hostname or email if the identifier is Subject Alt Name.
4. Click **OK**.
5. Select **Check Authorization List** to enforce the authorization list.

7. Select **Authorize Client Based on Serial Number** to have the server authenticate client based on the serial numbers of managed devices. The CN or subject in the client certificate must have the special keyword $UDID to enable this type of authentication.
8. Select the **Data Redistribution** option in the **Customize Communication** section to use a custom certificate to secure outgoing communication with data redistribute clients.
9. In **Disconnect Wait Time (min)**, specify how long Panorama should wait before terminating the current session and reestablishing the connection with its managed devices. This field is blank by default and the range is 0 to 44,640 minutes. Leaving this field blank is the same as setting it to 0.

   The disconnect wait time does not begin counting down until you commit the new configuration.

10. Click **OK**.
11. **Commit** your changes.

Configure Authentication Using Custom Certificates on Managed Devices

Complete the following procedure to configure the client side (firewall or Log Collector) to use custom certificates instead of predefined certificates for mutual authentication with managed devices in your deployment.

STEP 1 | Upgrade each managed firewall or Log Collector. All managed devices must be running PAN-OS 8.0 or later to enforce custom certificate authentication.

Upgrade the firewall. After upgrade, each firewall connects to Panorama using the default predefined certificates.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Obtain or generate the device certificate.

You can deploy certificates on Panorama or a server Log Collector by generating a self-signed certificate on Panorama or obtaining a certificate from your enterprise certificate authority (CA) or a trusted third-party CA.

Set the common name to $UDID or subject to CN=$UDID (in the SCEP profile) if authorizing client devices based on serial number.

- You can generate a self-signed certificate on Panorama or obtain a certificate from your enterprise CA or a trusted third-party CA.
- If you are using SCEP for the device certificate, configure a SCEP profile. SCEP allows you to automatically deploy certificates to managed devices. When a new client devices with a SCEP profile attempts to authenticate with Panorama, the certificate is sent by the SCEP server to the device.

STEP 3 | Configure the certificate profile for the client device.

You can configure this on each client device individually or you can push this configuration to the managed device as part of a template.

1. Select one of the following navigation paths:
   - For firewalls—Select **Device** > **Certificate Management** > **Certificate Profile**.
   - For Log Collectors—Select **Panorama** > **Certificate Management** > **Certificate Profile**.
2. Configure the certificate profile.

STEP 4 | Deploy custom certificates on each firewall or Log Collector.

1. Select one of the following navigation paths:
   - For firewalls: Select **Device** > **Setup** > **Management** and **Edit** the Panorama Settings
   - For Log Collectors: Select **Panorama** > **Managed Collectors** and **Add** a new Log Collector or select an existing one. Select **Communication**.
2. Select the **Secure Client Communication** check box (firewall only).
3. Select the **Certificate Type**.
   - If you are using a local device certificate, select the **Certificate** and **Certificate Profile**.
   - If you are using SCEP to deploy device certificate, select the **SCEP Profile** and **Certificate Profile**.
   - If you are using the default Panorama certificate, select **Predefined**.
4. (Optional) Enable **Check Server Identity**. The firewall or Log Collector checks the CN in the server certificate against Panorama's IP address or FQDN to verify its identity.
5. Click **OK**.
6. **Commit** your changes.

   After committing your changes, the managed device does not terminate its current session with Panorama until the Disconnect Wait Time is complete.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Select the incoming communication types for which you want to use a custom certificate:

- **HA Communication**
- **WildFire Communication**
- **Data Redistribution**

STEP 6 | After deploying custom certificates on all managed devices, enforce authentication using custom certificates.

*The WildFire appliance does not currently support custom certificates. If your Panorama is managing a WildFire appliance, do not select **Allow Custom Certificates Only**.*

1. Select **Panorama** > **Setup** > **Management** and **Edit** the Panorama settings.
2. Select **Allow Custom Certificate Only**.
3. Click **OK**.
4. **Commit** your changes.

   After committing this change, all devices managed by Panorama must use custom certificates. If not, authentication between Panorama and the device fails.

Add New Client Devices

When adding a new firewall or Log Collector to Panorama, the workflow depends on whether or not these devices are configured to use custom certificates only for mutual authentication.

- If the Custom Certificates Only is not selected on Panorama, you can add the device to Panorama and then deploy the custom certificate by following the process beginning in step Configure Authentication Using Custom Certificates on Managed Devices.
- If the Custom Certificates Only is selected on Panorama, you must deploy the custom certificates on the firewall before adding it to Panorama. If not, the managed device will not be able to authnticate with Panorama. This can be done manually through the firewall web interface or through bootstrapping as part of the bootstrap.xml file.

Change Certificates

If a custom certificate in your deployment has expired or been revoked and needs to be replaced, you can complete one of the tasks below.

- Change a Server Certificate
- Change a Client Certificate
- Change a Root or Intermediate CA Certificate

Change a Server Certificate

Complete the following task to replace a server certificate.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Deploy the new server certificate.

You can deploy certificates on Panorama or a server Log Collector by generating a self-signed certificate on Panorama or obtaining a certificate from your enterprise CA or a trusted third-party CA.

STEP 2 | Change the certificate in the SSL/TLS Service Profile.

1. Select **Panorama** > **Certificate Management** > **SSL/TLS Service Profile** and select the SSL/TLS service profile.
2. Select the **Certificate**.
3. Click **OK**.

STEP 3 | Reestablish the connection between the server (Panorama or a Log Collector) and client devices.

1. Select **Panorama** > **Setup** > **Management** and **Edit** the Panorama Settings for Panorama or select **Panorama** > **Managed Collectors** > **Add** > **Communication** for a Log Collector.
2. Set the **Disconnect Wait Time**.
3. Click **OK**.
4. **Commit** your changes.

Change a Client Certificate

Complete the following task to replace a client certificate.

STEP 1 | Obtain or generate the device certificate.

You can deploy certificates on Panorama or a server Log Collector by generating a self-signed certificate on Panorama or obtaining a certificate from your enterprise CA or a trusted third-party CA.

Set the common name to $UDID or subject to CN=$UDID (in the SCEP profile) if authorizing client devices based on serial number.

- You can generate a self-signed certificate on Panorama or obtain a certificate from your enterprise CA or a trusted third-party CA.
- If you are using SCEP for the device certificate, configure a SCEP profile. SCEP allows you to automatically deploy certificates to managed devices. When a new client devices with a SCEP profile attempts to authenticate with Panorama, the certificate is sent by the SCEP server to the device.

STEP 2 | Change the certificate in the certificate profile.

1. Select **Device** > **Certificate Management** > **Certificate Profile** and select the certificate profile.
2. Under CA Certificates, **Add** the new certificate to assign to the certificate profile.
3. Click **OK**.
4. **Commit** your changes.

Change a Root or Intermediate CA Certificate

Complete the following task to replace a root or intermediate CA certificate.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Configure the server to accept predefined certificates from clients.

1. Select **Panorama** > **Setup** > **Management** and **Edit** the Panorama Settings.
2. Uncheck **Custom Certificate Only**.
3. Select **None** from the Certificate Profile drop-down.
4. Click **OK**.
5. **Commit** your changes.

STEP 2 | Deploy the new root or intermediate CA certificate.

You can deploy certificates on Panorama or a server Log Collector by generating a self-signed certificate on Panorama or obtaining a certificate from your enterprise CA or a trusted third-party CA.

STEP 3 | Update the CA certificate in the server certificate profile.

1. Select **Panorama** > **Certificate Management** > **Certificate Profile** and select the certificate profile to update.
2. **Delete** the old CA certificate.
3. **Add** the new CA Certificate.
4. Click **OK**.

STEP 4 | Generate or import the new client certificate.

1. Select **Device** > **Certificate Management** > **Certificates**.
2. Create a self-signed root CA certificate or import a certificate from your enterprise CA.

STEP 5 | Update the CA certificate in the client certificate profile.

1. Select **Device** > **Setup** > **Management** and click the **Edit** icon in Panorama Settings for a firewall or Select **Panorama** > **Managed Collectors** > **Add** > **Communication** for a Log Collector and select the certificate profile to update.
2. **Delete** the old CA certificate.
3. **Add** the new CA Certificate.
4. Click **OK**.

STEP 6 | After updating the CA certificates on all managed devices, enforce custom-certificate authentication.

1. Select **Panorama** > **Setup** > **Management** and **Edit** the Panorama Settings.
2. Select **Custom Certificate Only**.
3. Click **OK**.
4. **Commit** your changes.

   After committing this change, all devices managed by Panorama must use custom certificates. If not, authentication between Panorama and the device fails.

 ©2025 Palo Alto Networks, Inc.

Manage Firewalls

To use the Panorama™ management server for managing Palo Alto Networks firewalls, you must add the firewalls as managed devices and then assign them to device groups and to templates or template stacks. The following tasks best suit a first-time firewall deployment. Before proceeding, review Plan Your Panorama Deployment to understand the deployment options.

- Add a Firewall as a Managed Device
- Install the Device Certificate for Managed Firewalls
- Set Up Zero Touch Provisioning
- Manage Device Groups
- Manage Templates and Template Stacks
- Manage the Master Key from Panorama
- Schedule a Configuration Push to Managed Firewalls
- Redistribute Data to Managed Firewalls
- Transition a Firewall to Panorama Management
- Device Monitoring on Panorama
- Use Case: Configure Firewalls Using Panorama

To view the **Objects** and **Policies** tabs on the Panorama web interface, you must first create at least one device group. To view the **Network** and **Device** tabs, you must create at least one template. These tabs contain the options by which you configure and manage the firewalls on your network.

Add a Firewall as a Managed Device

To use a Panorama™ management server to manage your firewalls, you need to enable a connection between the firewall and the Panorama management server. To strengthen your Security posture when onboarding a new firewall, you must create a unique device registration authentication key on the Panorama management server for mutual authentication between the new firewall and the server on first connection. A successful first connection requires that you add the Panorama IP address on each firewall the server will manage, add the serial number on the server for each firewall, and specify the device registration authentication key on both the server and the firewall. When you add a firewall as a managed device, you can also associate the new firewall with a device group, template stack, collector group, and Log Collector during the initial deployment. Additionally, you have the option to automatically push the configuration to your newly added firewall when the firewall first connects to the Panorama server, which ensures that firewalls are immediately configured and ready to secure your network.

If you are adding a firewall to Panorama in a high availability (HA) configuration, the device registration authentication key is required only to add the firewall to the primary peer. Panorama in HA configuration synchronize the Certificate Authority (CA) certificate that allows the secondary peer to manage firewalls in event of HA failover.

Adding a firewall as a managed device requires that the total count of managed firewalls not exceed the device management license activated on Panorama. Select ***Panorama*** *>* ***Licenses*** *to view the* `Device Management License` *active on Panorama and the maximum number of managed firewalls supported.*

If the firewall you are attempting to add exceeds the device management license limit, the operation is blocked and you are prompted with a warning indicating that adding the firewall to Panorama management failed.

The firewall uses the Panorama management server IP address for registration with the server. The Panorama server and the firewall authenticate with each other using 2,048-bit certificates and AES-256 encrypted SSL connections for configuration management and log collection.

To configure the device registration authentication key, specify the key lifetime and the number of times you can use the authentication key to onboard new firewalls. Additionally, you can specify one or more firewall serial numbers for which the authentication key is valid.

The authentication key expires 90 days after the key lifetime expires. After 90 days, you are prompted to re-certify the authentication key to maintain its validity. If you do not re-certify, then the authentication key becomes invalid. A system log is generated each time a firewall uses the Panorama-generated authentication key. The firewall uses the authentication key to authenticate the Panorama server when it delivers the device certificate that is used for all subsequent communications.

 ©2025 Palo Alto Networks, Inc.

For firewalls running a PAN-OS 10.2 release, Panorama running PAN-OS 10.2 or later release supports onboarding firewalls running PAN-OS 10.1.3 or later release only. You cannot add a firewall running PAN-OS 10.1.2 or earlier PAN-OS 10.2 release to Panorama management if Panorama is running PAN-OS 10.2 or later release.

Panorama supports onboarding firewalls running the following releases:

- ***Panorama running PAN-OS 10.2 or later release—** Firewalls running PAN-OS 10.1.3 or later release, and firewalls running PAN-OS 10.0 or earlier PAN-OS release.*

There is no impact to firewalls already managed by Panorama on upgrade to PAN-OS 10.2.

If you are experiencing issues adding a firewall to Panorama management, you may need to recover managed device connectivity to Panorama.

STEP 1 | Set up the firewall.

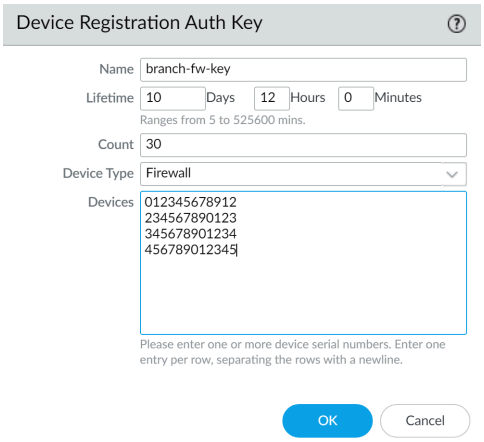
1. Perform initial configuration on the firewall so that it is accessible and can communicate with the Panorama server over the network.
2. Configure each data interface you plan to use on the firewall and attach it to a security zone so that you can push configuration settings and policy rules from the Panorama server.

 ©2025 Palo Alto Networks, Inc.

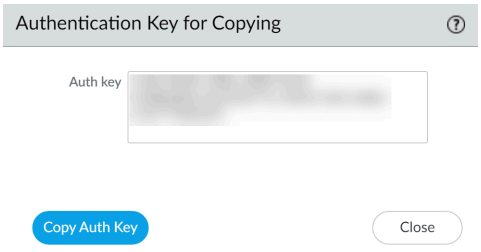
STEP 2 | Create a device registration authentication key.

1. Log in to the Panorama Web Interface.
2. Select **Panorama** > **Device Registration Auth Key** and **Add** a new authentication key.
3. Configure the authentication key.
   - **Name**—Add a descriptive name for the authentication key.
   - **Lifetime**—Specify the key lifetime to limit how long you can use the authentication key to onboard new firewalls.
   - **Count**—Specify how many times you can use the authentication key to onboard new firewalls.
   - **Device Type**—Specify that this authentication key is used to authenticate only a **Firewall**.

     *You can select **Any** to use the device registration authentication key to onboard firewalls, Log Collectors, and WildFire appliances.*
   - (Optional) **Devices**—Enter one or more device serial numbers to specify for which firewalls the authentication key is valid.
4. Click **OK**.



5. **Copy Auth Key** and **Close**.

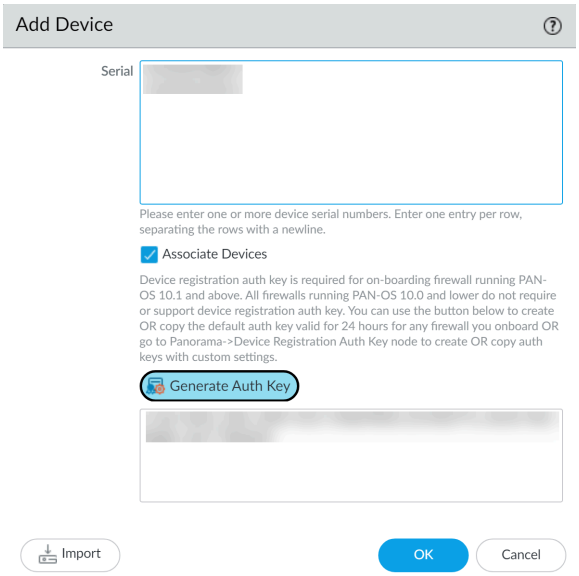


 ©2025 Palo Alto Networks, Inc.

STEP 3 | Add firewalls to a Panorama management server. You can manually add one or more firewalls or bulk import firewalls using a CSV file.

You can bulk import only single-vsys firewalls to the Panorama management server You cannot bulk import firewalls with more than one virtual system (vsys).

- Add one or more firewalls manually.
  1. Select **Panorama** > **Managed Devices** > **Summary** and **Add** a new firewall.
  2. Enter the firewall **Serial** number. If you are adding multiple firewalls, enter each serial number on a separate line.
  3. (Optional) Select **Associate Devices** to associate the firewall with a device group, template stack, Log Collector, or Collector group when the firewall first connect to the Panorama management server.
  4. Enter the device registration authentication key you created.



  5. Click **OK**.
  6. Associate your managed firewalls as needed.

     If you did not select **Associate Devices**, skip this step and continue to configure the firewall to communicate with Panorama.

     1. Assign the **Device Group**, **Template Stack**, **Collector Group**, and **Log Collector** as needed from the drop-down in each column.
     2. Enable **Auto Push on 1st connect** to automatically push the device group and template stack configuration to the new devices when they first successfully connect to the Panorama server.

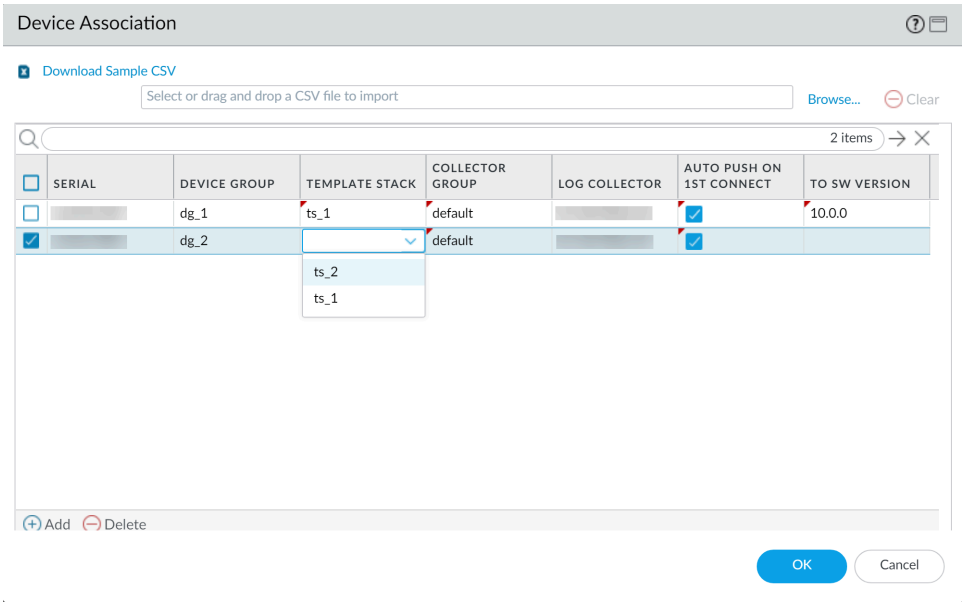
        *The **Auto Push on 1st Connect** option is supported only on firewalls running PAN-OS® 8.1 and later releases. The `commit all` job executes from Panorama to managed devices running PAN-OS 8.1 and later releases.*

 ©2025 Palo Alto Networks, Inc.

3. (Optional) Select a PAN-OS release version (**To SW Version** column) to begin automatically upgrading the managed firewall to the specified PAN-OS version upon successful connection to the Panorama management server.

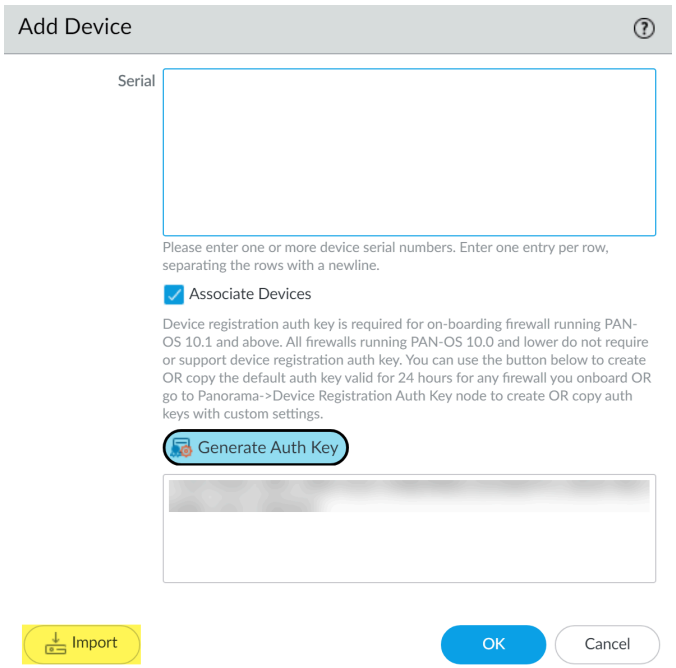
   To upgrade a managed firewall to a target PAN-OS release on first connection, you must install the minimum content release version required for that PAN-OS release before adding the firewall as a managed device. To do this, you must register the firewall, activate the support license, and install the content update before adding the firewall to Panorama management.

   Leave this column empty if you do not want to automatically upgrade the managed firewall.
4. Click **OK** to add the devices.



- Bulk import multiple firewalls using a CSV file.
  1. Select **Panorama** > **Managed Devices** > **Summary** and **Add** your new firewalls.
  2. Add the device registration authentication key you created.
  3. Click **Import**.

 ©2025 Palo Alto Networks, Inc.

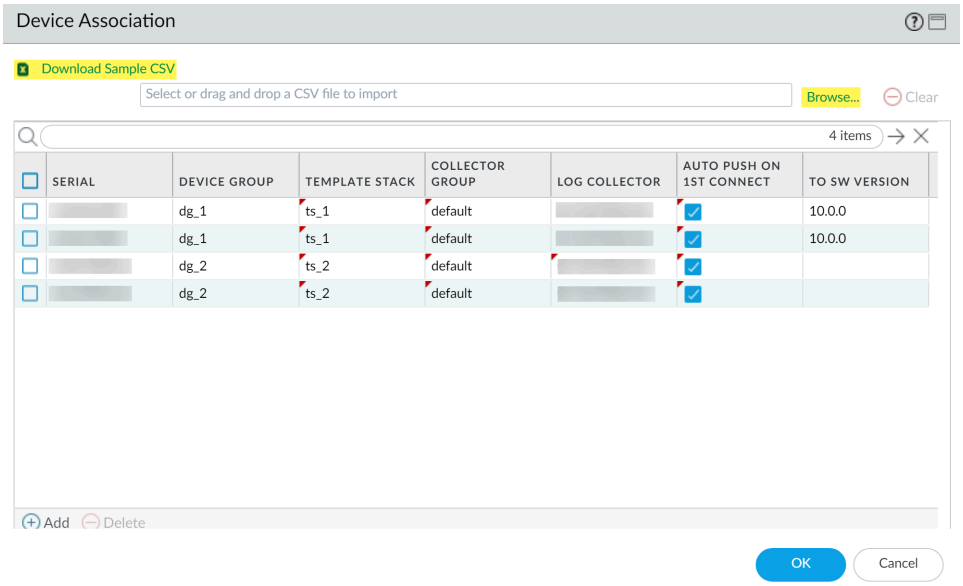


4. **Download Sample CSV** and edit the downloaded CSV file with the firewalls you are adding. You can choose to assign the firewalls to a device group, template stack,

 ©2025 Palo Alto Networks, Inc.

Collector Group, and Log Collector from the CSV or enter only the firewall serial numbers and assign them from the web interface. Save the CSV after you finish editing.

5. **Browse** to and select the CSV file you edited in the previous step.



6. If not already assigned in the CSV, assign the firewalls a **Device Group**, **Template Stack**, **Collector Group**, and **Log Collector** as needed from the drop-down in each column
7. If not already enabled in the CSV, enable **Auto Push on 1st connect** to automatically push the device group and template stack configuration to the new devices when they first successfully connect to the Panorama server.
8. (Optional) Select a PAN-OS release version (**To SW Version** column) to begin automatically upgrading the managed firewall to the specified PAN-OS version upon successful connection to the Panorama server.

   To upgrade a managed firewall to a target PAN-OS release on first connection, you must install the minimum content release version required for that PAN-OS release before adding the firewall as a managed device. To do this, you must register the firewall, activate the support license, and install the content update before adding the firewall to Panorama management.

   Leave this column empty if you do not want to automatically upgrade the managed firewall.
9. Click **OK** to add the firewalls.

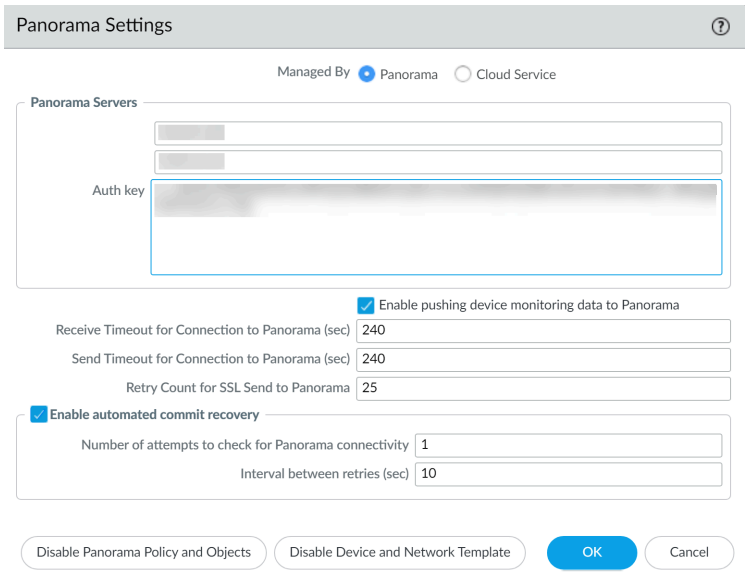
 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the firewall to communicate with the Panorama management server.

Repeat this step for each firewall the Panorama server will manage.

1. Log in to the firewall web interface.
2. Configure the Panorama Settings for the firewall.
   1. Select **Device** > **Setup** > **Management** and edit the Panorama Settings.
   2. (PAN-OS 10.2.3 and later releases only) For Managed By, select **Panorama**.
   3. Enter the Panorama IP address in the first field.

      Panorama issues a single IP address for device management, log collection, reporting, and dynamic updates. Enter the external, Internet-bound IP address to ensure Panorama can successfully access existing and new managed devices and Log Collectors. If an internal Panorama IP address is configured, you may be unable to manage some devices. For example, if you Install Panorama on AWS and enter the internal IP address, Panorama is unable to manage devices or Log Collectors outside of the AWS security group.
   4. (Optional) If you have configured a high availability (HA) pair in Panorama, enter the IP address of the secondary Panorama in the second field.
   5. Enter the **Auth key** you created on Panorama.
   6. Click **OK**.



   7. **Commit** your changes.

STEP 5 | (Optional) Add a **Tag**. Tags make it easier for you to find a firewall from a large list; they help you dynamically filter and refine the list of firewalls in your display. For example, if you add a tag called **`branch office`**, you can filter for all branch office firewalls across your network.

1. Select each firewall and click **Tag**.
2. Click **Add**, enter a string of up to 31 characters (no empty spaces), and click **OK**.

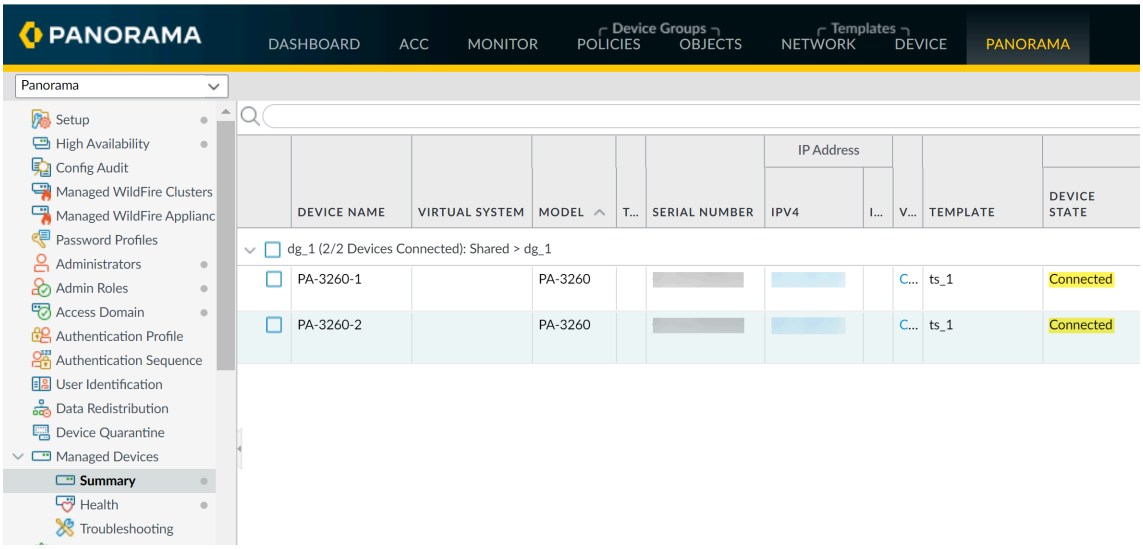
 ©2025 Palo Alto Networks, Inc.

STEP 6 | If your deployment is using custom certificates for authentication between Panorama and managed devices, deploy the custom client device certificate. For more information, see Set Up Authentication Using Custom Certificates and Add New Client Devices.

STEP 7 | Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 8 | Verify that the firewall is connected to Panorama.

1. Click **Panorama** > **Managed Devices** > **Summary**.
2. Verify that the **Device State** for the new device shows as `Connected`.



 ©2025 Palo Alto Networks, Inc.

Install the Device Certificate for Managed Firewalls

Install the device certificate on your managed firewalls to use one or more Palo Alto Networks cloud services. You can install the device certificate for a single managed firewall or multiple managed firewalls at once.

See Device Certificates to install the firewall device certificate locally.

- Install the Device Certificate for a Managed Firewall
- Install the Device Certificate for All Managed Firewalls Without a Device Certificate

Install the Device Certificate for a Managed Firewall

Where Can I Use This?	What Do I Need?
• NGFW (Managed by Panorama)	❑ Device management license ❑ Support license ❑ Outbound internet access ❑ Customer Support Portal (CSP) account with one of the following user roles: Super User, Standard User, Limited User, Threat Researcher, AutoFocus Trial Role, Group Super User, Group Standard User, Group Limited User, Group Threat Researcher, Authorized Support Center (ASC) User, and ASC Full Service User. ❑ Panorama superuser role

Select and install the device certificate for managed firewalls from the Panorama management server to use one or more cloud services. You only need to install a device certificate once. The device certificate has a 90-day lifetime. The firewall reinstalls the device certificate 15 days before the certificate expires. In the event the firewall is unable to reinstall the device certificate on its own, you may need to manually restore an expired device certificate.

To successfully install the device certificate on a firewall, the firewall must have outbound internet access and the following Fully Qualified Domain Names (FQDN) and ports must be allowed on your network in order to reach to the CSP. Additionally, the managed firewall must belong to the same CSP account as Panorama in order to generate the One Time Password (OTP) used to install the device certificate.

FQDN	Ports
• http://ocsp.paloaltonetworks.com • http://crl.paloaltonetworks.com	TCP 80

 ©2025 Palo Alto Networks, Inc.

FQDN	Ports
• http://ocsp.godaddy.com	
• https://api.paloaltonetworks.com • http://apitrusted.paloaltonetworks.com • https://certificatetrusted.paloaltonetworks.com • https://certificate.paloaltonetworks.com	TCP 443
• *.gpcloudservice.com	TCP 444 and TCP 443

The following Palo Alto Networks Next-Generation firewall models install the device certificate when they first connect to the Palo Alto Networks CSP during the initial registration process. You do not need to manually install the device certificate for these firewall models.

- *PA-410, PA-440, PA-450, and PA-460 firewalls*
- *PA-1400 Series firewalls*
- *PA-3400 Series firewalls*
- *PA-5410, PA-5420, and PA-5430 firewalls*
- *PA-5450 firewall*

STEP 1 | Log in to the Panorama Web Interface as a Superuser.

A Panorama admin with Superuser access privileges is required to generate OTP Request Token and apply the OTP used to install the device certificate on a managed firewall.

STEP 2 | (Best Practices) Configure the Network Time Protocol (NTP) server for Panorama.

An NTP server is required to validate the device certification expiration date, ensure the device certificate does not expire early or become invalid.

1. Select **Panorama** > **Setup** > **Services**.
2. Select **NTP** and enter the hostname or IP address of the **Primary NTP Server**.
3. (Optional) Enter a the hostname or IP address of the **Secondary NTP Server**.
4. (Optional) To authenticate time updates from the NTP server(s), for **Authentication Type**, select one of the following for each server.
   - **None** (default)—Disables NTP authentication.
   - **Symmetric Key**—Firewall uses symmetric key exchange (shared secrets) to authenticate time updates.
     - **Key ID**—Enter the Key ID (1-65534)
     - **Algorithm**—Select the algorithm to use in NTP authentication (**MD5** or **SHA1**)
5. Click **OK** to save your configuration changes.
6. Select **Commit** and **Commit to Panorama**.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Configure the Network Time Protocol (NTP) server for your firewalls

An NTP server is required to validate the device certification expiration date, ensure the device certificate does not expire early or become invalid.

1. Select **Device** > **Setup** > **Services** and select the **Template**.
2. Select one of the following depending on your platform:
   - For multi-virtual system platforms, select **Global** and edit the Services section.
   - For single virtual system platforms, edit the Services section.
3. Select **NTP** and enter the hostname or IP address of the **Primary NTP Server**.
4. (Optional) Enter a the hostname or IP address of the **Secondary NTP Server**.
5. (Optional) To authenticate time updates from the NTP server(s), for **Authentication Type**, select one of the following for each server.
   - **None** (default)—Disables NTP authentication.
   - **Symmetric Key**—Firewall uses symmetric key exchange (shared secrets) to authenticate time updates.
     - **Key ID**—Enter the Key ID (1-65534)
     - **Algorithm**—Select the algorithm to use in NTP authentication (**MDS** or **SHA1**)
6. Click **OK** to save your configuration changes.
7. Select **Commit** and **Commit and Push** your configuration changes to your managed firewalls.

STEP 4 | Generate the OTP Request Token on Panorama.

The OTP Request Token generated on Panorama is used to generate the OTP required to install the device certificate on a managed firewall.

1. Select **Panorama** > **Managed Devices** > **Summary**.
2. Select one or more managed firewalls.
3. Select **Request OTP From CSP** > **Custom selected devices**.
4. **Copy** the entire output in the OTP Request Token field.

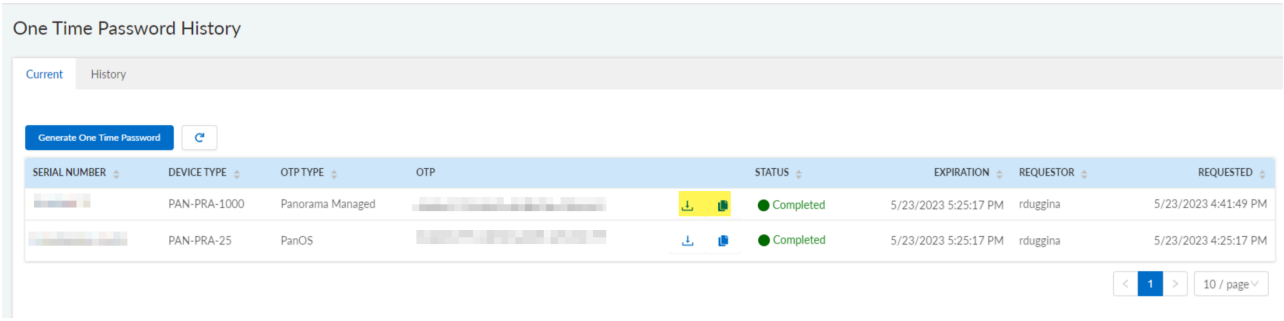
 ©2025 Palo Alto Networks, Inc.

STEP 5 | Generate the One Time Password (OTP) for managed firewalls.

OTP lifetime is 60 minutes and expires if not used within the 60 minute lifetime.

Firewall may only attempt to retrieve the OTP from the CSP one time. If the firewall fails for any reason to fetch the OTP, the OTP expires and you must generate a new OTP.

1. Log in to the Customer Support Portal with a user role that has permission to generate an OTP.
2. Select **Products** > **Device Certificates** and **Generate OTP**.
3. For the **Device Type**, select **Generate OTP for Panorama managed firewalls** and click **Next**.
4. Paste the OTP request you copied in the previous step and **Generate OTP**.
5. Click **Done** and wait a few minutes for the OTP to successfully generate.
6. **View OTP History**.
7. In the **Current** One Time Password History, copy or download the OTP
8. **Copy to Clipboard** or **Download** the OTP.



STEP 6 | Install the device certificate on managed firewalls.

Managed firewalls must have an outbound internet connection to successfully install the device certificate. After you upload the OTP from Panorama, each managed firewall connects to the Palo Alto Networks CSP to install the device certificate.

1. Log in to the Panorama web interface as a Superuser user.
2. Select **Panorama** > **Managed Devices** > **Summary** and **Upload OTP**.
3. Paste the OTP you generated and **Upload**.

You must still copy and paste the OTP generated from the Palo Alto Networks CSP even if you downloaded the OTP in the previous step. Uploading the file containing the OTP is not supported.

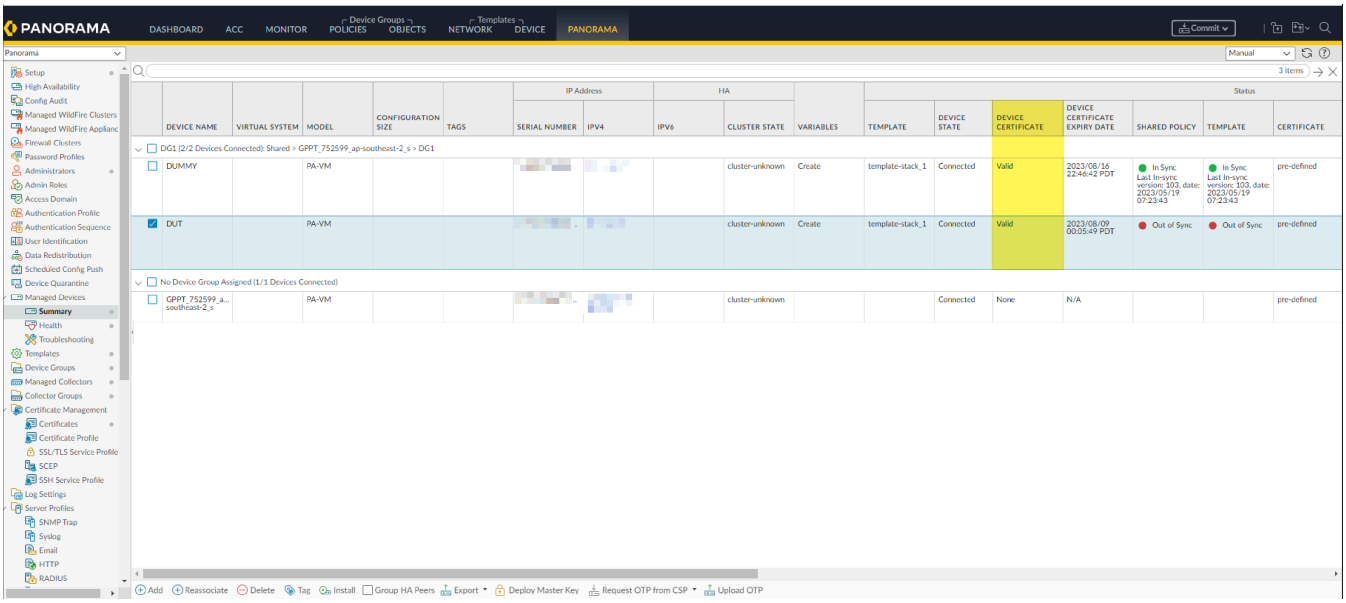
 ©2025 Palo Alto Networks, Inc.

STEP 7 | (WildFire and Advanced WildFire) Log in to the firewall CLI and refresh the firewall settings to establish a connection to the Advanced WildFire cloud with the updated device certificate.

Repeat this step for each managed firewall with an activate WildFire or Advanced WildFire subscription that is actively communicating with the Advanced WildFire cloud service.

```
admin>request wildfire registration channel public
```

STEP 8 | Verify that the **Device Certificate** column displays as `Valid` and that the **Device Certificate Expiry Date** displays an expiration date.



Install the Device Certificate for All Managed Firewalls Without a Device Certificate

Where Can I Use This?	What Do I Need?
• NGFW (Managed by Panorama)	☐ Device management license ☐ Support license ☐ Outbound internet access ☐ Customer Support Portal (CSP) account with one of the following user roles: Super User, Standard User, Limited User, Threat Researcher, AutoFocus Trial Role, Group Super User, Group Standard User, Group Limited User, Group Threat Researcher, Authorized Support Center (ASC) User, and ASC Full Service User. ☐ Panorama superuser role

 ©2025 Palo Alto Networks, Inc.

Install the device certificate for managed firewalls that do not already have a device certificate installed from the Panorama management server. The device certificate is required to successfully authenticate the managed firewall with the Palo Alto Networks CSP to leverage one or more cloud services. The device certificate has a 90-day lifetime. The firewall reinstalls the device certificate 15 days before the certificate expires. In the event the firewall is unable to reinstall the device certificate on its own, you may need to manually restore an expired device certificate.

To successfully install the device certificate for a managed firewall, managed firewalls must have outbound internet access and the following Fully Qualified Domain Names (FQDN) and ports must be allowed on your network. Additionally, the managed firewall must belong to the same CSP account as Panorama in order to generate the One Time Password (OTP) used to install the device certificate.

FQDN	Ports
• http://ocsp.paloaltonetworks.com • http://crl.paloaltonetworks.com • http://ocsp.godaddy.com	TCP 80
• https://api.paloaltonetworks.com • http://apitrusted.paloaltonetworks.com • https://certificatetrusted.paloaltonetworks.com • https://certificate.paloaltonetworks.com	TCP 443
• *.gpcloudservice.com	TCP 444 and TCP 443

The following Palo Alto Networks Next-Generation firewall models install the device certificate when they first connect to the Palo Alto Networks CSP during the initial registration process. You do not need to manually install the device certificate for these firewall models.

- *PA-410, PA-440, PA-450, and PA-460 firewalls*
- *PA-1400 Series firewalls*
- *PA-3400 Series firewalls*
- *PA-5410, PA-5420, and PA-5430 firewalls*
- *PA-5450 firewall*

STEP 1 | Log in to the Panorama Web Interface as a Superuser.

A Panorama admin with Superuser access privileges is required to generate OTP Request Token and apply the OTP used to install the device certificate on managed firewalls.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | (Best Practices) Configure the Network Time Protocol (NTP) server for Panorama.

An NTP server is required to validate the device certification expiration date, ensure the device certificate does not expire early or become invalid.

1. Select **Panorama** > **Setup** > **Services**.
2. Select **NTP** and enter the hostname or IP address of the **Primary NTP Server**.
3. (Optional) Enter a the hostname or IP address of the **Secondary NTP Server**.
4. (Optional) To authenticate time updates from the NTP server(s), for **Authentication Type**, select one of the following for each server.
   - **None** (default)—Disables NTP authentication.
   - **Symmetric Key**—Firewall uses symmetric key exchange (shared secrets) to authenticate time updates.
     - **Key ID**—Enter the Key ID (1-65534)
     - **Algorithm**—Select the algorithm to use in NTP authentication (**MDS** or **SHA1**)
5. Click **OK** to save your configuration changes.
6. Select **Commit** and **Commit to Panorama**.

STEP 3 | Configure the Network Time Protocol (NTP) server.

An NTP server is required to validate the device certification expiration date, ensure the device certificate does not expire early or become invalid.

1. Select **Device** > **Setup** > **Services** and select the **Template**.
2. Select one of the following depending on your platform:
   - For multi-virtual system platforms, select **Global** and edit the Services section.
   - For single virtual system platforms, edit the Services section.
3. Select **NTP** and enter the hostname or IP address of the **Primary NTP Server**.
4. (Optional) Enter a the hostname or IP address of the **Secondary NTP Server**.
5. (Optional) To authenticate time updates from the NTP server(s), for **Authentication Type**, select one of the following for each server.
   - **None** (default)—Disables NTP authentication.
   - **Symmetric Key**—Firewall uses symmetric key exchange (shared secrets) to authenticate time updates.
     - **Key ID**—Enter the Key ID (1-65534)
     - **Algorithm**—Select the algorithm to use in NTP authentication (**MDS** or **SHA1**)
6. Click **OK** to save your configuration changes.
7. Select **Commit** and **Commit and Push** your configuration changes to your managed firewalls.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Generate the OTP Request Token on Panorama.

The OTP Request Token generated on Panorama is used to generate the OTP required to install the device certificate on managed firewalls.

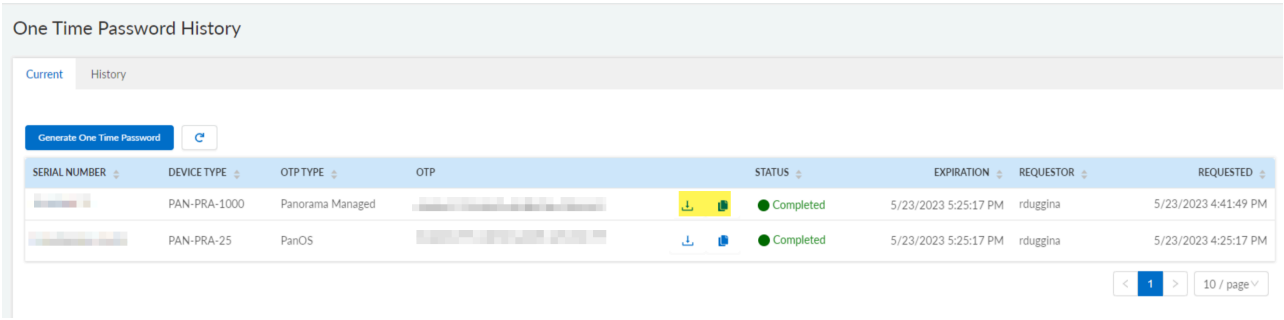
1. Select **Panorama** > **Managed Devices** > **Summary**.
2. Select **Request OTP From CSP** > **Select all devices without a certificate**.
3. **Copy** the entire output in the OTP Request Token field.

STEP 5 | Generate the One Time Password (OTP) for managed firewalls.

OTP lifetime is 60 minutes and expires if not used within the 60 minute lifetime.

Firewall may only attempt to retrieve the OTP from the CSP one time. If the firewall fails for any reason to fetch the OTP, the OTP expires and you must generate a new OTP.

1. Log in to the Customer Support Portal with a user role that has permission to generate an OTP.
2. Select **Products** > **Device Certificates** and **Generate OTP**.
3. For the **Device Type**, select **Generate OTP for Panorama managed firewalls** and click **Next**.
4. Paste the OTP request you copied in the previous step and **Generate OTP**.
5. Click **Done** and wait a few minutes for the OTP to successfully generate.
6. **View OTP History**.
7. In the **Current** One Time Password History, copy or download the OTP
8. **Copy to Clipboard** or **Download** the OTP.



 ©2025 Palo Alto Networks, Inc.

STEP 6 | Install the device certificate on managed firewalls.

Managed firewalls must have an outbound internet connection to successfully install the device certificate. After you upload the OTP from Panorama, each managed firewall connects to the Palo Alto Networks CSP to install the device certificate.

1. Log in to the Panorama web interface as a Superuser user.
2. Select **Panorama** > **Managed Devices** > **Summary** and **Upload OTP**.
3. Paste the OTP you generated and **Upload**.

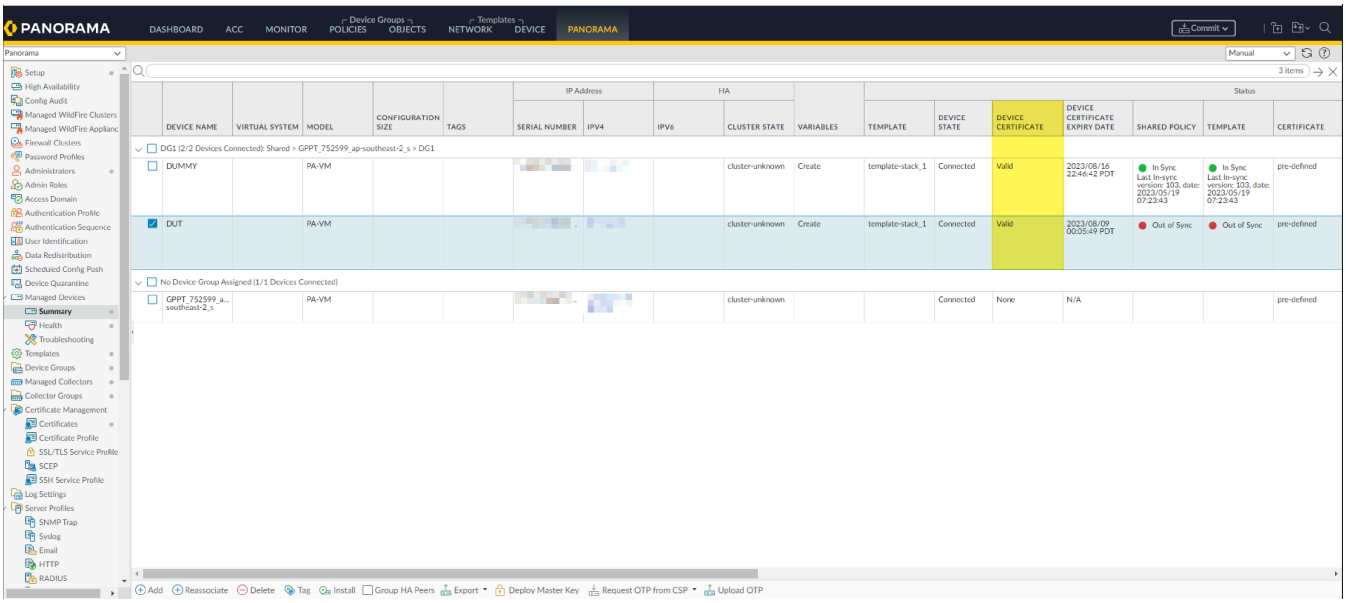
You must still copy and paste the OTP generated from the Palo Alto Networks CSP even if you downloaded the OTP in the previous step. Uploading the file containing the OTP is not supported.

STEP 7 | (WildFire and Advanced WildFire) Log in to the firewall CLI and refresh the firewall settings to establish a connection to the Advanced WildFire cloud with the updated device certificate.

Repeat this step for each managed firewall with an activate WildFire or Advanced WildFire subscription that is actively communicating with the Advanced WildFire cloud service.

```
admin>request wildfire registration channel public
```

STEP 8 | Verify that the **Device Certificate** column displays as `Valid` and that the **Device Certificate Expiry Date** displays an expiration date.



 ©2025 Palo Alto Networks, Inc.

Set Up Zero Touch Provisioning

Set up Zero Touch Provisioning (ZTP) to simplify and streamline initial firewall deployments by automating the new managed firewall on-boarding without the need for network administrators to manually provision the firewall.

ZTP onboarding requires on the ZTP firewall, you cable the Eth1/1 interface with an outbound internet connection before the ZTP firewall is powered on. This is required to successfully onboard the ZTP firewall to Panorama management, register your ZTP firewall with the CSP, and push the policy and network configurations from Panorama.

Only Panorama administrators with Superuser privileges can access the ZTP settings required to set up ZTP.

To successfully leverage the ZTP service, on-board your ZTP firewalls with the factory default PAN-OS version before upgrading to PAN-OS 10.0.0 or later release.

The ZTP plugin is supported on PAN-OS 9.1.4 and later releases.

- ZTP Overview
- Install the ZTP Plugin
- Configure the ZTP Installer Administrator Account
- Add ZTP Firewalls to Panorama
- Use the CLI for ZTP Tasks
- Uninstall the ZTP Plugin

ZTP Overview

Learn more about Zero Touch Provisioning (ZTP) and its configuration elements.

- About ZTP
- ZTP Configuration Elements

About ZTP

Zero Touch Provisioning (ZTP) is designed to simplify and automate the on-boarding of new firewalls to the Panorama™ management server. ZTP streamlines the initial firewall deployment process by allowing network administrators to ship managed firewalls directly to their branches and automatically add the firewall to the Panorama™ management server after the ZTP firewall successfully connects to the Palo Alto Networks ZTP service. This allows businesses to save on time and resources when deploying new firewalls at branch locations by removing the need for IT administrators to manually provision the new managed firewall. After successful on-boarding, Panorama provides the means to configure and manage your ZTP configuration and firewalls.

The ZTP cloud service supports a direct internet connection to successfully onboard a ZTP firewall to Panorama management. The ZTP cloud service does not support an explicit web proxy and is unable to onboard a ZTP firewall to Panorama management if an explicit web proxy is configured as a gateway to the internet for your ZTP firewalls and Panorama.

 ©2025 Palo Alto Networks, Inc.

Review and subscribe to ZTP Service Status events to be notified about scheduled maintenance windows, outages, and workarounds.

ZTP is supported on the following ZTP firewalls:

- PA-220-ZTP and PA-220R-ZTP
- PA-410, PA-440, PA-450, and PA-460
- PA-820-ZTP and PA-850-ZTP
- PA-3220-ZTP, PA-3250-ZTP, and PA-3260-ZTP
- PA-3400 Series Firewalls
- PA-5410, PA-5420, and PA-5430
- PA-5450

Before you begin setting up ZTP on Panorama, review the Firewall Hardware Quick Start and Reference Guides to understand how to correctly install your firewall to successfully leverage ZTP.

ZTP Configuration Elements

The following elements work together to allow you to quickly on-board newly deployed ZTP firewalls by automatically adding them to the Panorama management server using the ZTP service.

- **ZTP Plugin**—The ZTP plugin allows Panorama to connect to the ZTP service and claim a ZTP firewall for simplified on-boarding.
- **Customer Support Portal (CSP)**—The Palo Alto Networks Customer Support Portal is used to register your Panorama to connect to the CSP to automatically register newly added ZTP firewalls.
- **One-time Password (OTP)**—A one-time password provided by Palo Alto Networks used to retrieve and install a certificate on Panorama for it to communicate with the CSP and ZTP service.
- **Installer**—An administrator user created using the `installeradmin` admin role for ZTP firewall on-boarding. This admin user has limited access to the Panorama web interface, only allowing access to enter the ZTP firewall serial number and claim key to register firewalls on the CSP and Panorama. The installer admin can be created on Panorama or created using remote authentication such as RADIUS, SAML, or TACACS+.
- **Claim Key**—Eight digit numeric key physically attached to the ZTP firewall used to register the ZTP firewall with the CSP.
- **To-SW-Version**—Designate the PAN-OS software version of the ZTP firewall (**Panorama** > **Managed Devices** > **Summary**). Select the target PAN-OS release, and if the firewall is running an earlier release than the indicated version, the firewall begins an upgrade loop until the target release is successfully installed.

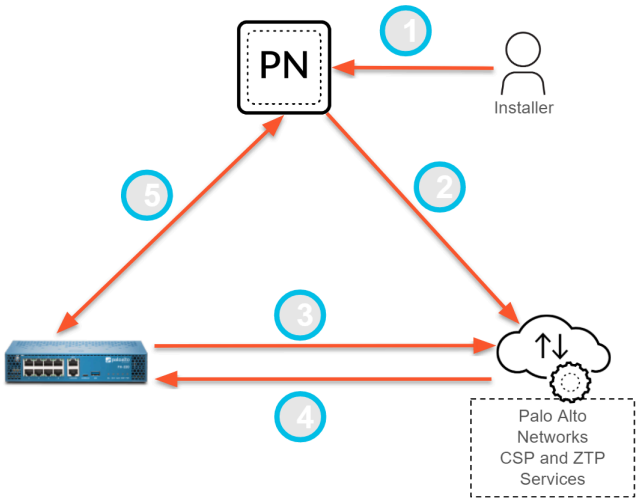
  Panorama can only manage firewalls running a PAN-OS release equal to or less than that installed on the Panorama.

After you successfully install the ZTP plugin on Panorama and register Panorama with the ZTP service, the ZTP on boarding process continues as follows:

 ©2025 Palo Alto Networks, Inc.

1. Installer or IT administrator registers ZTP firewalls by adding them to Panorama using the firewall serial number and claim key.
2. Panorama registers the firewalls with the CSP. After the firewalls are successfully registered, the firewall is associated with the same ZTP tenant as the Panorama in the ZTP service.

   ZTP firewalls successfully registered with the ZTP service are automatically added as managed firewalls (**Panorama** > **Managed Devices**) on Panorama.
3. When the firewall connects to the Internet, the ZTP firewall requests a device certificate from the CSP in order to connect to the ZTP service.
4. The ZTP service pushes the Panorama IP or FQDN to the ZTP firewalls.
5. The ZTP firewalls connect to Panorama and the device group and template configurations are pushed from Panorama to the ZTP firewalls.



Install the ZTP Plugin

Install the ZTP plugin on your Panorama™ management server to register Panorama with the ZTP service in order to claim ZTP firewalls for simplified on-boarding.

If your Panorama is in a high availability (HA) configuration, install the ZTP plugin and register both Panorama HA peers with the ZTP service.

- Install the ZTP Plugin on Panorama
- Register Panorama with the ZTP Service

Install the ZTP Plugin on Panorama

Simplify the on-boarding and management of ZTP firewalls by installing the ZTP plugin on your Panorama management server.

STEP 1 | Install the Panorama Device Certificate.

STEP 2 | Log in to the Panorama web interface as a superuser or Panorama administrator with access to Panorama plugins (**Panorama** > **Plugins**).

STEP 3 | Select **Panorama** > **Plugins** and search for the **`ztp`** plugin.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | **Download** and **Install** the most recent version of the ZTP plugin.

Register Panorama with the ZTP Service

Register the Panorama™ management server with the ZTP service for new and existing deployments.

- Register Panorama with the ZTP Service for New Deployments
- Register Panorama with the ZTP Service for Existing Deployments

Register Panorama with the ZTP Service for New Deployments

After you install the ZTP plugin on the Panorama™ management server, you must register the Panorama with the ZTP service to enable the ZTP service to associate firewalls with the Panorama. As part of the registration process for ZTP new deployment, automatically generate the device group and template configurations required to connect your ZTP firewalls to the ZTP service. After the device group and template are automatically generated, you must add your ZTP firewalls to the device group and template so they can connect to the ZTP service after they first connect to Panorama.

STEP 1 | Install the Panorama Device Certificate.

STEP 2 | Log in to the Palo Alto Networks Customer Support Portal (CSP).

STEP 3 | Associate your Panorama with the ZTP Service on the Palo Alto Networks CSP.

The ZTP Service supports associating up to two Panoramas only if they are in a high availability (HA) configuration. If Panorama is not in an HA configuration, only a single Panorama can be associated.

1. Select **Assets** > **ZTP Service** and **Associate Panorama(s)**.
2. Select the serial number of the Panorama managing your ZTP firewalls.
3. (HA only) Select the serial number of the Panorama HA peer.
4. Click **OK**.

STEP 4 | Log in to the Panorama Web Interface.

STEP 5 | Select **Panorama** > **Zero Touch Provisioning** > **Setup** and edit the **General** ZTP settings.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Register Panorama with the ZTP service.

1. **Enable ZTP Service**.
2. Enter the **Panorama FQDN or IP Address**.

   This is the FQDN or public IP address of the Panorama the ZTP plugin is installed on and that the CSP pushes to the ZTP firewalls.

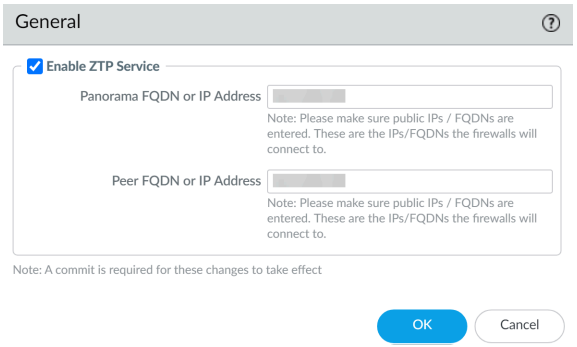
   (Managed firewalls running PAN-OS 10.1.4 and earlier releases) Enter the Panorama IP address to avoid the managed firewall disconnecting from Panorama on reboot or after a successful PAN-OS upgrade.

   If you need to use the Panorama FQDN, configure a static destination route to avoid the managed firewall disconnecting from Panorama on reboot or after a successful PAN-OS upgrade.
3. (HA only) Enter the **Peer FQDN or IP Address**.

   This is the FQDN or public IP address of the Panorama peer on which the ZTP plugin is installed and that the CSP pushes to the ZTP firewalls in case of failover.

   (Managed firewalls running PAN-OS 10.1.4 and earlier releases) Enter the Panorama IP address to avoid the managed firewall disconnecting from Panorama on reboot or after a successful PAN-OS upgrade.

   If you need to use the Panorama FQDN, configure a static destination route to avoid the managed firewall disconnecting from Panorama on reboot or after a successful PAN-OS upgrade.
4. Click **OK** to save your configuration changes.



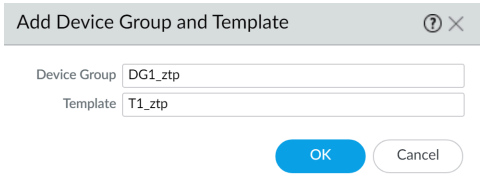
 ©2025 Palo Alto Networks, Inc.

STEP 7 | Create the default device group and template to automatically generate the required configuration to connect your ZTP firewalls to Panorama.

Adding the device group and template automatically generates a new device group and template that contain the default configuration to connect the Panorama and the ZTP firewalls.

Palo Alto Networks recommends giving the ZTP device group and template a descriptive name that makes their purpose clear. Unintentionally modifying the default ZTP configuration results in connectivity issues if you want to re-use the device group and template to onboard new ZTP firewalls in the future.

1. **Add Device Group and Template**.
2. Enter the **Device Group** name.
3. Enter the **Template** name.
4. Click **OK** to save your configuration changes.



STEP 8 | Modify the ZTP device group, templates, and template stack as needed.

Moving a ZTP firewall to a different device group or template stack is not supported. You must keep the ZTP firewalls in the ZTP device group and template stack that includes the ZTP template that were created. This is required for the firewall to maintain connectivity with Panorama and prevent any unintended configuration reverts on the firewall.

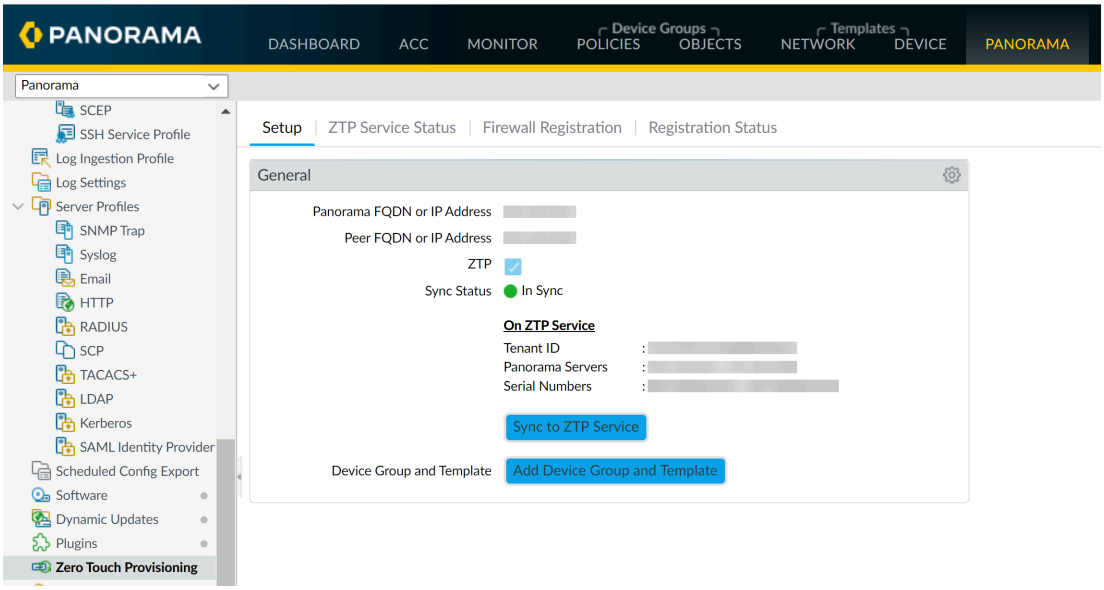
When considering your device group hierarchy and template priority in your template stack, ensure that the device group and template containing the required ZTP configuration that allows the ZTP firewall and Panorama to communicate have priority such that the configuration is not overridden in the event of conflicting configurations.

If modifying the ZTP device group and template used to onboard the ZTP firewall, be careful to not modify any of the ZTP configuration that was automatically populated when you created the device group and template in the previous step. This includes configurations like the Panorama IP address, virtual router, the `ethernet1/1` interface, Security zone of the `ethernet1/1` interface, the `loopback.900` loopback interface, the `rule1` Security policy rule, `ztp-nat` NAT policy rule, and the service route. These configurations are required to connect your ZTP firewall to Panorama and can lead to connectivity issues if modified.

STEP 9 | Select **Commit** and **Commit to Panorama**

 ©2025 Palo Alto Networks, Inc.

STEP 10 | **Sync to ZTP Service** and verify that the Panorama Sync Status displays as `In Sync`.



STEP 11 | Add ZTP Firewalls to Panorama.

Register Panorama with the ZTP Service for Existing Deployments

After you install the ZTP plugin on the Panorama™ management server, you must register Panorama with the ZTP service to enable the ZTP service to associate firewalls with the Panorama. As part of the registration process, add your ZTP firewalls to the existing ZTP device group and template stack that contain the required ZTP configuration to connect your ZTP firewalls with the ZTP service after they first connect to Panorama.

STEP 1 | Install the Panorama Device Certificate.

STEP 2 | Log in to the Palo Alto Networks Customer Support Portal (CSP).

STEP 3 | Associate your Panorama with the ZTP Service on the Palo Alto Networks CSP.

The ZTP Service supports associating up to two Panoramas only if they are in a high availability (HA) configuration. If Panorama is not in an HA configuration, only a single Panorama can be associated.

1. Select **Assets** > **ZTP Service** and **Modify Association**.
2. Select the serial number of the Panorama managing your ZTP firewalls.
3. (HA only) Select the serial number of the Panorama HA peer.
4. Click **OK**.

STEP 4 | Log in to the Panorama Web Interface.

STEP 5 | Select **Panorama** > **Zero Touch Provisioning** > **Setup** and edit the **General** ZTP settings.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Register Panorama with the ZTP service.

1. **Enable ZTP Service**.
2. Enter the **Panorama FQDN or IP Address**.

   This is the FQDN or public IP address of the Panorama the ZTP plugin is installed on and that the CSP pushes to the ZTP firewalls.

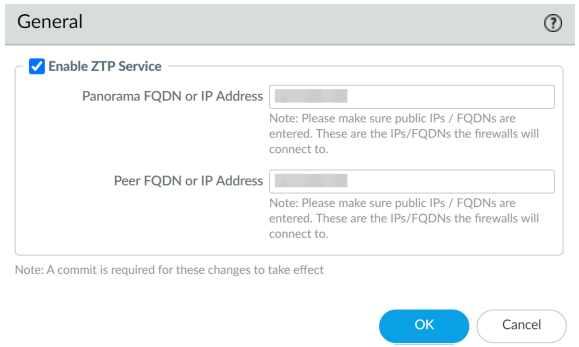
   (Managed firewalls running PAN-OS 10.1.4 and earlier releases) Enter the Panorama IP address to avoid the managed firewall disconnecting from Panorama on reboot or after a successful PAN-OS upgrade.

   If you need to use the Panorama FQDN, configure a static destination route to avoid the managed firewall disconnecting from Panorama on reboot or after a successful PAN-OS upgrade.
3. (HA only) Enter the **Peer FQDN or IP Address**.

   This is the FQDN or public IP address of the Panorama peer on which the ZTP plugin is installed and that the CSP pushes to the ZTP firewalls in case of failover.

   (Managed firewalls running PAN-OS 10.1.4 and earlier releases) Enter the Panorama IP address to avoid the managed firewall disconnecting from Panorama on reboot or after a successful PAN-OS upgrade.

   If you need to use the Panorama FQDN, configure a static destination route to avoid the managed firewall disconnecting from Panorama on reboot or after a successful PAN-OS upgrade.
4. Click **OK** to save your configuration changes.



STEP 7 | Modify the ZTP device group, templates, and template stack as needed.

Moving a ZTP firewall to a different device group or template stack is not supported. You must keep the ZTP onboarded firewalls in the ZTP device group and templates that were created. This is required for the firewall to maintain connectivity with Panorama and prevent any unintended configuration reverts on the firewall.

When considering your device group hierarchy and template priority in your template stack, ensure that the device group and template containing the required ZTP configuration

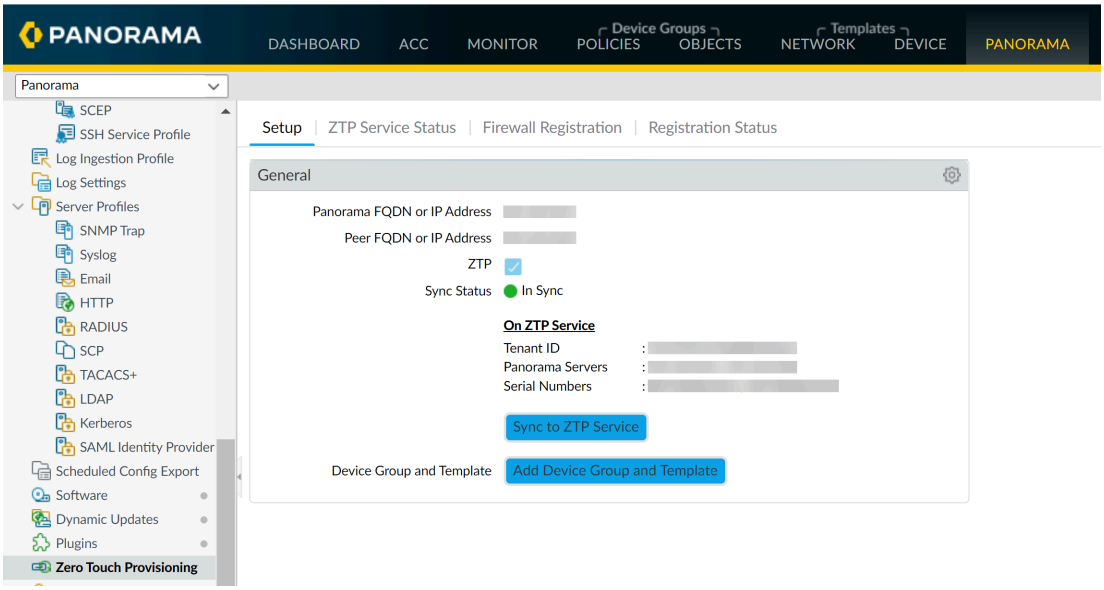
 ©2025 Palo Alto Networks, Inc.

that allows the ZTP firewall and Panorama to communicate have priority such that the configuration is not overridden in the event of conflicting configurations.

If modifying the ZTP device group and template used to onboard the ZTP firewall, be careful to not modify any of the ZTP configuration that was automatically populated when you created the device group and template in the previous step. This includes configurations like the Panorama IP address, virtual router, the `ethernet1/1` interface, Security zone of the `ethernet1/1` interface, the `loopback.900` loopback interface, the `rule1` Security policy rule, `ztp-nat` NAT policy rule, and the service route. These configurations are required to connect your ZTP firewall to Panorama and can lead to connectivity issues if modified.

STEP 8 | Select **Commit** and **Commit to Panorama**

STEP 9 | **Sync to ZTP Service** and verify that the Panorama Sync Status displays as `In Sync`.



STEP 10 | Add ZTP Firewalls to Panorama.

Configure the ZTP Installer Administrator Account

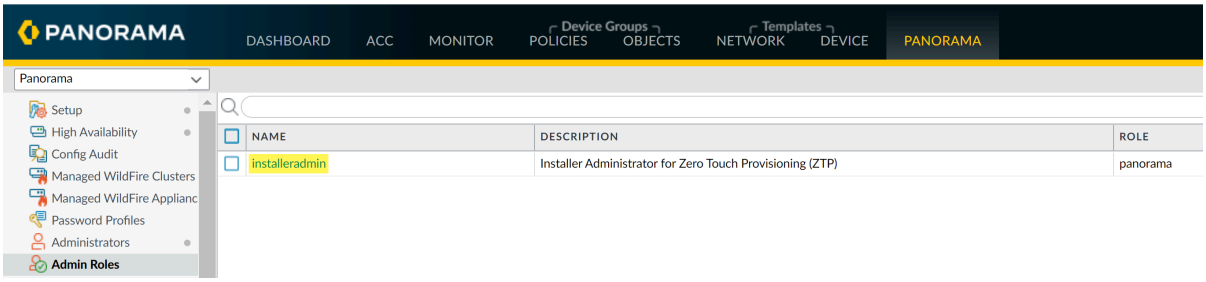
The ZTP installer admin user is an administrator account created for non-IT staff or installation contractor to on-board new ZTP firewalls. The installer admin uses an automatically created `installeradmin` admin role to limit visibility into the Panorama web interface and only allow the installer the ability to enter the ZTP firewall claim key and serial number on Panorama.

STEP 1 | Log in to the Panorama Web Interface.

 ©2025 Palo Alto Networks, Inc.

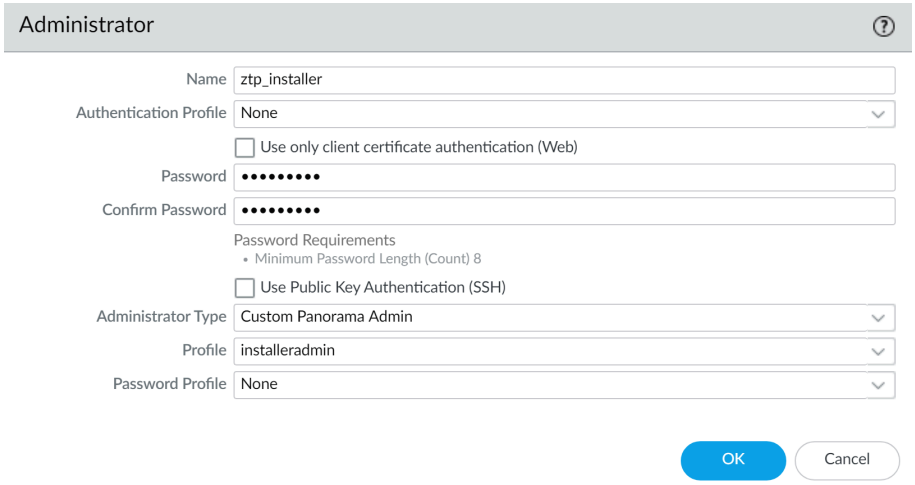
STEP 2 | Select **Panorama** > **Admin Roles** and verify that the `installeradmin` admin role is created.

The `installeradmin` is automatically created after you successfully install the ZTP plugin on Panorama.



STEP 3 | Configure the ZTP installer administrator user.

1. Select **Panorama** > **Administrators** and **Add** a new admin user.
2. Enter a descriptive **Name** for the ZTP installer admin user.
3. Enter a secure **Password** and **Confirm Password**.
4. For the **Administrator Type**, select **Custom Panorama Admin**.
5. For the **Profile**, select **installeradmin**
6. Click **OK** to save your configuration changes.



STEP 4 | Select **Commit** and **Commit to Panorama**.

Add ZTP Firewalls to Panorama

You can add a single ZTP firewall or import multiple ZTP firewalls to the Panorama™ management server.

- Add a ZTP Firewall to Panorama
- Import Multiple ZTP Firewalls to Panorama

 ©2025 Palo Alto Networks, Inc.

Add a ZTP Firewall to Panorama

Log in to the web interface of the Panorama™ management server as a Superuser, Panorama admin, or as the ZTP installer admin to add a ZTP firewall to Panorama. To add the ZTP firewall, you must enter the firewall serial number and claim key provided by Palo Alto Networks and then register the firewall with the ZTP service. Registering the firewall claims the firewall as an asset in your account in the Customer Support Portal and allows the ZTP service to associate the firewall with the Panorama.

Before you can successfully add a ZTP firewall to Panorama, you must ensure that a Dynamic Host Configuration Protocol (DHCP) server is deployed on the network. A DHCP server is required to successfully onboard a ZTP firewall to Panorama. The ZTP firewall is unable to connect to the Palo Alto Networks ZTP service to facilitate onboarding without a DHCP server.

If a ZTP firewall is initially assigned to a ZTP plugin-created template and device group, you cannot reassign it to a standard template or device group.

To associate a ZTP firewall with a regular (non-ZTP) device group and template, you must configure target templates and template stacks, including the network interfaces within the template. For details, see step 5.

By default, a ZTP firewall uses eth1/1 for management operations. To configure the ZTP firewall to use eth0 for all management operations, configure the Management Interface. For details, see step 6.

Migrating *a firewall added to Panorama management using ZTP from one Panorama to another is not supported.*

While adding ZTP firewalls to Panorama, do not perform any commits on the ZTP firewall before you verify that the firewall is successfully added to Panorama in Step 4. Performing a local commit on the ZTP firewall disables ZTP functionality and results in the failure to successfully add the firewall to Panorama.

STEP 1 | Log in to the Panorama Web Interface.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Add a ZTP firewall to Panorama.

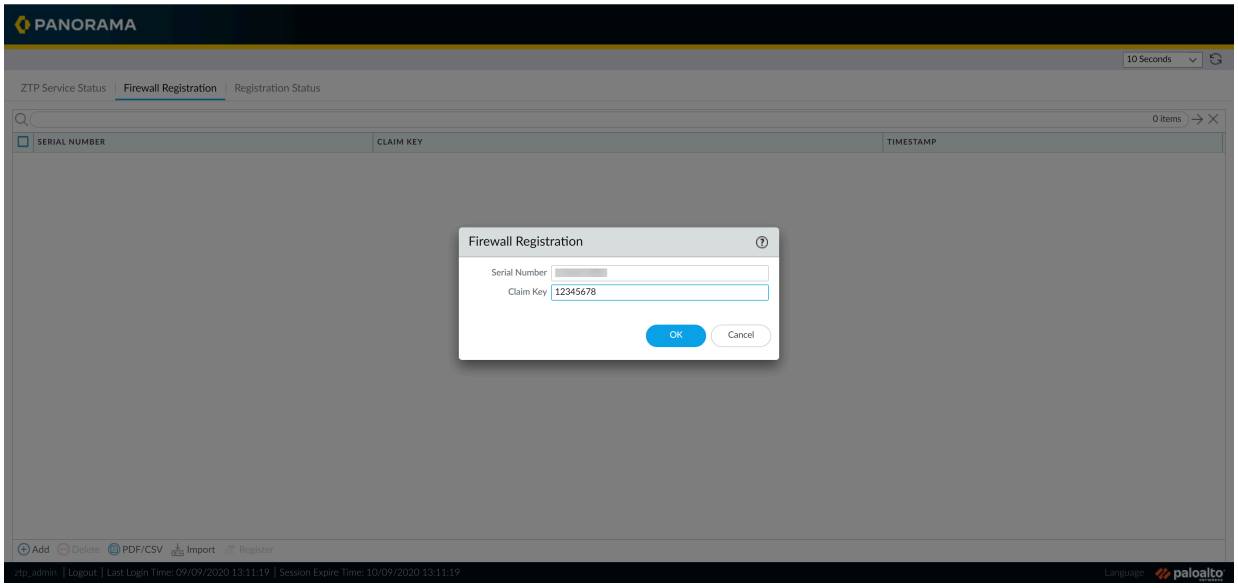
You must connect the Eth1/1 interface on ZTP firewalls to successfully register ZTP firewalls with the CSP and push the policy and network configurations.

1. Select **Firewall Registration** and **Add** a new ZTP firewall.
2. Enter the **Serial Number** of the ZTP firewall.
3. Enter the **Claim Key** for the ZTP firewall provided by Palo Alto Networks.

   The eight digit numeric claim key is printed on a physical label attached to the back of the ZTP firewall you received from Palo Alto Networks.



4. Click **OK** to save your configuration changes.



STEP 3 | Register the ZTP firewall.

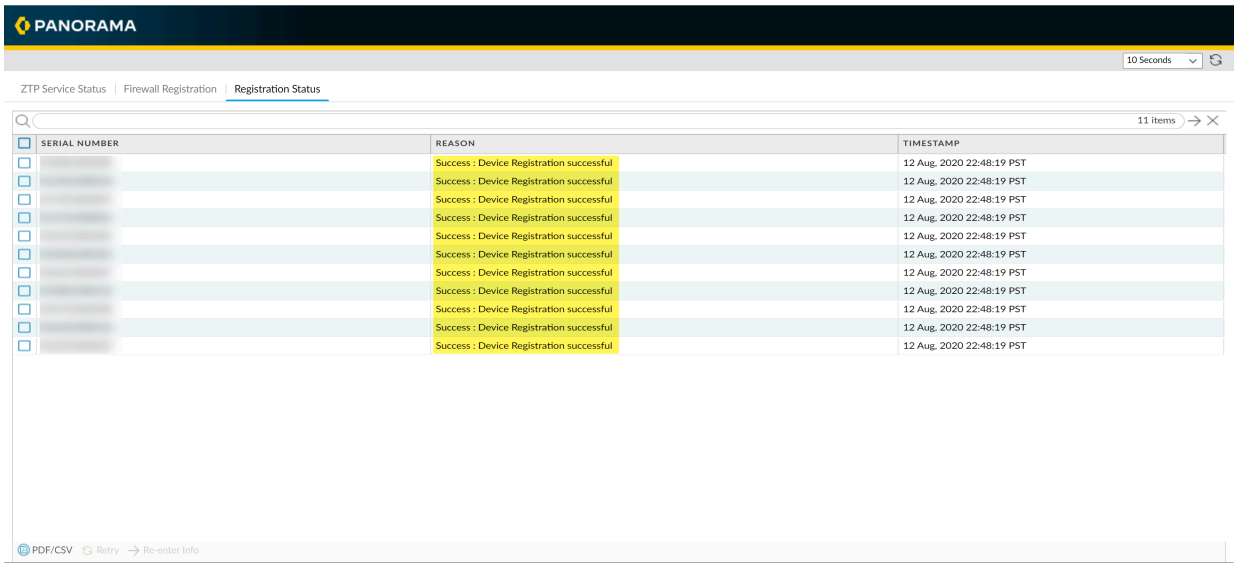
1. Select the newly added ZTP firewall and **Register** the firewall.
2. When prompted, click **Yes** to confirm registering the ZTP firewall.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Verify the firewall successfully registered with the CSP.

The firewall must successfully register with the CSP to successfully obtain device certificate.

1. Select **Registration Status** and verify that the ZTP firewall successfully registered with the CSP.



2. Log in to the Panorama Web Interface using admin credentials.
3. Select **Panorama** > **Managed Devices** > **Summary** and verify that the ZTP firewall is successfully added as a managed firewall.

*Ensure that the **To SW Version** column is configured to the correct PAN-OS version so that the firewall does not upgrade or downgrade unintentionally. ZTP functionality is supported only for PAN-OS 10.0.1 and later releases. Additionally, the PAN-OS version must be the same or an earlier version of the PAN-OS version running on Panorama.*

For more information, see Upgrade a ZTP Firewall.

*Using the **To SW Version** column when upgrading ZTP firewalls from PAN-OS 10.2 to PAN-OS 11.1 or later releases causes the upgrade flow to download and install an intermediary base version with an expired root certificate. This causes the firewall to lose its connection with Panorama. After onboarding the ZTP firewall, ensure that you manually upgrade to the target PAN-OS version that includes a valid root certificate*

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Configure Target Template(s) and Template Stacks to use non-ZTP device groups and templates.

1. Navigate to **Panorama** > **Templates**.
2. Configure the target templates and their respective template stacks that you plan to assign to the ZTP firewall. For more information, see Manage Templates and Template Stacks.
3. Select **Network** > **Interfaces** and then from the **Template** drop-down select the target Template.
4. Configure at least one interface, preferably eth1/1. If you do not require this interface, you can remove it after the initial successful.

STEP 6 | Configure Management Interface within the Template to use eth0 for management operations.

1. Select **Device** > **Setup** > **Interfaces**.
2. From the **Template** drop-down, select the target Template.
3. Select the **Management** interface and configure a static IP or choose **DHCP Client** to configure an IP address on the interface. For more information on how to configure connection settings, see Device > Setup > Interfaces.

   Ensure that the interface has connectivity to Panorama.
4. Perform a configuration push from Panorama to the ZTP Firewall.

STEP 7 | Add the ZTP firewall to the device group and template stack to be associated with the ZTP firewall.

You must add the ZTP firewall to a device group and template stack for your firewalls to display as `Connected` to push policy and network configurations.

You must keep the ZTP firewall in the ZTP device group and template stack that the ZTP template is associated with. This is required for the firewall to maintain connectivity with Panorama and prevent any unintended configuration reverts on the firewall.

1. Log in to the Panorama Web Interface using admin credentials.
2. Select **Panorama** > **Device Groups** and add the ZTP firewall to the device group created when you registered Panorama with the ZTP service.

   This is required for the ZTP firewall to successfully connect to Panorama.
3. Select **Panorama** > **Templates**, add the ZTP firewall to the template stack you created when you registered Panorama with the ZTP service.

   This is required for the ZTP firewall to successfully connect to Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Complete setting up the newly onboarded firewall.

1. Log in to the firewall web interface and activate the Support license.
2. Log in to the Panorama web interface and activate any additional licenses on your managed firewall.
3. Install the latest dynamic content updates on the managed firewall.
   1. Select **Panorama** > **Device Deployment** > **Dynamic Updates** and **Check Now** for the latest updates
   2. **Download** the latest dynamic content release version.
   3. **Install** and select the newly added firewalls.

      Click **OK** when prompted.
4. (Optional) Upgrade the managed firewall as needed.

Import Multiple ZTP Firewalls to Panorama

Log in to the web interface of the Panorama™ management server as a Superuser, Panorama admin, or as the ZTP installer admin to import multiple ZTP firewalls to Panorama. To import multiple ZTP firewalls, you must import a CSV file of the ZTP firewall serial number and corresponding claim key provided by Palo Alto Networks and then register the firewalls with the ZTP service. Registering the firewall claims the firewalls as assets in your account in the Customer Support Portal and allows the ZTP service to associate the firewalls with the Panorama.

Before you can successfully add a ZTP firewall to Panorama, you must ensure that a Dynamic Host Configuration Protocol (DHCP) server is deployed on the network. A DHCP server is required to successfully onboard a ZTP firewall to Panorama. The ZTP firewall is unable to connect to the Palo Alto Networks ZTP service to facilitate onboarding without a DHCP server.

Migrating *a firewall added to Panorama management using ZTP from one Panorama to another is not supported.*

Firewalls onboarded to Panorama management using ZTP do not support high availability (HA) configuration.

You must disable ZTP on your firewalls *to configure them in an HA configuration. After disabling ZTP,* add your firewalls as managed devices *and set up your firewalls in an* active/passive *or* active/active *HA configuration.*

While adding ZTP firewalls to Panorama, do not perform any commits on the ZTP firewall before you verify that the firewall is successfully added to Panorama in Step 5. Performing a local commit on the ZTP firewall disables ZTP functionality and results in the failure to successfully add the firewall to Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Gather the serial numbers and claim keys for your ZTP firewalls.

The eight digit numeric claim key is printed on a physical label attached to the back of the ZTP firewall you received from Palo Alto Networks.



STEP 2 | Create a CSV file containing the ZTP firewall serial numbers and claim keys. The first column must contain the serial numbers and the second column must contain the corresponding claim key for that firewall. Refer to the following example for reference.

	A	B
1	Serial Number	Claim Key
2	abcd1234	123456789
3	xyz7890	987654321

STEP 3 | Import the ZTP firewalls to Panorama.

You must connect the Eth1/1 interface on ZTP firewalls to successfully register ZTP firewalls with the CSP and push the policy and network configurations.

1. Log in to the Panorama Web Interface using the ZTP installer admin credentials.
2. Select **Panorama** > **Zero Touch Provisioning** > **Firewall Registration** and **Import** the ZTP firewalls.
3. **Browse** and select the CSV file containing the ZTP firewall information and click **OK**.

STEP 4 | Register the ZTP firewalls.

1. Select the newly added ZTP firewalls and **Register** the firewalls.
2. When prompted, click **Yes** to confirm registering the ZTP firewalls.

STEP 5 | Verify the firewall successfully registered with the ZTP service.

1. Select **Registration Status** and verify that the ZTP firewalls successfully registered with the ZTP service.
2. Log in to the Panorama Web Interface using admin credentials.
3. Select **Panorama** > **Managed Devices** > **Summary** and verify that the ZTP firewalls are successfully added as a managed firewall.

*Ensure that the **To SW Version** column is configured to the correct PAN-OS version so that the firewall does not upgrade or downgrade unintentionally. ZTP functionality is supported only for PAN-OS 10.0.1 and later releases. Additionally, the PAN-OS version must be the same or an earlier version of the PAN-OS version running on Panorama.*

For more information, see Upgrade a ZTP Firewall.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Add the ZTP firewall to the device group and template stack that contain the required ZTP configuration.

You must add the ZTP firewall to a device group and template stack for your firewalls to display as `Connected` to push policy and network configurations.

You must keep the ZTP firewall in the ZTP device group and template stack that the ZTP template is associated with. This is required for the firewall to maintain connectivity with Panorama and prevent any unintended configuration reverts on the firewall.

1. Log in to the Panorama Web Interface using admin credentials.
2. Select **Panorama** > **Device Groups** and add the ZTP firewall to the device group created when you registered Panorama with the ZTP service.

   This is required for the ZTP firewall to successfully connect to Panorama.
3. Select **Panorama** > **Templates**, add the ZTP firewall to the template stack you created when you registered Panorama with the ZTP service.

   This is required for the ZTP firewall to successfully connect to Panorama.

STEP 7 | Complete setting up the newly onboarded firewalls.

1. Log in to the firewall web interface and activate the Support license.

   You must activate the Support license locally for each managed firewall you added to Panorama management.
2. Log in to the Panorama web interface and activate any additional licenses on your managed firewalls.
3. Install the latest dynamic content updates on the managed firewalls.
   1. Select **Panorama** > **Device Deployment** > **Dynamic Updates** and **Check Now** for the latest updates
   2. **Download** the latest dynamic content release version.
   3. **Install** and select the newly added firewalls.

      Click **OK** when prompted.
4. (Optional) Upgrade the managed firewalls as needed.

Configure ZTP for Firewalls in HA

Where Can I Use This?	What Do I Need?
• NGFW (Managed by Panorama)	❑ Device management license ❑ Support license ❑ Claim key ❑ Auth code

 ©2025 Palo Alto Networks, Inc.

For high availability environments, ZTP supports configurations with both the paired Panorama appliances and firewalls, ensuring continuity of management capabilities. The system requires proper network connectivity between Panorama, the ZTP Service, and the firewalls being provisioned.

Before you can successfully add a ZTP firewall to Panorama, you must ensure that a Dynamic Host Configuration Protocol (DHCP) server is deployed on the network. A DHCP server is required to successfully onboard a ZTP firewall to Panorama. The ZTP firewall is unable to connect to the Palo Alto Networks ZTP service to facilitate onboarding without a DHCP server.

For ZTP firewalls in an HA configuration, you must use the ethernet0 interface for management operations. ZTP supports only active/passive HA configurations and not active/active HA configurations.

STEP 1 | Configure Panorama for the ZTP Service.

1. Log in to the Panorama web interface as a Superuser or Panorama administrator.
2. Install the Panorama Device Certificate.
3. Install the ZTP Plugin.

STEP 2 | Log in to the Palo Alto Networks Customer Support Portal (CSP).

STEP 3 | Associate your Panorama with the ZTP Service on the Palo Alto Networks CSP.

STEP 4 | Log in to the Panorama web interface again.

STEP 5 | Select **Panorama** > **Zero Touch Provisioning** > **Setup** and edit the **General ZTP** settings.

STEP 6 | Register Panorama with the ZTP service.

STEP 7 | Add the ZTP firewalls (active peer) to Panorama.

1. Select **Zero Touch Provisioning** > **Firewall Registration** and click **Add**.
2. Enter the Serial Number of the ZTP firewall.
3. Enter the Claim Key for the ZTP firewall provided by Palo Alto Networks. The eight-digit numeric claim key is printed on a physical label attached to the back of the ZTP firewall you received from Palo Alto Networks.
4. Click **OK** to save your configuration changes.
5. Select the newly added ZTP firewall and **Register** the firewall.
6. When prompted, click **Yes** to confirm registering the ZTP firewall.

STEP 8 | Add the ZTP firewalls (passive peer) to Panorama.

Repeat step 7.

STEP 9 | Verify the firewall successfully registered with the CSP.

Select **Registration Status** and verify that the ZTP firewall successfully registered with the CSP.

Registration is required to obtain the device certificate.

 ©2025 Palo Alto Networks, Inc.

STEP 10 | (Optional) Add a Device Group.

You can either add a new Device Group or use an existing one.

1. Select **Panorama** > **Device Groups** and click **Add**.
2. Enter a unique **Name** and a **Description** to identify the device group and click **OK**.
3. Add Security Policy Rules to the Device Group.

STEP 11 | (Optional) Configure the template and template stack for the active primary firewall. You can configure new templates and template stacks or use existing ones.

Add a template.

1. Select **Panorama** > **Templates**.
2. Click **Add** and enter a unique **Name** to identify the template and click **OK**.

Create a template stack and add the template generated in the previous step.

1. Select **Panorama** > **Templates** and click **Add Stack**.
2. Enter a descriptive Name for the template stack.
3. (PAN-OS 11.2 and later releases) Check (enable) Automatically push content when software device registers to Panorama.
4. In the **Templates**, add the template generated in the previous step and click **OK**.
5. Select **Device** > **Setup** > **Interfaces**.
6. From the **Template** drop-down, select the Template you created.
7. Select the **Management** interface and choose **IPv4** as the **DHCP Client**. For more information on how to configure connection settings, see Device > Setup > Interfaces.
8. Select **Network** > **Interfaces** > **Template** and then from **Template** the drop-down select target Template.
9. Configure the eth1/1 interface and set the interface type to **Layer3** or **HA**. For more information on configuring interfaces, see Network > Interfaces.
10. Select **Device** > **High Availability** and configure the HA settings for the active template.

STEP 12 | (Optional) Configure the template and template stack for the passive firewall.

Repeat step 11.

STEP 13 | Add the ZTP firewalls (HA Pair) to the device group and template stack.

1. Select **Panorama** > **Managed Devices** > **Summary**.
2. Select the ZTP firewalls and click **Reassociate**.
3. Select the respective device group and template stacks. Select the active template stack for the active firewall and the passive template for the passive firewall.
4. Check (enable) **Auto Push on 1st Connect** to automatically push the device group and template stack configurations when the ZTP firewall successfully connects to Panorama for the first time and click OK.
5. Select **Commit** > **Commit to Panorama** and commit your changes.

STEP 14 | Connect the eth1/1 interface and the management interface of the ZTP firewall to the management network.

 ©2025 Palo Alto Networks, Inc.

STEP 15 | Power on the ZTP firewall. Wait for the ZTP firewall to finish powering on.

STEP 16 | Monitor the ZTP process.

Use the CLI for ZTP Tasks

Use the following CLI commands to perform Zero Touch Provisioning (ZTP) tasks and view the ZTP service status.

If you want to ...	Use ...
Administer the firewall from the firewall CLI	
Display the connection status to the ZTP service.	`> show system ZTP status`
Display the connection status to the Panorama management server.	`> show panorama status`
Display the ZTP model number and firewall system information.	`> show system info`
Enable the ZTP state machine on the firewall. PA-5400, PA-400, PA-410, PA-1400, and PA-3400 only.	`> set system ztp enable` *Re-enabling the ZTP state machine initiates a soft factory reset that results in the deletion of the existing firewall configuration.*
Disable the ZTP state machine on the firewall. *Disabling the ZTP state machine initiates a soft factory reset that results in the deletion of the existing firewall configuration.*	`> request disable-ztp` PA-220-ZTP, PA-220R-ZTP, PA-800-ZTP, PA-850-ZTP, PA-3220-ZTP, PA-3250-ZTP, and PA-3260-ZTP only *You cannot re-enable the ZTP state machine on the firewall after it is disabled from the CLI.* *To re-enable, you must reset the firewall to factory default settings.*

 ©2025 Palo Alto Networks, Inc.

If you want to ...	Use ...
	`> set system ztp disable` PA-5400, PA-400, PA-410, PA-1400, and PA-3400 only.
Register, configure, and manage your ZTP firewalls from Panorama	
Create a device group or template containing the necessary configurations to connect managed firewalls with Panorama using the ZTP service on the Eth1/1 interface.	`> request plugins ztp create dgroup-template device-group <device group name>` `> request plugins ztp create dgroup-template template <template name>`
Add a ZTP firewall to the list of firewalls for future registration with the ZTP service.	`> request plugins ztp firewall-add <serial number> claim-key <claim key>`
Modify the serial number of a ZTP firewall that has already been added to the list of firewalls for future registration with the ZTP service.	`> request plugins ztp firewall-add-modify firewall <old serial number> claim-key <claim key> new-serial <new serial number>`
Delete a ZTP firewall from the list of firewalls for future registration with the ZTP service.	`> request plugins ztp firewall-delete firewall <serial number>`
Add a ZTP firewall to the list of firewalls for future re-registration with the ZTP service. Use this command when a ZTP firewall initially fails registration with the ZTP service and needs.	`> request plugins ztp firewall-re-enter-info firewall <serial number> claim-key <claim key>`
Register your Panorama™ management server with the ZTP service.	`> request plugins ztp panorama-registration`

 ©2025 Palo Alto Networks, Inc.

If you want to ...	Use ...
Register a ZTP firewall with the ZTP service.	`> request plugins ztp firewall-registration firewall <serial number> claim-key <claim key>`
Re-register ZTP firewalls with the ZTP service. Use this command to start the re-registration process for a ZTP firewall that failed initial registration with the ZTP service.	`> request plugins ztp firewall-register-retry firewall <serial number> claim-key <claim key>`
Import ZTP firewall serial number and claim key information. The specified file must be in CSV format.	`> request plugins ztp ztp-add-import import-path <file path>`
View ZTP firewall information and ZTP service status from Panorama	
Retrieve the list of ZTP firewalls registered to the Panorama from the ZTP service.	`> request plugins ztp ztp-service-info` The following details are displayed: • `first-firewall-connect-time`—Timestamp of when the ZTP firewall first connected to the ZTP service. • `last-firewall-connect-time`—Timestamp of when the ZTP firewall last connected to the ZTP service. • `registration-time`—Timestamp of when the ZTP firewall registered with the ZTP service. • `isZTPFirewall`—Whether the firewall is a ZTP firewall. • `created_by`—Administrative user that added the ZTP firewall. • `IP address`—IP address of the ZTP firewall.
View the list of ZTP firewalls in the list of firewalls to be registered with the ZTP service.	`> show plugins ztp device-add-list`

 ©2025 Palo Alto Networks, Inc.

If you want to ...	Use ...
View the registration status of your ZTP firewalls.	`> show plugins ztp device-reg-status`
View the ZTP service synchronization status for ZTP firewalls.	`> request plugins ztp ztp-sync-status`
Show the full management plane ZTP connectivity history. This is helpful for troubleshooting connectivity to the ZTP service.	`> tail follow yes mp-log ms.log`

Uninstall the ZTP Plugin

Follow the procedure to remove the ZTP configuration from your Panorama™ management server and uninstall the ZTP plugin. If your Panorama is in a high availability (HA) configuration, repeat these steps on both Panorama HA peers.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Delete the ZTP installer administrator account.

1. Select **Panorama** > **Administrators** and select the ZTP installer administrator account you previously configured.
2. **Delete** the ZTP installer administrator account.
3. Select **Panorama** > **Administrators** and select the `installeradmin` admin role.
4. **Delete** the `installeradmin` admin role.
5. Select **Commit** and **Commit to Panorama**.

STEP 3 | Uninstall the ZTP plugin

1. Select **Panorama** > **Plugins** and navigate to the ZTP plugin installed on Panorama.
2. In the Actions column, **Remove Config** to delete ZTP related configurations from Panorama
3. Click **OK** when prompted to confirm removing the ZTP configuration from Panorama.
4. Select **Commit** and **Commit to Panorama**.
5. **Uninstall** the ZTP plugin.
6. Click **OK** when prompted to uninstall the ZTP plugin from Panorama.

 ©2025 Palo Alto Networks, Inc.

Manage Device Groups

- Add a Device Group
- Create a Device Group Hierarchy
- Create Objects for Use in Shared or Device Group Policy
- Revert to Inherited Object Values
- Manage Unused Shared Objects
- Manage Precedence of Inherited Objects
- Move or Clone a Policy Rule or Object to a Different Device Group
- Push a Policy Rule to a Subset of Firewalls
- Device Group Push to a Multi-VSYS Firewall
- Manage the Rule Hierarchy

Add a Device Group

After adding firewalls (see Add a Firewall as a Managed Device), you can group them into Device Groups (up to 1,024), as follows. Be sure to assign both firewalls in an active-passive high availability (HA) configuration to the same device group so that Panorama will push the same policy rules and objects to those firewalls. PAN-OS doesn't synchronize pushed rules across HA peers. To manage rules and objects at different administrative levels in your organization, Create a Device Group Hierarchy.

STEP 1 | Select **Panorama** > **Device Groups**, and click **Add**.

STEP 2 | Enter a unique **Name** and a **Description** to identify the device group.

STEP 3 | In the Devices section, select check boxes to assign firewalls to the group. To search a long list of firewalls, use the Filters.

You can assign any firewall to only one device group. You can assign each virtual system on a firewall to a different device group.

STEP 4 | In the Reference Template section, **Add** any templates or template stacks with objects referenced by the device group configuration.

You must assign the appropriate template or template stack references to the device group in order to successfully associate the template or template stack to the device group. This allows you to reference objects configured in a template or template stack without adding an unrelated device to a template stack.

Skip this step if the device group configuration does not reference any objects configured in a template or template stack.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | (Optional) Select **Group HA Peers** for firewalls that are HA peers.

You can only group managed firewall HA peers if they are in the same device group.

The firewall name of the passive or active-secondary peer is in parentheses. Grouping HA peers is a visual change and no configuration change occurs.

STEP 6 | Select the **Parent Device Group** (default is **Shared**) that will be just above the device group you are creating in the device group hierarchy.

STEP 7 | If your policy rules will reference users and groups, assign a **Master** firewall.

This will be the only firewall in the device group from which Panorama gathers username and user group information.

STEP 8 | Click **OK** to save your changes.

STEP 9 | Select **Commit** > **Commit and Push** and then **Commit and Push** your changes to the Panorama configuration and to the device group you added.

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 and the Shared location. 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 firewalls to it.
2. Remove firewall or vsys assignments from existing device groups if those assignments don't fit your planned hierarchy.
   1. Select **Panorama** > **Device Groups** and select the device group.
   2. In the Devices section, clear the check boxes of firewalls and virtual systems you want to remove, and click **OK**.
3. If necessary, add more firewalls that you will assign to device groups: see Add a Firewall as a Managed Device.
4. If you are using multiple Panorama plugins to perform endpoint monitoring, a device group containing firewalls deployed in a particular hypervisor cannot be the child or parent of a device group containing firewalls deployed in a different hypervisor. See Device Group Hierarchy for more information.

STEP 2 | For each top-level device group, Add a Device Group.

1. In the **Panorama** > **Device Groups** page, click **Add** and enter a **Name** to identify the device group.
2. In the Devices section, select check boxes to assign firewalls and virtual systems to the device group.
3. Leave the **Parent Device Group** option at **Shared** (the default) and click **OK**.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | For each lower-level device group, Add a Device Group.

- For new device groups at each lower level, repeat the previous step, but set the **Parent Device Group** to a device group at the next level above.
- For each existing device group, in the **Device Groups** page, select the device group to edit it, select a **Parent Device Group**, and click **OK**.

If you move a device group to a different parent, all its descendant device groups move with it, along with all firewalls, policy rules, 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.

STEP 4 | Configure, move, and clone objects and policy rules as needed to account for inheritance in the device group hierarchy.

- Create Objects for Use in Shared or Device Group Policy, or edit existing objects.

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. However, you can optionally see Step Override inherited object values.

- Create or edit policies.
- Move or Clone a Policy Rule or Object to a Different Device Group.

STEP 5 | Override inherited object values.

Applicable only if object values in a particular device group must differ from the values inherited from an ancestor device group.

After overriding an object, you can override it again in descendant device groups. However, you can never override shared or predefined (default) objects.

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**).
2. Select the **Device Group** that will have the override instance.
3. Select the object and click **Override**.
4. Edit the values. You can't edit the **Name** or **Shared** settings.
5. Click **OK**. The Name column displays a yellow-overlapping-green icon for the object to indicate it is overridden.

If necessary, you can later Revert to Inherited Object Values.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Save and commit your changes.

Commit to Panorama and push to device groups after any change to the hierarchy.

You must also push changes to templates if a template references objects in a device group (such as interfaces referencing addresses), and a firewall assigned to the template is no longer assigned to that device group because of a hierarchy change.

Select **Commit** > **Commit and Push** and then **Commit and Push** your changes to the Panorama configuration and to the device groups you added or changed.

Create Objects for Use in Shared or Device Group Policy

You can use an object in any policy rule that is in the Shared location, or in the same device group as the object, or in descendants of that device group (for details, see Device Group Objects).

Shared device group objects can be viewed and referenced in a specific device group. Changing the name of a Shared device group object in one device group changes the name of the Shared object in all device groups. This includes any configuration the Shared object is referenced, such as in Policy rules. Changing the name of a Shared device group object may cause the configuration push to managed firewalls to fail.

For example, you create a Shared object named ObjectA and create a Security policy rule in the DG1 device group where ObjectA is referenced. This configuration is pushed to your managed firewalls. Later in the DG1 device group, you change the name of ObjectA to ObjectB and try to push the configuration to your managed firewalls. This push fails because your managed firewalls have the Shared object with the name ObjectA as part of their configuration, and are expecting that configuration object to have the same name.

See Use Dynamic Address Groups in Policy to verify the number of supported registered IP addresses on Panorama if you intended to leverage dynamic address groups in order to create policies that automatically adapt to changes in your network.

 ©2025 Palo Alto Networks, Inc.

- Create a shared object.

  In this example, we add a shared object for URL Filtering categories for which we want to trigger alerts.

  1. Select the **Objects** > **Security Profiles** > **URL Filtering** tab and click **Add**.

     The **Objects** tab appears only after you Add a Device Group (at least one).
  2. Enter a **Name** and a **Description**.
  3. Select **Shared**.
  4. The **Disable Override** option is cleared by default, which means you can override inherited instances of the object in all device groups. To disable overrides for the object, select the check box.
  5. In the **Categories** tab, select every Category for which you want notification.
  6. In the **Action** column, select **Alert**.
  7. Click **OK** to save your changes to the object.
  8. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

- Create a device group object.

  In this example, we add an address object for specific web servers on your network.

  1. Select **Objects** > **Addresses** and select the **Device Group** in which you will use the object.
  2. Click **Add** and enter a **Name** to identify the object.
  3. Be sure to leave the **Shared** option cleared.
  4. The **Disable Override** option is cleared by default, which means you can override inherited instances of the object in device groups that are descendants of the selected **Device Group**. To disable overrides for the object, select the **Disable Override** option.
  5. Select the **Type** of address object and the associated value. For example, select **IP Range** and enter the IP address range for the web servers.
  6. Click **OK** to save your changes to the object.
  7. Select **Commit** > **Commit and Push** and then **Commit and Push** your changes to the Panorama configuration and to the device group where you added the object.

     When you activate an antivirus license on a firewall, a list of predefined IP lists are automatically added to the firewall. As a result, this reduces the total number of individual address objects, dynamic groups, external IP lists, predefined IP block lists, and external predefined IP lists you can push from Panorama.

 ©2025 Palo Alto Networks, Inc.

- View shared objects and device group objects in Panorama.

  In the pages of the **Objects** tab, the Location column indicates whether an object is shared or is specific to a device group.

  1. In the **Objects** tab, select the object type (**Objects** > **Addresses**, in this example).
  2. Select the **Device Group** to which you added the object.

     *The **Objects** tab only displays objects that are in the selected **Device Group** or are inherited from an ancestor device group or the Shared location.*
  3. Verify that the device group object appears. Note that the device group name in the Location column matches the selection in the **Device Group** drop-down.

Revert to Inherited Object Values

After overriding the values that a device group object inherits from an ancestor device group, you can revert the object to its ancestor values at any time. In the **Objects** tab, overridden objects have a yellow-overlapping-green icon ( ) in the Name column.

If you want to push ancestor values to all overridden objects instead of reverting a specific object, see Manage Precedence of Inherited Objects.

For the steps to override values, see Step 5

For details on object inheritance and overrides, see Device Group Objects.

STEP 1 | In the **Objects** tab, select the object type (for example, **Objects** > **Addresses**) and select the **Device Group** that has an override instance of the object.

STEP 2 | Select the object, click **Revert**, and click **Yes**. The Name column displays a green icon for the object, indicating that it now inherits all values from an ancestor device group.

STEP 3 | Select **Commit** > **Commit and Push** and then **Commit and Push** your changes to the Panorama configuration and to the device group where you reverted the object.

Manage Unused Shared Objects

When you push configuration changes Device Groups, by default Panorama pushes all shared objects to firewalls whether or not any shared or device group policy rules reference the objects. However, you can configure Panorama to push only the shared objects that rules reference in the device groups. The **Share Unused Address and Service Objects with Devices** option enables you to limit the objects that Panorama pushes to the managed firewalls.

*When **Share Unused Address and Service Objects with Devices** is disabled, Panorama ignores the **Target** firewalls when you* Push a Policy Rule to a Subset of Firewalls*. This means that all objects referenced by any rules are pushed to all firewalls in the device group.*

To limit the number of objects pushed to a set of managed firewalls, add the policy rules to a child device group and reference shared objects as needed. See Create a Device Group Hierarchy *for more information on creating a child device group.*

 ©2025 Palo Alto Networks, Inc.

On lower-end models, such as the PA-220, consider pushing only the relevant shared objects to the managed firewalls. This is because the number of objects that can be stored on the lower-end models is considerably lower than that of the mid- to high-end models. Also, if you have many address and service objects that are unused, clearing **Share Unused Address and Service Objects with Devices** reduces the commit times significantly on the firewalls because the configuration pushed to each firewall is smaller. However, disabling this option might increase the commit time on Panorama because Panorama has to dynamically check whether policy rules reference a particular object.

STEP 1 | Select **Panorama** > **Setup** > **Management**, and edit the Panorama Settings.

STEP 2 | Clear the **Share Unused Address and Service Objects with Devices** option to push only the shared objects that rules reference, or select the option to re-enable pushing all shared objects.

Unchecking this option forces Panorama to check all of its policies for references to the objects and may increase commit times depending upon the configuration.

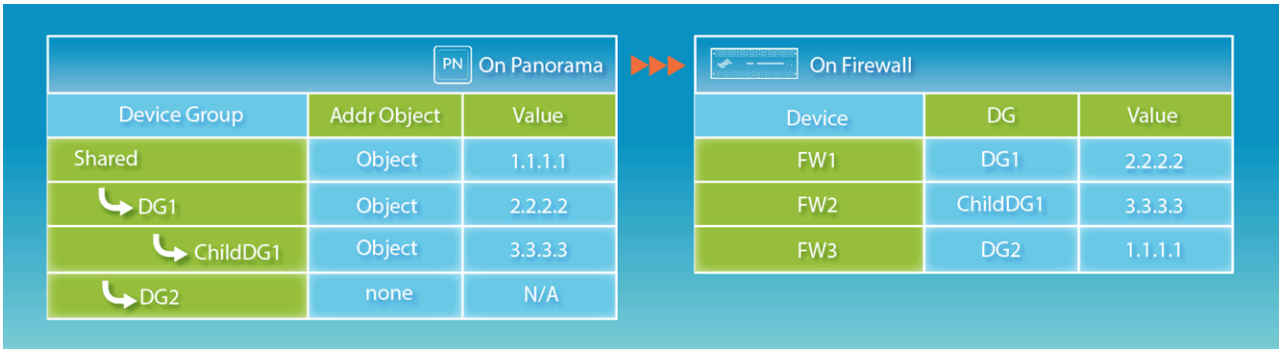
(Best Practices) If you plan to uncheck this option for all future commits, limit the number of shared configuration objects to help reduce commit times.

STEP 3 | Click **OK** to save your changes.

STEP 4 | Select **Commit** > **Commit to Panorama** and **Commit** your changes.

Manage Precedence of Inherited Objects

By default, when device groups at different levels in the Device Group Hierarchy have an object with the same name but different values (because of overrides, as an example), policy rules in a descendant device group use the object values in that descendant instead of using object values inherited from ancestor device groups. Optionally, you can reverse this order of precedence to push values from the highest ancestor containing the object to all descendant device groups. After you enable this option, the next time you push configuration changes to device groups, the values of inherited objects replace the values of any overridden objects in the descendant device groups. The figure below demonstrates the precedence of inherited objects in a device group:



If a firewall has locally defined objects with the same name as shared or device group objects that Panorama pushes, a commit failure occurs.

If you want to revert a specific overridden object to its ancestor values instead of pushing ancestor values to all overridden objects, see Revert to Inherited Object Values.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Select **Panorama** > **Setup** > **Management** and edit the Panorama Settings.

STEP 2 | If you want to reverse the default order of precedence, select **Objects defined in ancestors will take higher precedence**. The dialog then displays the **Find Overridden Objects** link, which provides the option to see how many overridden (shadowed) objects will have ancestor values after you commit this change. You can hover over the quantity message to display the object names.

If you want to revert to the default order of precedence, clear **Objects defined in ancestors will take higher precedence**.

__Find Overridden Objects__ only detects a Shared device group object that shares a name with another object in the device group.

STEP 3 | Click **OK** to save your changes.

STEP 4 | Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 5 | (Optional) If you selected **Objects defined in ancestors will take higher precedence**, Panorama does not push the ancestor objects until you push configuration changes to device groups: select **Commit** > **Push to Devices** and **Push** your changes.

Move or Clone a Policy Rule or Object to a Different 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**), you must move or clone the referenced objects and the referencing rule or object in the same operation. 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.

Audit comments do not move with the rule when you move the rule to a different device group.

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 Create Objects for Use in Shared or Device Group Policy).

When cloning multiple policy rules, the order by which you select the rules will determine the order they are copied to the device group. For example, if you have rules 1-4 and your selection order is 2-1-4-3, the device group where these rules will be cloned will display the rules in the same order you selected. However, you can reorganize the rules as you see fit once they have been successfully copied.

STEP 1 | Log in to Panorama and select the rulebase (for example, **Policy** > **Security** > **Pre Rules**) or object type (for example, **Objects** > **Addresses**).

STEP 2 | Select the **Device Group** and select one or more rules or objects.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Perform one of the following steps:

- (Rules only) **Move** > **Move to other device group**
- (Objects only) **Move**
- (Rules or objects) **Clone**

STEP 4 | In the **Destination** drop-down, select the new device group or **Shared**. The default is previously selected **Device Group**.

STEP 5 | (Rules only) Select the **Rule order**:

- **Move top** (default)—The rule will come before all other rules.
- **Move bottom**—The rule will come after all other rules.
- **Before rule**—In the adjacent drop-down, select the rule that comes after the Selected Rules.
- **After rule**—In the adjacent drop-down, select the rule that comes before the Selected Rules.

STEP 6 | The **Error out on first detected error in validation** check box is selected 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 is referenced in the rule you are moving. When you move or clone many items at once, selecting this check box can simplify troubleshooting. 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 | Select **Commit** > **Commit and Push**, **Edit Selections** in the Push Scope, select **Device Groups**, select the original and destination device groups, click **OK**, and then **Commit and Push** your changes to the Panorama configuration and to the device groups.

Ensure that you push all the changes when the original or the destination device groups do not have any devices. Selective push is not supported for device groups that do not have any devices.

Push a Policy Rule to a Subset of Firewalls

A policy *target* allows you to specify the firewalls in a device group to which to push policy rules. It allows you to exclude one or more firewalls or virtual systems, or to apply a rule only to specific firewalls or virtual systems in a device group.

As your rulebase evolves and you push new or modified rules to firewalls, changes and audit information get lost over time unless they are archived at the time the rule is created or modified. Use the audit comment archive to view the audit comment and configuration log history of a selected rule, as well to compare two policy rule versions to see how the rule changed. The audit comment history for a rule pushed from Panorama is viewable only from the Panorama management server. However, you can view the audit comments in the configurations logs forwarded to Panorama from managed firewalls. However, the audit comment archive is not viewable for rules created or modified locally on the firewall. To ensure that audit comments are

 ©2025 Palo Alto Networks, Inc.

captured at the time a rule is created or modified, Enforce Policy Rule, Description, Tag and Audit Comment.

The ability to target a rule enables you to keep policies centralized on Panorama. Targeted rules allow you to define the rules (as either shared or device group pre- or post-rules) on Panorama and improve visibility and efficiency when managing the rules (see Device Group Policies). The audit comment archive adds further visibility by allowing you to track how and why your policy rules change over time so you can audit the rule evolution over the course of the rule lifecycle.

STEP 1 | (Best Practice) Enforce audit comments for policy rules.

Although this step is optional, it is a best practice to enforce audit comments for policy rules to ensure that you capture the reason for creating or modifying the rule. This also helps maintain an accurate rule history for auditing purposes.

1. Select **Panorama** > **Setup** > **Management** and edit the Policy Rulebase Settings.
2. Enable the option to **Require audit comment on policies**.
3. Configure the Audit Comment Regular Expression to specify the audit comment format.

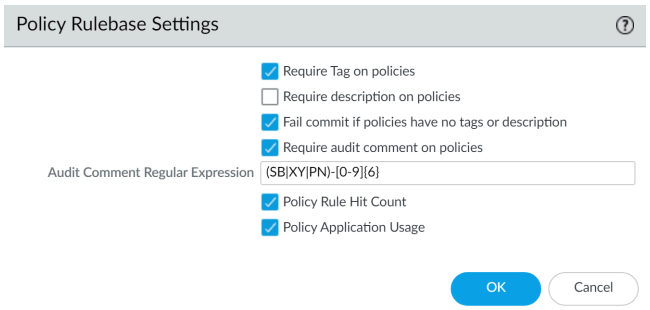
   When creating or modifying a rule, require audit comments to adhere to a specific format based on your business and auditing needs by specifying letter and number expressions. For example, you can use this setting to specify regular expressions to match your ticketing number formats:

   - **`[0-9]{<Number of digits>}`**—Requires the audit comment to contain a minimum number of digits ranging from 0 to 9. For example, **`[0-9]{6}`** requires a minimum of 6 digit numerical expression with numbers 0 to 9. Configure the minimum number of digits as needed.
   - <**`Letter Expression`**>—Requires the audit comment to contain a letter expression. For example, **`Reason for Change-`** requires that the administrator to begin the audit comment with this letter expression.
   - <**`Letter Expression`**>-**`[0-9]{<Number of digits>}`**—Requires the audit comment to contain a set character prefix with a minimum number of digits ranging from 0 to 9. For example, **`SB-[0-9]{6}`** requires the audit comment format to begin with **`SB-`**, followed by a minimum 6 digit numerical expression with numbers 0 to 9 such as **`SB-012345`**.
   - **`(`**<**`Letter Expression`**>**`)|(`**<**`Letter Expression`**>**`)|(`**<**`Letter Expression`**>**`)|-[0-9]{<Number of digits>}`**—Requires the audit comment to contain a prefix using one of the configured set of letter expressions with a minimum number of digits ranging from 0 to 9. For example, **`(SB|XY|PN)-[0-9]{6}`** requires the audit comment format begin with **`SB-`**, **`XY-`**, or **`PN-`** followed by a minimum 6

 ©2025 Palo Alto Networks, Inc.

digit numerical expression with numbers 0 to 9 such as **SB-012345**, **XY-654321**, or **PN-012543**.

4. Click **OK** to apply the new policy rulebase settings.



5. Select **Commit** and **Commit to Panorama**.

STEP 2 | Create a rule.

In this example, we define a pre-rule in the Security rulebase that permits users on the internal network to access the servers in the DMZ.

1. On the **Policies** tab and select the **Device Group** for which you want to define a rule.
2. Select the rulebase. For this example, select **Policies** > **Security** > **Pre-Rules** and **Add** a rule.
3. In the **General** tab, enter a descriptive rule **Name** and enter an **Audit Comment**.
4. In the **Source** tab, set the **Source Zone** to **Trust**.
5. In the **Destination** tab, set the **Destination Zone** to **DMZ**.
6. In the **Service/ URL Category** tab, set the **Service** to **application-default**.
7. In the **Actions** tab, set the **Action** to **Allow**.
8. Leave all the other options set to their default values.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Target the rule to include or exclude a subset of firewalls.

To apply the rule to a selected set of firewalls:

1. Select the **Target** tab in the Policy Rule dialog.
2. Select the firewalls to which you want to apply the rule.

   Ensure that you select either the target tags or target devices. If you select both target tags and target devices, the configuration push might fail.

   If you do not select firewalls to target, the rule is added to all of the (unchecked) firewalls in the device group.

   By default, although the check box for the virtual systems in the device group is disabled, all virtual systems will inherit the rule on commit unless you select one or more virtual systems to which you want the rule to apply.
3. (Optional) To exclude a subset of firewalls from inheriting the rule, **Install on all but specified devices** and select the firewalls you want to exclude.

   *If you **Install on all but specified devices** and do not select any firewalls, the rule is not added to any of the firewalls in the device group.*
4. Click **OK** to add the rule.

STEP 4 | Commit and push the configuration changes.

1. Select **Commit** > **Commit and Push** and **Edit Selections** in the Push Scope.
2. Select **Device Groups**, select the device group where you added the rule, and click **OK**.
3. (Optional) Disable **Merge with Device Candidate Config** if you manage local firewall configuration changes independently of configuration changes from Panorama.

   This setting is enabled by default and merges any pending local firewall configurations with the configuration push from Panorama. The local firewall configuration is merged and committed regardless of the admin pushing the changes from Panorama or the admin who made the local firewall configuration changes.
4. **Commit and Push** your changes to the Panorama configuration and to device groups.

STEP 5 | Troubleshoot Policy Rule Traffic Match to verify that the rules allow and deny traffic as the intended.

Device Group Push to a Multi-VSYS 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-vsys 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-vsys firewalls are included in a push from Panorama and one push fails, then the entire push fails to all

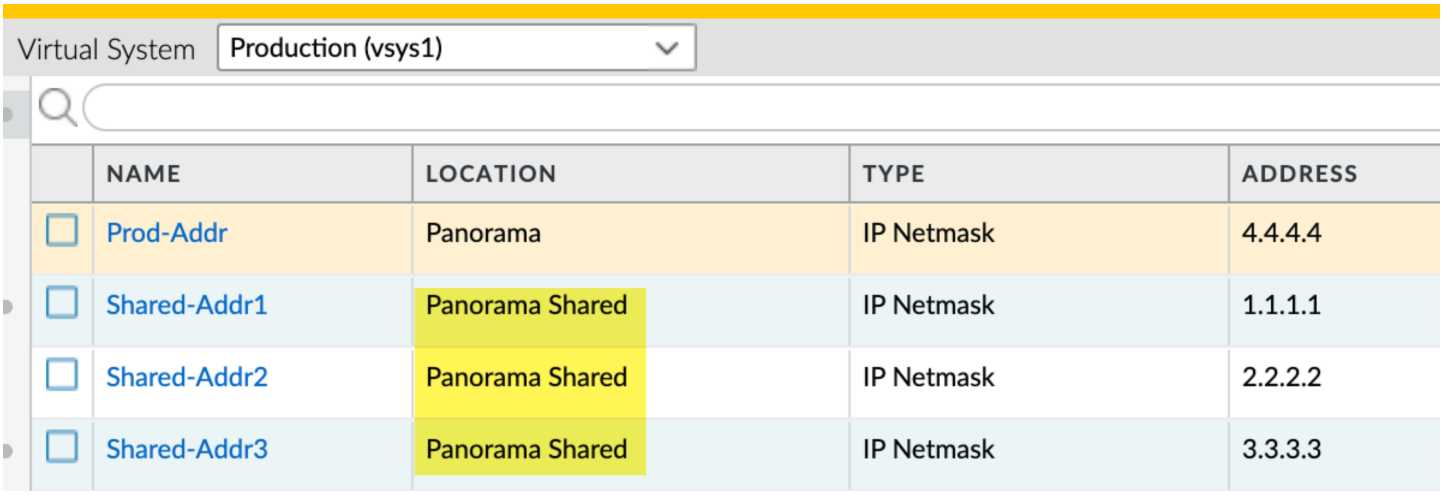
 ©2025 Palo Alto Networks, Inc.

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 are failures occur, an error description indicating the impacted vsys is displayed.

This functionality is supported for multi-vsys firewalls managed by Panorama running PAN-OS 10.2 and later releases by default. Palo Alto Networks recommends that all vsys of a multi-vsys managed firewall be managed by Panorama. After a successful upgrade to PAN-OS 10.2, a full commit and push from Panorama to managed firewalls is required to perform an administrator-level push which optimizes shared object pushes to multi-vsys firewalls as described below. If a full commit and push is not performed after upgrade, then all subsequent pushes to multi-vsys firewall fail due to duplicate objects and all shared configuration objects are saved to the `Panorama` location, rather than the optimized `Panorama Shared` location.

Shared Objects Pushed to a Multi-VSYS Firewall

To reduce the operational burden of scaling configurations for multi-vsys firewalls, Shared configuration objects pushed to a multi-vsys firewall are pushed to the Panorama Shared location on the managed multi-vsys firewall. The Panorama Shared location is available to all vsys of the firewall, meaning that Shared objects are not replicated to each vsys.



The following configurations cannot be added to the Shared Panorama location and are replicated to the Panorama location of each vsys of a multi-vsys firewall.

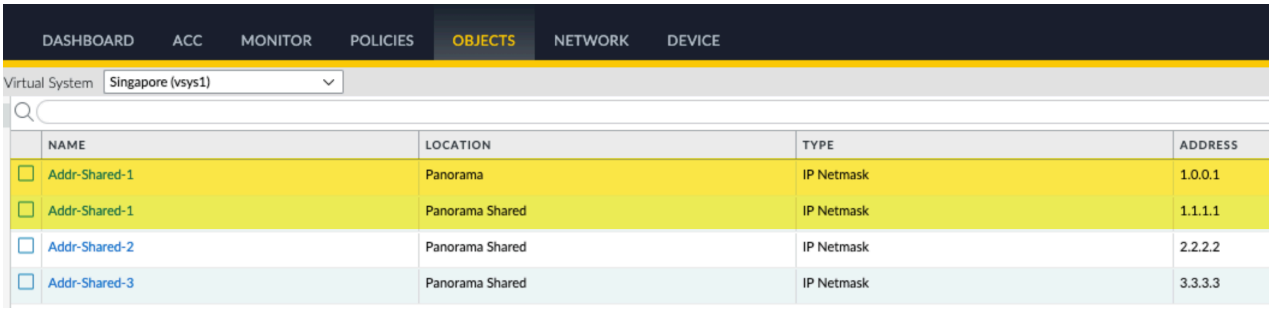
- *Pre and Post Rules*
- *External Dynamic Lists (EDL)*
- *Security Profile Groups*
- *HIP objects and profiles*
- *Custom URL objects*
- *Decryption Profiles*
- *SD-WAN Link Management Profiles*

If a Panorama Shared object is overridden in a device group, a new object with the same name but with the overridden value is created in the Panorama location of that device group and pushed to all vsys of a multi-vsys firewall. If the a configuration object with the same name is present in both

 ©2025 Palo Alto Networks, Inc.

the Panorama and the Panorama Shared locations, preference in the configuration given to the object in the Panorama location as because it is specific to that vsys on the firewall.

For example, the vsys below shows the `Addr-Shared-1` address object in both the Panorama Shared and Panorama locations. If the `Addr-Shared-1` object is used in a policy rule, the `1.0.0.1` IP address is used.



(Best Practice) When making configuration changes to shared objects on a multi-vsys firewall, ensure that you push the configuration changes to all device groups and vsys associated with the target firewall.

For example, in a multi-vsys firewall where vsys1 is in device group 1, vsys2 is in device group 2, and vsys3 is in device group 3, after making a configuration change to any shared object, ensure that you push the configuration change to all three device groups.

Manage the Rule Hierarchy

The order of policy rules is critical for the security of your network. Within any policy layer (shared, device group, or locally defined rules) and rulebase (for example, shared Security pre-rules), the firewall evaluates rules from top to bottom in the order they appear in the pages of the **Policies** tab. The firewall matches a packet against the first rule that meets the defined criteria and ignores subsequent rules. Therefore, to enforce the most specific match, move the more specific rules above more generic rules.

To understand the order in which the firewall evaluates rules by layer and by type (pre-rules, post-rules, and default rules) across the Device Group Hierarchy*, see* Device Group Policies.

STEP 1 | View the rule hierarchy for each rulebase.

1. Select the **Policies** tab and click **Preview Rules**.
2. Filter the preview by **Rulebase** (for example, **Security** or **QoS**).
3. Filter the preview to display the rules of a specific **Device Group** and the rules it inherits from the Shared location and ancestor device groups. You must select a device group that has firewalls assigned to it.
4. Filter the preview by **Device** to display its locally defined rules.
5. Click the green arrow icon to apply your filter selections to the preview (see Device Group Policies).
6. Close the Combined Rules Preview dialog when you finish previewing rules.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Delete or disable rules, if necessary.



*To determine which rules a firewall doesn't currently use, select that firewall in the **Context** drop-down on Panorama, select the rulebase (for example, **Policies** > **Security**), and select the **Highlight Unused Rules** check box. A dotted orange background indicates the rules that the firewall doesn't use.*

1. Select the rulebase (for example, **Policies** > **Security** > **Pre Rules**) that contains the rule you will delete or disable.
2. Select the **Device Group** that contains the rule.
3. Select the rule, and click **Delete** or **Disable** as desired. Disabled rules appear in italicized font.

STEP 3 | Reposition rules within a rulebase, if necessary.



*To reposition local rules on a firewall, access its web interface by selecting that firewall in the **Context** drop-down before performing this step.*

1. Select the rulebase (for example, **Policies** > **Security** > **Pre Rules**) that contains the rule you will move.
2. Select the **Device Group** that contains the rule.
3. Select the rule, select **Move**, and select:
   - **Move Top**—Moves the rule above all other rules in the device group (but not above rules inherited from Shared or ancestor device groups).
   - **Move Up**—Moves the rule above the one that precedes it (but not above rules inherited from Shared or ancestor device groups).
   - **Move Down**—Moves the rule below the one that follows it.
   - **Move Bottom**—Moves the rule below all other rules.
   - **Move to other device group**—See Move or Clone a Policy Rule or Object to a Different Device Group.

STEP 4 | If you modified the rules, commit and push the changes.

1. Select **Commit** > **Commit and Push** and **Edit Selections** in the Push Scope
2. Select **Device Groups**, select the device group that contains the rules you changed or deleted, and click **OK**.
3. **Commit and Push** your changes to the Panorama configuration and to device groups.

 ©2025 Palo Alto Networks, Inc.

Manage Templates and Template Stacks

Use templates and template stacks to define the common base configurations that enable firewalls to operate in your network. See Templates and Template Stacks for an overview of the issues you should consider when deciding which firewalls to add to which templates, ordering templates in a stack to manage layers of common and firewall group-specific settings, and overriding template settings with firewall-specific values.

To delete a template, you must first locally Disable/Remove Template Settings on the firewall. Only administrators with the superuser role can disable a template.

- Template Capabilities and Exceptions
- Add a Template
- Configure a Template Stack
- Configure a Template or Template Stack Variable
- Import and Overwrite Existing Template Stack Variables
- Override a Template Setting
- Disable/Remove Template Settings

Template Capabilities and Exceptions

You can use Templates and Template Stacks to define a wide array of settings but you can perform the following tasks only locally on each managed firewall:

- Configure a device block list.
- Clear logs.
- Enable operational modes such as normal mode, multi-vsys mode, or FIPS-CC mode.
- Configure the IP addresses of firewalls in an HA pair.
- Configure a master key and diagnostics.
- Compare configuration files (Config Audit).

  *To Manage Licenses and Updates (software or content) for firewalls, use the **Panorama** > **Device Management** tab options; do not use templates.*
- Renaming a vsys on a multi-vsys firewall.

Add a Template

You must add at least one template before Panorama™ displays the **Device** and **Network** tabs required to define the network setup and device configuration elements for firewalls. Panorama supports up to 1,024 templates. Every managed firewall must belong to a template stack. While templates contain managed device configurations, template stacks allow you to manage and push the template configurations to all managed firewalls assigned to the template stack.

 ©2025 Palo Alto Networks, Inc.

Combine templates in to a template stack to avoid duplicating many configurations among templates (see Templates and Template Stacks *and* Configure a Template Stack*).*

STEP 1 | Add a template.

1. Select **Panorama** > **Templates**.
2. Click **Add** and enter a unique **Name** to identify the template.
3. (Optional) Enter a **Description** for the template.
4. Click **OK** to save the template.
5. If the template has a virtual system (vsys) with configurations (for example, interfaces) that you want Panorama to push to firewalls that don't have virtual systems, select the template you created, select the vsys from the **Default VSYS** drop-down and click **OK**.
6. Select **Commit** > **Commit and Push** and then **Commit and Push** your changes to the Panorama configuration and to the template.

STEP 2 | Verify that the template is available.

After you add the first template, Panorama displays the **Device** and **Network** tabs. These tabs display a **Template** drop-down. Check that the drop-down displays the template you just added.

STEP 3 | Configure a Template Stack and add the template to the template stack.

 ©2025 Palo Alto Networks, Inc.

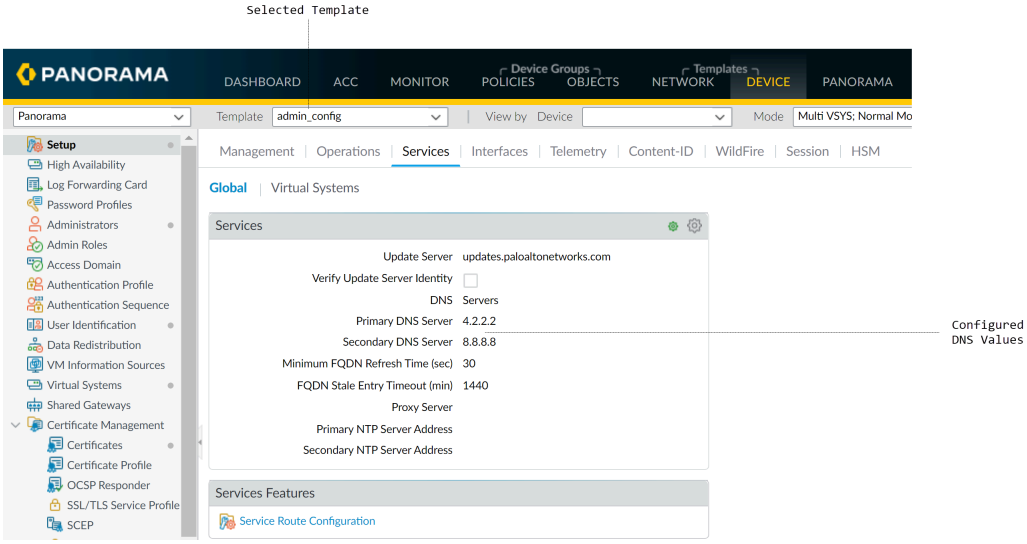
STEP 4 | Use the template to push a configuration change to firewalls.

Renaming a vsys is allowed only on the local firewall, not on Panorama the result is an entirely new vsys or the new vsys name gets mapped to the wrong vsys on the firewall.

For example, define a primary Domain Name System (DNS) server for the firewalls in the template.

You can also Configure a Template or Template Stack Variable to push device-specific values to managed devices.

1. In the **Device** tab, select the **Template** from the drop-down.
2. Select **Device** > **Setup** > **Services** > **Global**, and edit the Services section.
3. Enter an IP address for the **Primary DNS Server**.

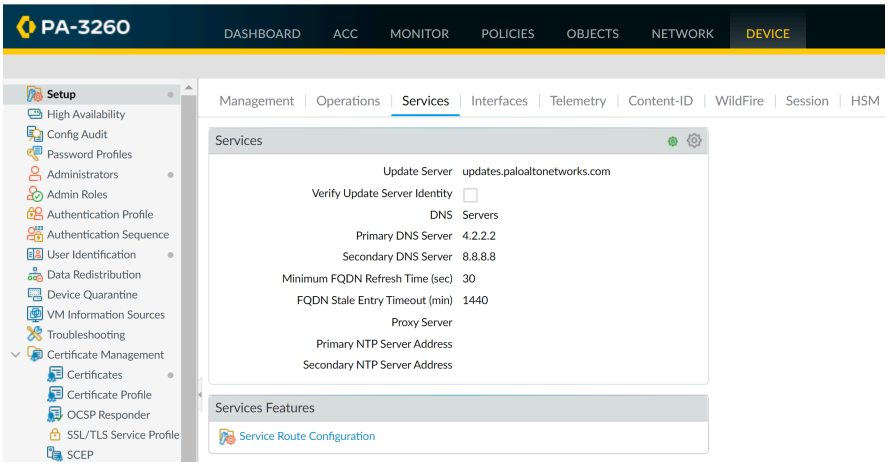


4. Select **Commit** > **Commit and Push** and then **Commit and Push** your changes to the Panorama configuration and to the template.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Verify that the firewall is configured with the template settings that you pushed from Panorama.

1. In the **Context** drop-down, select one of the firewalls to which you pushed the template setting.
2. Select **Device** > **Setup** > **Services** > **Global**. The IP address that you pushed from the template appears. The Services section header displays a template icon ( ) to indicate that settings in the section have values pushed from a template.



STEP 6 | Troubleshoot Connectivity to Network Resources to verify your firewalls can access your network resources.

Configure a Template Stack

A template stack is configurable and allows you to combine multiple templates to push full configurations to your managed firewalls. While templates are modular portions of your firewall configuration that you can reuse across different stacks, you can also configure the template stack to fill in the remaining configurations that you need to apply across all firewalls assigned to the stack. Panorama supports up to 1,024 template stacks and each stack can have up to 8 templates assigned to it. You can reference objects configured in a template stack from a template belonging to the temple stack. The template stack inherits configuration objects from the templates you add and is based on how you order templates in the template stack. You can also override template setting in the template stack to create a template stack configuration object. For details and planning, see Templates and Template Stacks.

Add a Template *to configure interfaces, VLANs, Virtual Wires, IPSec Tunnels, DNS Proxy and Virtual Systems. These objects must be configured and pushed from a template, and not a template stack. Once pushed from a template, you can override these objects, except for Virtual Systems, in the template stack.*

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Plan the templates and their order in the stack.

Add a Template you plan to assign to the template stack.

When planning the priority order of templates within the stack (for overlapping settings), you must check the order to prevent misconfiguration. For example, consider a stack in which the ethernet1/1 interface is of type Layer 3 in Template_A but of type Layer 2 with a VLAN in Template_B. If Template_A has a higher priority, Panorama will push ethernet1/1 as type Layer 3 but assigned to a VLAN.

Also note that a template configuration can't reference a configuration in another template even if both templates are in the same stack. For example, a zone configuration in Template_A can't reference a zone protection profile in Template_B.

STEP 2 | Create a template stack.

1. Select **Panorama** > **Templates** and **Add Stack**.

   *Panorama supports only **Add Stack** to create a new template stack. You cannot clone an existing template stack.*

2. Enter a unique **Name** to identify the stack.
3. (Optional) Add a **Description** for the template stack.
4. (Optional) Check (enable) **Automatically push content when software device registers to Panorama**.

   This setting is supported for VM-Series and CN-Series firewalls only. You must add the Panorama **Public IP** address to the Management Interface (**Panorama** > **Setup** > **Interfaces** > **Management**) to automatically push the Antivirus and Application and Threats content versions to VM-Series and CN-Series firewalls.

   VM-Series firewalls deployed on NSX and hardware firewalls are not supported.

   Enable this setting to automatically push the Antivirus and Applications and Threats content versions installed on Panorama to your VM-Series and CN-Series firewalls on first connection to Panorama. Panorama attempts to push the installed dynamic content versions one time and does not attempt any subsequent pushes of the installed Antivirus and Application and Threats content versions if the initial push fails for any reason.

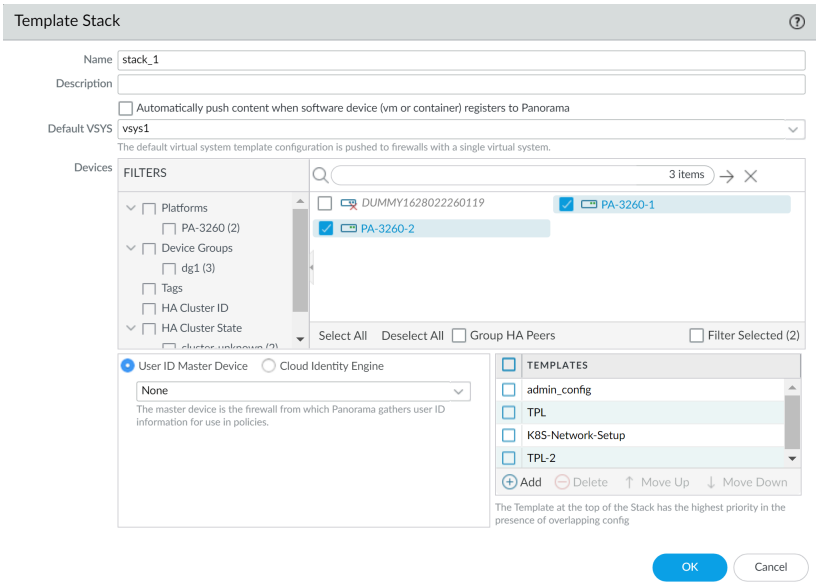
   For example, you add a VM-Series firewall to Panorama management and enable **Auto Push on 1st Connect** to automatically push the device group and template stack configuration to the VM-Series firewall on first connection. However, the template stack contains an invalid configuration and the push to the VM-Series firewall fails. In this scenario, the automatic content push to the VM-Series firewall also fails because the

 ©2025 Palo Alto Networks, Inc.

configuration push and dynamic content version push are included in the same push operation to the VM-Series firewall.

When leveraging auto-scale, enabling this setting allows you to maintain existing images for VM-Series and CN-Series firewalls leveraging dynamic content in their configurations, such as in policies and AppID. This helps eliminate the operational overhead required to update VM-Series and CN-Series firewall images when new dynamic content update versions are introduced.

5. For each of the templates the stack will combine (up to 8), **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 **Move Up** or **Move Down**.



6. In the Devices section, select firewalls to assign them to the stack. For firewalls with multiple virtual systems, you can't assign individual virtual systems, only an entire firewall. You can assign any firewall to only one template stack.

Whenever you add a new managed firewall to Panorama, you must assign it to the appropriate template stack; Panorama does not automatically assign new firewalls to a template or template stack. When you push configuration changes to a template, Panorama pushes the configuration to every firewall assigned to the template stack.

7. (Optional) Select **Group HA Peers** to display a single check box for firewalls that are in a high availability (HA) configuration. Icons indicate the HA state: green for active and yellow for passive. The firewall name of the secondary peer is in parentheses.

   For active/passive HA, add both peers to the same template so that both will receive the configurations. For active/active HA, whether you add both peers to the same template depends on whether each peer requires the same configurations. For a list of the configurations that PAN-OS synchronizes between HA peers, see High Availability Synchronization.

8. Click **OK** to save the template stack.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | (Optional) Configure a Template or Template Stack Variable.

STEP 4 | Edit the **Network** and **Device** settings, as necessary.

Renaming a vsys is allowed only on the local firewall. If you rename a vsys on Panorama, the result is an entirely new vsys or the new vsys name gets mapped to the wrong vsys on the firewall.

In an individual firewall context, you can override settings that Panorama pushes from a stack in the same way you override settings pushed from a template, see Override a Template or Template Stack Value.

1. Filter the tabs to display only the mode-specific settings you want to edit:

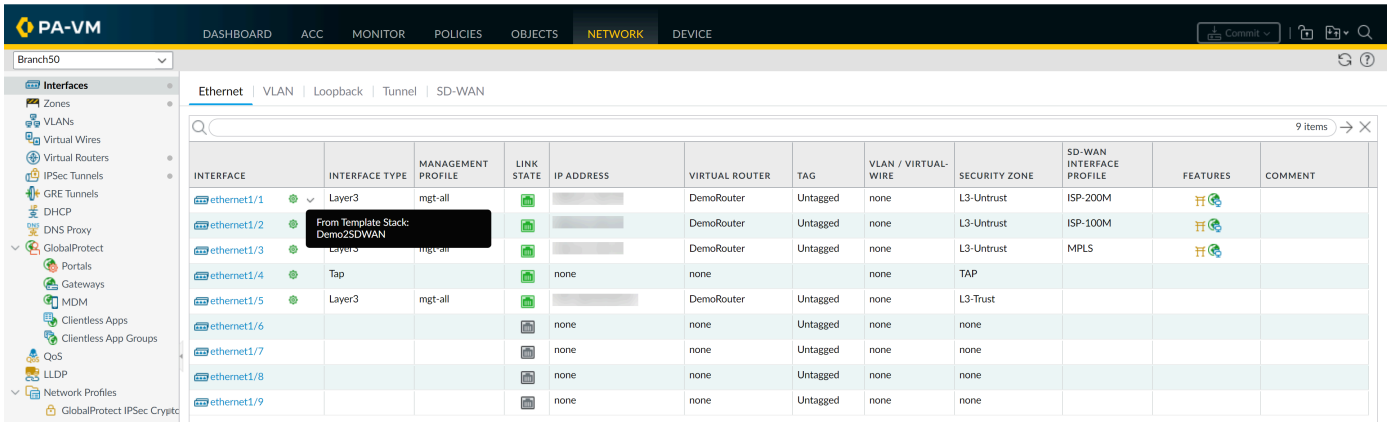
   *While Panorama pushes mode-specific settings only to firewalls that support those modes, this selective push doesn't adjust mode-specific values. For example, if a template has firewalls in Federal Information Processing Standards (FIPS) mode and an IKE Crypto profile that uses non-FIPS algorithms, the template push 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.*

   - In the **Mode** drop-down, select or clear the **Multi VSYS**, **Operational Mode**, and **VPN Mode** filter options.
   - Set all the **Mode** options to reflect the mode configuration of a particular firewall by selecting it in the **Device** drop-down.
2. Set up your interfaces and network connectivity. For example, Configure Zones and Interfaces to segment your network to manage and control traffic passing through your firewall.
3. Edit the settings as needed.
4. Select **Commit** > **Commit and Push**, **Edit Selections** in the Push Scope, select **Templates**, select the firewalls assigned to the template stack, and then **Commit and Push** your changes to the Panorama configuration and to the template stack.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Verify that the template stack works as expected.

1. Select a device assigned to the template stack from the **Context** drop-down.
2. Select a tab to which you pushed configuration changes using the template stack.
3. Values pushed from the template stack display a template icon ( ) to indicate that settings in the section have values pushed from a template stack. Hover your mouse over the stack to view which template stack from which the value was pushed.



STEP 6 | Troubleshoot Connectivity to Network Resources to verify your firewalls can access your network resources.

Configure a Template or Template Stack Variable

You can use template and template stack variables to replace IP addresses, IP ranges, FQDN, interfaces in IKE and VPN, and Group IDs in your configurations. If multiple templates in the template stack use different variables for the same configuration object, the variable value inherited by the template stack is based on the order of inheritance described in Templates and Template Stacks. Additionally, you can override a template value using a template stack variable to manage a configuration object from the template stack.

Variables allow you to reduce the total number of templates and template stacks you need to manage, while allowing you to keep any firewall- or appliance-specific values. For example, if you have a template stack with a base configuration, you can use variables to create values that do not apply to all firewalls in the template or template stack. This allows you to manage and push configurations from fewer templates and template stacks while accounting for any firewall- or appliance specific values that you would otherwise need before you can create a new template or template stack.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Create a template and template stack.

1. Add a Template
2. Configure a Template Stack.

STEP 3 | Select **Panorama** > **Templates** and **Manage** (Variables column) the template or template stack for which you want to create a variable.

 ©2025 Palo Alto Networks, Inc.

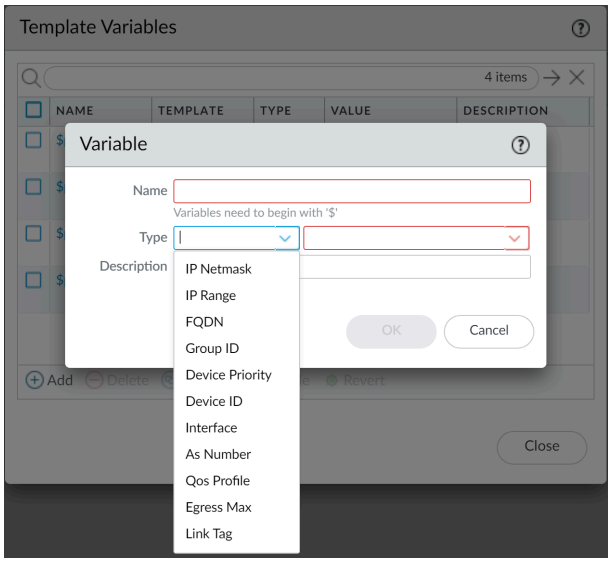
STEP 4 | **Add** the new variable.

A variable name must start with the dollar ($) symbol.

1. Name the new variable. In this example, the variables are named `$DNS-primary` and `$DNS-secondary`.
2. Select the variable **Type** and enter the corresponding value for the selected variable type.

   For this example, select **IP Netmask**.
3. (Optional) Enter a description for the variable.
4. Click **OK** and **Close**

Variables can also be created inline where variables are supported.



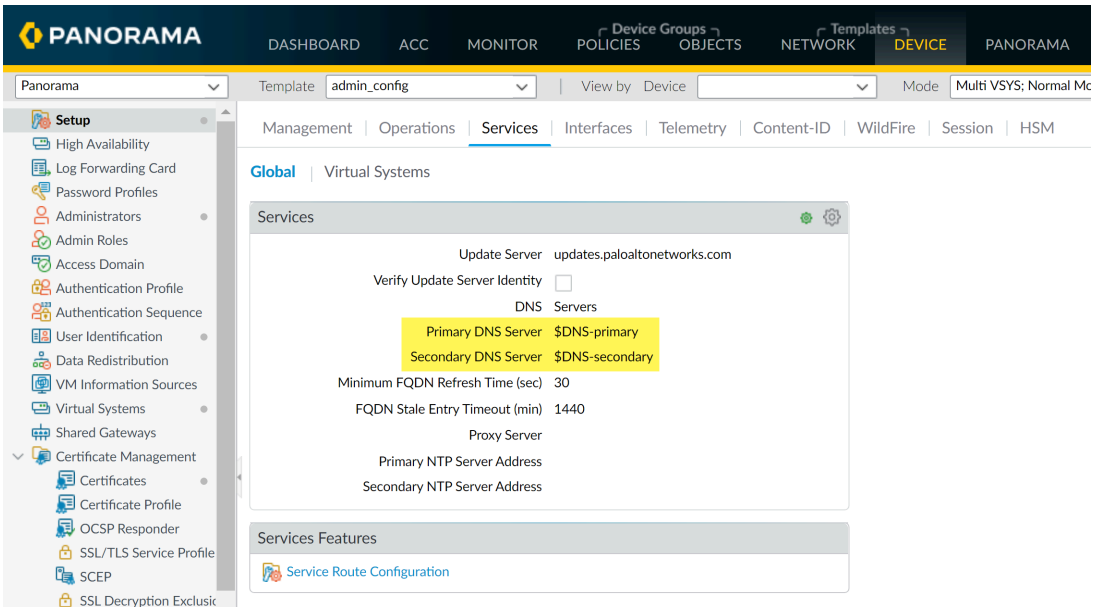
STEP 5 | From the **Template** drop-down, select the template or template stack to which the variable belongs.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Enter the variable in the appropriate location.

For this example, reference the previously defined DNS value.

1. Select **Device**.
2. From the **Template** drop-down, select the template or template stack to which the variable belongs.
3. Select **Setup** > **Services**.
4. Edit the Services.
5. Type **$DNS-primary** or select it from the drop-down for **Primary DNS Server**.
6. Type **$DNS-secondary** or select it from the drop-down for **Secondary DNS Server**.
7. Click **OK**.



STEP 7 | Click **Commit** and **Commit and Push** your changes to managed firewalls.

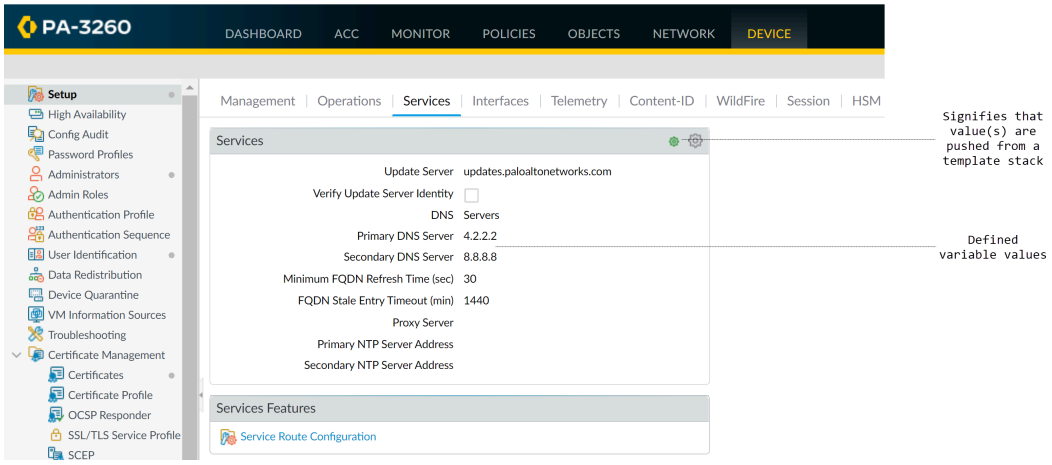
*When you push a device group configuration with references to template or template stack variables, you must **Edit Selections** and **Include Device and Network Templates**.*

STEP 8 | Verify that the values for all variables were pushed to the managed devices.

1. From the **Context** drop-down, select a firewall that belongs to the template stack for which the variable was created.
2. Select **Device** > **Setup** > **Services**.
3. Settings with values defined by a template or template stack are indicated by a template symbol ( ). Hover over the indicator to view to which template or template stack the

 ©2025 Palo Alto Networks, Inc.

variable definition belongs. When viewing from the firewall context, the variables display as the IP address you configured for the variable.



STEP 9 | Troubleshoot Connectivity to Network Resources to verify your firewalls can access your network resources.

Import and Overwrite Existing Template Stack Variables

Use template stack variables to replace IP addresses, IP ranges, FQDN, interfaces, or group ID in your firewall configurations. Variables allow you to reduce the total number of templates and template stacks you need to manage, while allowing you to preserve any firewall-specific values.

Importing template stack variables allows you to overwrite the values of multiple existing variables, and you cannot create new template stack variables when importing. For more information how on how to create new template or template stack variable, see Configure a Template or Template Stack Variable.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Export the existing template stack variables.

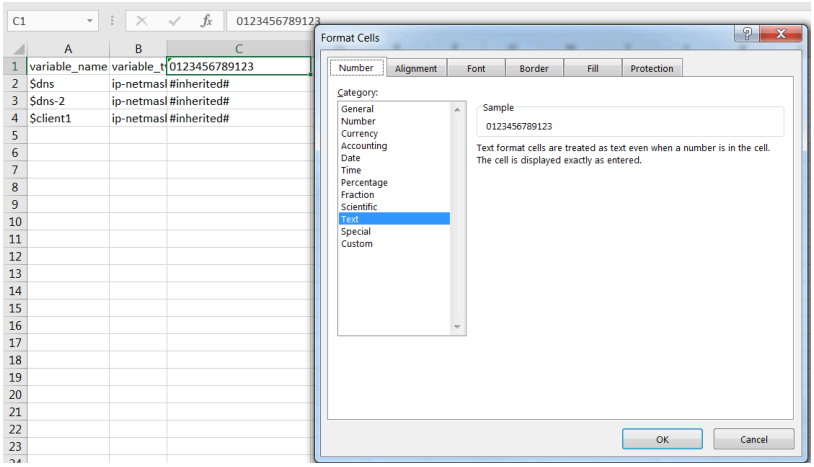
1. Select **Panorama** > **Templates** and select a template or template stack.
2. Select **Variable CSV** > **Export**. The configured template stack variables are downloaded locally as a CSV file.
3. Open the exported CSV.

 ©2025 Palo Alto Networks, Inc.

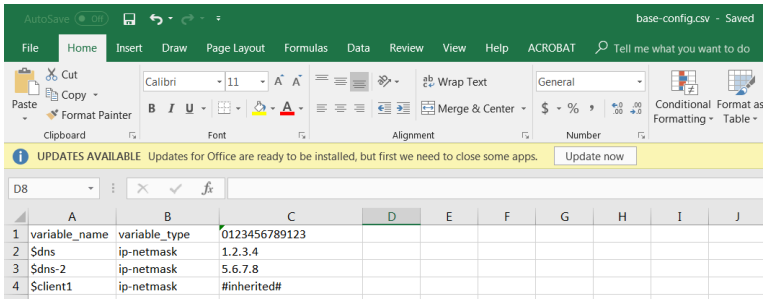
STEP 3 | Edit the CSV file containing the template stack variables to import to Panorama in the following format:

Values that display as `#inherited#` are values that are defined in the template stack.

1. Correct the number of the cells containing the firewall serial number. Repeat this step for all firewalls in the CSV file.
   1. Right-click the cell containing the firewall serial number and select **Format Cells**.
   2. Select **Number** > **Text** and click **OK**.
   3. Add a **0** at the beginning of the serial number.



2. Enter a new value for the desired template variable.
3. Select **File** > **Save As** and save the file in **CSV UTF-8** format.



	A	B	C
1	variable_name	variable_type	0123456789123
2	$dns	ip-netmask	1.2.3.4
3	$dns-2	ip-netmask	5.6.7.8
4	$client1	ip-netmask	#inherited#

STEP 4 | Import the CSV file to the template stack.

1. Log in to the Panorama Web Interface.
2. Select **Panorama** > **Templates** and select the template stack for which you exported the variables in Step 2.
3. Select **Variable CSV** > **Import** and **Browse** for the CSV file edited in Step 3.
4. Click **OK** to import the template stack variables.

STEP 5 | Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 6 | Enter the variables in the appropriate locations.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Click **Commit** and **Commit and Push** your changes to managed firewalls.

*When you push a device group configuration with references to template or template stack variables, you must **Edit Selections** and **Include Device and Network Templates**.*

Override a Template or Template Stack Value

While Templates and Template Stacks enable you to apply a base configuration to multiple firewalls, you might want to configure firewall-specific settings that don't apply to all the firewalls in a template or template stack. Conversely, you may want to override the template settings to create a template stack configuration that you can apply as a base configuration to all your managed firewalls. Overrides allow for exceptions or modifications to meet your configuration needs. For example, if you use a template to create a base configuration but a few firewalls in a test lab environment need different settings for the Domain Name System (DNS) server IP address or the Network Time Protocol (NTP) server, you can override the template and template stack settings.

If you want to disable or remove all the template or stack settings on a firewall instead of overriding a single value, see Disable/Remove Template Settings.

You can override a template or template stack value in one of the following ways:

- Override a Template Value on the Firewall or Override a Template or Template Stack Value Using Variables—There are two ways to override values pushed from a template or template stack. The first is to define a value locally on the firewall to override a value pushed from a template or template stack. The second is to define firewall-specific variables to override values pushed from a template or template stack.
- Override a Template Value Using a Template Stack—Define values or variables on the template stack to override values pushed from a template.

Override a Template Value on the Firewall

Override a setting on the local firewall that was pushed from a template or template stack to create firewall-specific configurations. This allows you to manage the base template or template stack configuration from Panorama™, while maintaining any firewall-specific configurations that do not apply to other firewalls.

STEP 1 | Access the firewall web interface.

Directly access the firewall by entering its IP address in the URL field of your browser or use the **Context** drop-down in Panorama to switch to the firewall context.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Override a value pushed from a template or template Stack.

In this example, you override the DNS server IP address that you assigned using a template in Add a Template

1. Select **Device** > **Setup** > **Services** and edit the Services section.
2. Click the template icon ( ) for the **Primary DNS Server** to enable overrides for that field.
3. Enter the new IP address for the **Primary DNS Server**. A template override symbol ( ) indicates that the template value was overridden.
4. Click **OK** and **Commit** your changes.

Override a Template Value Using a Template Stack

You can use template stack values to override configurations pushed to the managed firewall from a template to create a template stack configuration that you can use to manage the base configuration of your managed firewalls from Panorama™. This enables you to leverage the management capabilities of Panorama to push configuration changes to multiple devices from a single location. In this example, you will use a template stack to override the Primary DNS server IP address variable called $DNS that was pushed from a template.

Panorama supports using a template stack to override interfaces configured in a template except for Layer2 sub-interfaces of an aggregated interface.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | From the **Template** drop-down, select the template stack that will override the template configuration.

STEP 3 | Override the pushed template configuration.

1. Select **Device** > **Setup** > **Services** and edit the Services section.
2. Configure the **Primary DNS** with the IP address to override the pushed template configuration and click **OK**.

STEP 4 | **Commit and Push** the configuration change.

Override a Template Value Using a Template Stack Variable

You can use template stack values and variables to override configurations pushed to the managed firewall from a template to create a template stack configuration that you can use to manage the base configuration of your managed firewalls from Panorama™. This enables you to leverage the management capabilities of Panorama to push configuration changes to multiple firewalls from a single location. In this example, you will create a template stack variable by overriding the Primary DNS server IP address variable called $DNS that was pushed from a template.

Panorama supports using a template stack to override interfaces configured in a template except for Layer2 sub-interfaces of an aggregated interface.

STEP 1 | Log in to the Panorama Web Interface.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Override the template variable.

1. Select **Panorama** > **Templates**.
2. **Manage** (Variables column) the template stack containing the template variable you need to override.
3. Locate and select the **$DNS** variable.
4. Select **Override**.
5. Enter the new variable value and click **OK**.

STEP 3 | **Commit and Push** your changes.

Override a Template or Template Stack Value Using Variables

You can use firewall-specific variables to override variables pushed to the managed firewall from a template or template stack to create firewall-specific configurations. This allows you to manage the base template or template stack configuration while maintaining any firewall-specific configurations that do not apply to other firewalls—all from Panorama™. This allows you to leverage the management capabilities of Panorama while accounting for any specific configurations required for individual firewalls. In this example, the Primary DNS server IP address variable called $DNS that has been pushed from a template will be overridden to create a firewall-specific variable.

*You can override template or template stack variables that have not been overridden. If a template or template stack variable is already overridden, **Revert** the override to create a firewall-specific variable.*

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Override the template or template stack variable.

1. Select **Panorama** > **Managed Devices** > **Summary**.
2. **Edit** (Variables column) the firewall containing the variable you need to override.
3. Locate and select the **$DNS** variable.
4. Select **Override**.
5. Enter the new firewall-specific IP address and click **OK**.

STEP 3 | **Commit and Push** your changes.

Disable/Remove Template Settings

If you want to stop using a template or template stack for managing the configuration on a managed firewall, you can disable the template or stack. When disabling, you can copy the template/stack values to the local configuration of the firewall or delete the values.

If you want to override a single setting instead of disabling or removing every template or stack setting, see Override a Template Setting.

See Templates and Template Stacks *for details on how to use these for managing firewalls.*

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Access the web interface of the managed firewall as an administrator with the Superuser role. You can directly access the firewall by entering its IP address in the browser URL field or, in Panorama, select the firewall in the **Context** drop-down.

STEP 2 | Select **Device** > **Setup** > **Management** and edit the Panorama Settings.

STEP 3 | Click **Disable Device and Network Template**.

STEP 4 | (Optional) Select **Import Device and Network Template before disabling**, to save the configuration settings locally on the firewall. If you do not select this option, PAN-OS will delete all Panorama-pushed settings from the firewall.

STEP 5 | Click **OK** twice and then **Commit** the changes.

 ©2025 Palo Alto Networks, Inc.

Manage the Master Key from Panorama

Panorama, firewalls, Log Collectors, and WF-500 appliances use a master key to encrypt sensitive elements in the configuration and they have a default master key they use to encrypt passwords and configuration elements. As part of a standard security practice, you should replace the default master key and change the key on each individual firewall, Log Collector, WildFire appliance, and Panorama before it expires.

To strengthen your security posture, configuring a unique master key for Panorama and for each managed firewall. By configuring unique master keys, you can ensure that a compromised master key does not compromise the configuration encryption of 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.

The default encryption algorithm that the master key uses to encrypt data is AES-256-CBC—the same algorithm that the master key used prior to PAN-OS 10.0. AES-256-CBC is the default encryption level because when you manage firewalls with Panorama, the managed firewalls may be on different PAN-OS releases, and firewalls on PAN-OS releases earlier than PAN-OS 10.0 do not support AES-256-GCM. This is why Panorama must use the lowest level of encryption that its managed devices can use. For example, if some managed devices run PAN-OS 10.0 and some run earlier versions, Panorama must use AES-256-CBC. However, if all managed devices run PAN-OS 10.0 or later, then Panorama and all of its managed devices can use AES-256-GCM.

Palo Alto Networks recommends using AES 256-GCM level 2 for master key encryption.

When a master key expires, you must enter the current master key in order to configure a new master key.

Be sure to keep track of the master key you deploy to your managed firewalls, Log Collectors, and WildFire appliances because master keys cannot be recovered. you must reset to factory default if you cannot provide the current master key when it expires.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | (Best Practice) Select **Commit** and **Commit and Push** any pending configuration changes.

Panorama must re-encrypt data using the new master key. To ensure all configuration elements are encrypted with the new master key, you should commit all pending changes before deploying the new master key.

 ©2025 Palo Alto Networks, Inc.

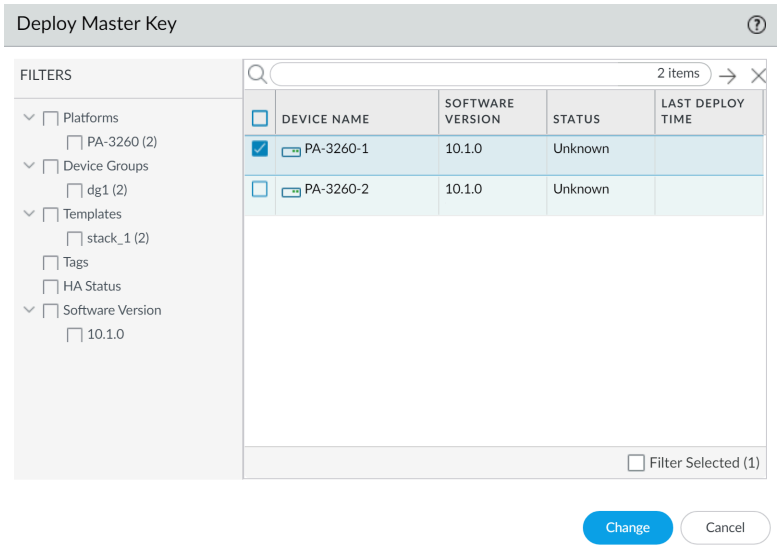
STEP 3 | Configure a unique master key for a managed firewall.

1. (HA only) Disable Config Sync for managed firewalls.

   This step is required before deploying a new master key to a firewall HA pair

   1. Log in to the Panorama Web Interface.
   2. Select **Device** > **High Availability** > **General** and select the **Template** containing the managed firewall HA configuration.
   3. Edit the HA Pair Settings **Setup**.
   4. Disable (clear) **Enable Config Sync** and click **OK**.
   5. **Commit** and **Commit and Push** your configuration changes.
2. Select **Panorama** > **Managed Devices** > **Summary** and **Deploy Master Key**.
3. 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.

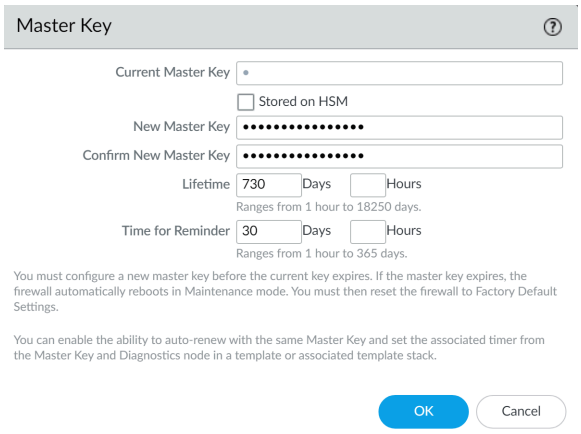


4. Configure the master key:

   1. If renewing a master key, enter the **Current Master Key**. If you are replacing the default master key with a new master key, do not specify a **Current Master Key**.
   2. (Optional) Enable (check) **Stored on HSM** if the master key is encrypted on a Hardware Security Module (HSM).
   3. Specify the **New Master Key** and **Confirm Master Key**.
   4. Configure the master key **Lifetime** and **Time for Reminder**.
   5. Click **OK**.

      *The new master key is automatically pushed to your managed firewalls after you click **OK**. Proceed only if you are certain you are ready to change the master key for your managed firewalls.*

 ©2025 Palo Alto Networks, Inc.



5. 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.

6. (Optional) Configure the master key to automatically renew 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.

   1. Select **Device** > **Master Key and Diagnostic** and select the **Template** containing the target managed firewalls.
   2. Edit the **Master Key** settings and configure the **Auto Renew With Same Master Key** setting.
   3. Click **OK**.
   4. **Commit** and **Commit and Push** your configuration changes.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the master key on Panorama.

1. (HA only) Disable the HA configuration for Panorama.

   This step is required to successfully change the master for both Panorama HA peers. You are unable to commit configuration changes on the secondary HA peer when Panorama is in an HA configuration.

   1. Log in to the Panorama Web Interface.
   2. Select **Panorama** > **High Availability** > **General** and edit the HA Setup.
   3. Disable (uncheck) **Enable HA** and click **OK**.
   4. **Commit** and **Commit to Panorama**.
2. Select **Panorama** > **Master Key and Diagnostics** and configure the master key.

   1. If renewing a master key, enter the **Current Master Key**. If you are replacing the default master key with a new master key, o not specify a **Current Master Key**.
   2. Configure the **New Master Key** and **Confirm Master Key**.
   3. Configure the master key **Lifetime** and **Time for Reminder**.
   4. Click **OK**.

   > *The new master key is automatically committed to Panorama after you click **OK**. Proceed only if you are certain you are ready to change the master key on Panorama.*
3. (Optional) Configure the Panorama master key to automatically renew.

   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.

   1. Select **Panorama** > **Master Key and Diagnostic** and edit the **Master Key** setting.
   2. Configure the **Auto Renew With Same Master Key** setting.
   3. Click **OK**.
4. Select **Commit** > **Commit to Panorama** and **Commit** your changes.
5. (HA only) Repeat this step to configure an identical master key on the secondary HA peer.

   You must manually configure an identical master key on the primary and secondary HA peers when Panorama is in an HA configuration. The master key is not synchronized between the primary and secondary HA peers.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Deploy the master key to Log Collectors.

The master key configured for your Log Collectors must be identical to the master key configured for Panorama.

1. Select **Panorama** > **Managed Collectors** and **Deploy Master Key**.
2. Select all devices and **Change** the master key.
3. Configure the master key:
   1. If renewing a master key, enter the **Current Master Key**. If you are replacing the default master key with a new master key, do not specify a **Current Master Key**.
   2. Specify the **New Master Key** and **Confirm Master Key**.
   3. Configure the master key **Lifetime** and **Time for Reminder**.
   4. Click **OK**.

   *The new master key is automatically pushed to your Log Collectors after you click **OK**. Proceed only if you are certain you are ready to change the master key for your Log Collectors.*
4. Verify that the master key was deployed successfully to all selected devices.

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

STEP 6 | Deploy the master key to managed WildFire appliances.

The master key configured your WildFire appliances must be identical to the master key configured for Panorama.

1. Select **Panorama** > **Managed WildFire Appliances** and **Deploy Master Key**.
2. Select all devices and **Change** the master key.
3. Configure the master key:
   1. If renewing a master key, enter the **Current Master Key**. If you are replacing the default master key with a new master key, do not specify a **Current Master Key**.
   2. Specify the **New Master Key** and **Confirm Master Key**.
   3. Configure the master key **Lifetime** and **Time for Reminder**.
   4. Click **OK**.

   *The new master key is automatically pushed to your WildFire appliances after you click **OK**. Proceed only if you are certain you are ready to change the master key for your WildFire appliances.*
4. Verify that the master key was deployed successfully to all selected devices.

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

STEP 7 | (HA Panorama only) Reconfigure the Panorama HA configuration.

Repeat this step for both the primary and secondary Panorama HA peers.

1. Select **Panorama** > **High Availability** > **General** and edit the HA Setup.
2. Enable (check) **Enable HA** and click **OK**.
3. **Commit** and **Commit to Panorama**.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | (HA Firewalls only) Enable config sync for managed firewalls.

1. Select **Device** > **High Availability** > **General** and select the **Template** containing the managed firewall HA configuration.
2. Edit the HA Pair Settings **Setup**.
3. Enable (check) **Enable Config Sync** and click **OK**.
4. **Commit** and **Commit and Push** your configuration changes.

 ©2025 Palo Alto Networks, Inc.

Schedule a Configuration Push to Managed Firewalls

Reduce the operational overhead of pushing configuration changes to managed firewalls by creating 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 on a regularly occurring schedule. This allows you to push configuration made by multiple administrators to multiple firewalls without the need for involvement of any administrator. A scheduled configuration push is supported for a target managed firewall running any PAN-OS release.

Superusers and custom Panorama admins with an appropriately defined admin role profile can create a scheduled configuration push to managed firewalls. To create a scheduled configuration push, 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 5 minute interval to allow for the Panorama management server to validate the configuration. Scheduled configuration pushes that are within 5 minutes of each other may fail due to Panorama being unable to validate the first scheduled configuration push changes.

Panorama performs the scheduled device group and template configuration push to managed firewalls if the **Device Groups** or **Templates** Last Commit Status is `out-of-sync`. 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 configuration change that interrupted the connection between the managed firewall and Panorama.

STEP 1 | Log in to the Panorama Web Interface.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Create a scheduled configuration push.

1. Select **Panorama** > **Scheduled Config Push** and **Add** a new scheduled configuration push.

   *You can also schedule a configuration push to managed firewalls when you push to devices (**Commit** > **Push to Devices**).*

2. Configure name and frequency of the scheduled configuration push.
   - **Name**—Name of the configuration push schedule.
   - **Admin Scope—**The admin for which the configuration changes will be pushed.

     The name of the logged in admin creating the scheduled config push is displayed by default. Click the admin name to add more Panorama admins to the scheduled config push.
   - **Date**—Date on which the configuration push is scheduled to occur next.
   - **Time**—Time (`hh:mm:ss`) at which the configuration push is scheduled to occur on the scheduled configuration push **Date**.
   - **Recurrence**—Whether the scheduled configuration push is a one time push or a recurring scheduled push (`monthly`, `weekly`, or `daily`).
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 config push.

   All managed firewalls associated with the selected device groups, templates, or template stacks are included in the scheduled config 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.

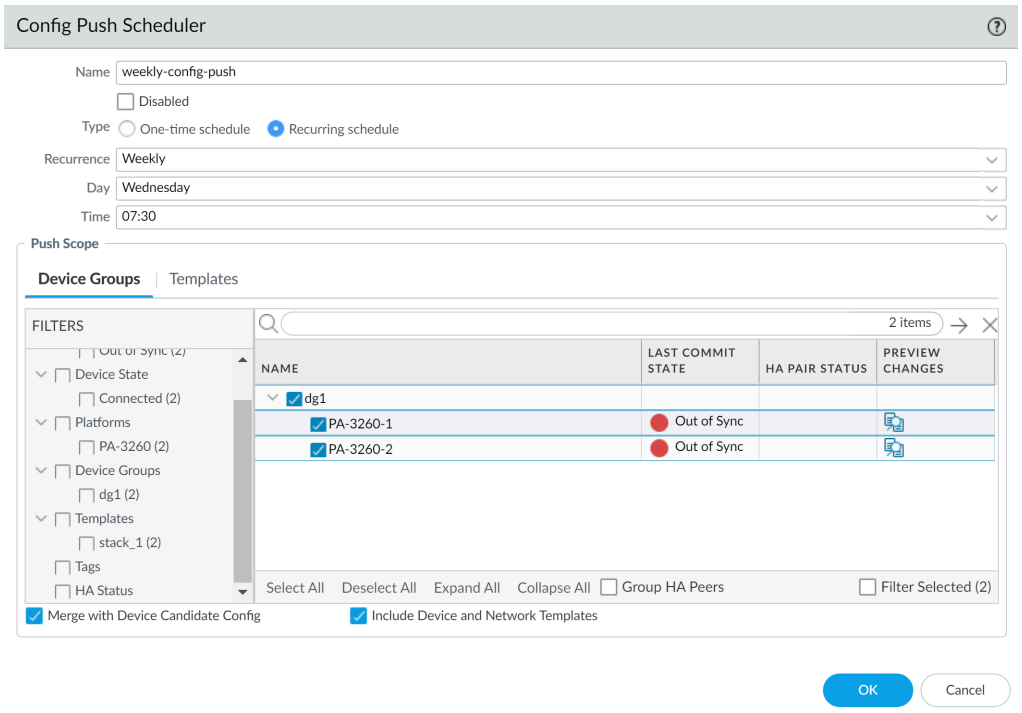
   3. Verify whether to **Merge with Device Candidate config** to merge the configuration changes pushed from Panorama with any pending configuration changes implemented locally on the firewall.

      This setting is enabled by default.
   4. Verify whether to **Include Device and Network Templates** to push both device group changes and the associate template changes in a single operation.

      This setting is enabled by default. If disabled, Panorama pushes the device group and associated template changes as separate operations.

   ***Force Template Values** is not supported for a scheduled configuration push to prevent outages during off hours caused by a configuration push that overwrites the local firewall configuration.*
4. Click **OK**.
5. Click **Commit** and **Commit to Panorama**.

 ©2025 Palo Alto Networks, Inc.



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.

   This includes the last time the scheduled configuration push occurred and the total number of impacted managed firewalls. Of the total number of impacted firewalls, you can view how many scheduled configuration pushes were successful, how many failed, and how many of the managed firewalls automatically reverted their configuration due to a configuration change that caused a disconnect between the managed firewall on Panorama.
3. Click **Tasks** to view the full operation details for the latest scheduled configuration push.

 ©2025 Palo Alto Networks, Inc.

Redistribute Data to Managed Firewalls

To ensure all the firewalls that enforce policies and generate reports have the required data and authentication timestamps for your policy rules, you can leverage your Panorama infrastructure to redistribute the mappings and timestamps.

- Configure the Panorama management server to redistribute data.
  1. Add firewalls, virtual systems, or Windows User-ID agents as redistribution agents to Panorama:
     1. Select **Panorama** > **Data Redistribution** and **Add** each redistribution agent.
     2. Enter a **Name** to identify the redistribution agent.
     3. Confirm that the agent is **Enabled**.
     4. Enter the **Host** name or IP address of the MGT interface on firewall.
     5. Enter the **Port** number on which the firewall will listen for data redistribution queries (default is 5007).
     6. If the redistribution agent is a firewall or virtual system, enter the **Collector Name** and **Collector Pre-Shared Key**.
     7. Select the **Data type** that you want to redistribute. You can select all data types, but you must select at least one of the following data types:
        - **IP User Mappings**
        - **IP Tags**
        - **User Tags**
        - **HIP**
        - **Quarantine List**
     8. Click **OK** to save the configuration.
  2. Enable the Panorama MGT interface to respond to data redistribution queries from firewalls:

     If the Panorama management server has a high availability (HA) configuration, perform this step on each HA peer as a best practice so that redistribution continues if Panorama fails over.

     1. Select **Panorama** > **Setup** > **Interfaces** and **Management**.
     2. Select **User-ID** in the Network Services section and click **OK**.
  3. Select **Commit** > **Commit to Panorama** to activate your changes on Panorama.

 ©2025 Palo Alto Networks, Inc.

- Configure firewalls to receive data that Panorama redistributes.
  1. Select **Device** > **Data Redistribution** > **Agents** then select the **Template** to which the firewalls are assigned.
  2. **Add** an agent and enter a **Name**.
  3. Select how you want to add the agent:
     - **Serial Number**—Select the **Serial Number** of the Panorama you want to use from the list:
       - **panorama**—The active or solitary Panorama
       - **panorama2**—(HA only) The passive Panorama
     - **Host and Port**—Specify the following information:
       - Select the **Host** name or IP address of the MGT interface on firewall.
       - Select whether the host is an **LDAP Proxy**.
       - Enter the **Port** number on which the firewall will listen for data redistribution queries (default is 5007).
       - If the redistribution agent is a firewall or virtual system, enter the **Collector Name** and **Collector Pre-Shared Key**.
       - Select the **Data type** that you want to redistribute.
  4. Confirm that the agent is **Enabled** and click **OK** to save the configuration.
  5. Select **Commit** > **Commit and Push** to activate your changes on Panorama and push the changes to the firewalls.

- Verify that Panorama and firewalls receive redistributed data.
  1. View the agent statistics **Panorama** > **Data Redistribution** > **Agents** and select **Status** to view a summary of the activity for the redistribution agent, such as the number of mappings that the client firewall has received.
  2. Confirm the **Source Name** in the User-ID logs (**Monitor** > **Logs** > **User-ID**) to verify that the firewall receives the mappings from the redistribution agents.
  3. View the IP-Tag log (**Monitor** > **Logs** > **IP-Tag**) to confirm that the client firewall receives data.
  4. Access the CLI of a firewall or Panorama management server that redistributes data.
  5. Display all the user mappings by running the following command:

     ```
     > show user ip-user-mapping all
     ```
  6. Record the IP address associated with any one username.
  7. Access the CLI of a firewall or Panorama management server that receives redistributed data.
  8. Display the mapping information and authentication timestamp for the *<IP-address>* you recorded:

     ```
     > show user ip-user-mapping ip <IP-address>
     IP address:    192.0.2.0 (vsys1)
     ```

 ©2025 Palo Alto Networks, Inc.

```
User:           corpdomain\username1
From:           UIA
Idle Timeout:   10229s
Max. TTL:       10229s
MFA Timestamp: first(1) - 2016/12/09 08:35:04
Group(s):       corpdomain\groupname(621)
```

This example output shows the timestamp for a response to one authentication challenge (factor). For Authentication rules that use multi-factor authentication (MFA), the output shows multiple timestamps.

 ©2025 Palo Alto Networks, Inc.

Transition a Firewall to Panorama Management

If you have already deployed Palo Alto Networks firewalls and configured them locally, but now want to use Panorama for centrally managing them, you must perform pre-migration planning. The migration involves importing firewall configurations into Panorama and verifying that the firewalls function as expected after the transition. If some settings are unique to individual firewalls, you can continue accessing the firewalls to manage the unique settings. You can manage any given firewall setting by pushing its value from Panorama or by configuring it locally on the firewall, but you cannot manage the setting through both Panorama and the firewall. If you want to exclude certain firewall settings from Panorama management, you can either:

- Migrate the entire firewall configuration and then, on Panorama, delete the settings that you will manage locally on firewalls. You can also Override a Template or Template Stack Value that Panorama pushes to a firewall instead of deleting the setting on Panorama.
- Load a partial firewall configuration, including only the settings that you will use Panorama to manage.



Firewalls do not lose logs during the transition to Panorama management.

- Plan the Transition to Panorama Management
- Migrate a Firewall to Panorama Management and Reuse Existing Configuration
- Migrate a Firewall to Panorama Management and Push a New Configuration
- Migrate a Firewall HA Pair to Panorama Management and Reuse Existing Configuration
- Migrate a Firewall HA Pair to Panorama Management and Push a New Configuration
- Load a Partial Firewall Configuration into Panorama
- Localize a Panorama Pushed Configuration on a Managed Firewall

Plan the Transition to Panorama Management

The following tasks are a high-level overview of the planning required to migrate firewalls to Panorama management:

- ☐ Decide which firewalls to migrate.
- ☐ Plan a maintenance window and ensure there are no pending configuration changes on Panorama or the firewalls.
- ☐ If you are migrating the firewall from one Panorama to another, localize the Panorama pushed configuration on the firewall.
- ☐ Preserve your known working Panorama and firewall configurations prior to migration.
  - Export the device state of your firewalls.
  - Export a named Panorama configuration shapshot of the running Panorama configuration.
- ☐ Determine the Panorama and firewall software and content versions, and how you will manage licenses and software upgrades. For important details, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.

 ©2025 Palo Alto Networks, Inc.

- Plan how to manage shared settings.

  Plan the Device Group Hierarchy, Templates and Template Stacks in a way that will reduce redundancy and streamline the management of settings that are shared among all firewalls or within firewall sets. During the migration, you can select whether to import objects from the Shared location on the firewall into Shared on Panorama, with the following exceptions:

  - If a shared firewall object has the same name and value as an existing shared Panorama object, the import excludes that firewall object.
  - If the name or value of the shared firewall object differs from an existing shared Panorama object, Panorama imports the firewall object into each new device group that is created for the import.
  - If a configuration imported into a template references a shared firewall object, or if a shared firewall object references a configuration imported into a template, Panorama imports the object as a shared object regardless of whether you select the **Import devices' shared objects into Panorama's shared context** check box.
- Determine if the firewall has configuration elements (policies, objects, and other settings) that you don't want to import, either because Panorama already contains similar elements or because those elements are firewall-specific (for example, timezone settings) and you won't use Panorama to manage them. You can perform a global find to determine if similar elements exist on Panorama.
- Decide the common zones for each device group. This includes a zone-naming strategy for the firewalls and virtual systems in each device group. For example, if you have zones called Branch LAN and WAN, Panorama can push policy rules that reference those zones without being aware of the variations in port or media type, model, or logical addressing schema.
- Create a post-migration test plan.

  You will use the test plan to verify that the firewalls work as efficiently after the migration as they did before. The plan might include tasks such as:

  - Monitor the firewalls for at least 24 hours after the migration.
  - Monitor Panorama and firewall logs for anomalies.
  - Check administrator logins on Panorama.
  - Test various types of traffic from multiple sources. For example, check bandwidth graphs, session counts, and deny-rule traffic log entries (see Use Panorama for Visibility). The testing should cover a representative sample of policy configurations.
  - Check with your network operations center (NOC) and security operations center (SOC) for any user-reported issues.
  - Include any other test criteria that will help verify firewall functionality.

Migrate a Firewall to Panorama Management and Reuse Existing Configuration

Migrate a firewall to Panorama management and import the existing firewall configuration to Panorama to reuse it. When you import a firewall configuration, Panorama automatically creates a template to contain the imported network and device settings. To contain the imported policies and objects, Panorama automatically creates one device group for each firewall or one device group for each virtual system (vsys) in a multi-vsys firewall.

 ©2025 Palo Alto Networks, Inc.

When you perform the following steps, Panorama imports the entire firewall configuration. Alternatively, you can Load a Partial Firewall Configuration into Panorama.

To migrate a firewall to Panorama management and create a new configuration, see Migrate a Firewall to Panorama Management and Push a New Configuration. To migrate a firewall HA pair to Panorama management, see Migrate a Firewall HA Pair to Panorama Management and Reuse Existing Configuration.

Panorama can import configurations from firewalls that run PAN-OS 5.0 or later releases and can push configurations to those firewalls. The exception is that Panorama 6.1 and later releases cannot push configurations to firewalls running PAN-OS 6.0.0 through 6.0.3.

Panorama can import configurations from firewalls that are already managed devices but only if they are not already assigned to device groups or templates.

STEP 1 | Plan the migration.

See the checklist in Plan the Transition to Panorama Management.

STEP 2 | Add the firewall as a managed device.

See Add a Firewall as a Managed Device for more information on adding a firewall to Panorama management.

1. Log in to the Panorama Web Interface
2. Select **Panorama** > **Device Registration Auth Key** and **Add** a new authentication key.
   Copy Auth Key after you successfully create the device registration authentication key.
3. Select **Panorama** > **Managed Devices** > **Summary** to **Add** a firewall as a managed device.
4. Enter the serial number of the firewall and click **OK**.

   If you will import multiple firewall configurations, enter the serial number of each one on a separate line. Optionally, you can copy and paste the serial numbers from a Microsoft Excel worksheet.
5. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 3 | Set up a connection from the firewall to Panorama.

1. Log in to the firewall web interface
2. Select **Device** > **Setup** > **Management** and edit the Panorama Settings.
3. In the **Panorama Servers** fields, enter the IP addresses of the Panorama management server.
4. Paste the **Auth Key** you copied in the previous step.
5. Click **OK** and **Commit**.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Import the firewall configuration into Panorama.

*If you later decide to re-import a firewall configuration, first remove the firewall device groups and template to which it is a member. If the device group and template names are the same as the firewall hostname, then you can delete the device group and template before re-importing the firewall configuration or use the **Device Group Name Prefix** fields to define new names for the device group and template created by the re-import. Additionally, firewalls don't lose logs when you remove them from device groups or templates.*

1. From Panorama, select **Panorama** > **Setup** > **Operations**, click **Import device configuration to Panorama**, and select the **Device**.

   Panorama can't import a configuration from a firewall that is assigned to an existing device group or template.

2. (Optional) Edit the **Template Name**. The default value is the firewall name. You can't use the name of an existing template or template stack.
3. (Optional) Edit the **Device Group** names. For a multi-vsys firewall, each device group has a vsys name by default, so add a character string as a Device Group Name Prefix for each. Otherwise, the default value is the firewall name. You can't use the names of existing device groups.

   *The **Import devices' shared objects into Panorama's shared context** check box is selected by default, which means Panorama compares imports objects that belong to the Shared location in the firewall to Shared in Panorama. If an imported object is not in the Shared context of the firewall, it is applied to each device group being imported. If you clear the check box, Panorama copies will not compare imported objects, and apply all shared firewall objects into device groups being imported instead of Shared. This could create duplicate objects, so selecting the check box is a best practice in most cases. To understand the consequences of importing shared or duplicate objects into Panorama, see* Plan how to manage shared settings.

4. 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.

   If Panorama has a rule with the same name as a firewall rule that you import, Panorama displays both rules. Delete one of the rules before performing a Panorama commit to prevent a commit error.

5. Click **OK**. Panorama displays the import status, result, details about your selections, details about what was imported, and any warnings. Click **Close**.
6. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Push the configuration bundle from Panorama to the newly added firewall to remove all policy rules and objects from its local configuration.

This step is necessary to prevent duplicate rule or object names, which would cause commit errors when you push the device group configuration from Panorama to the firewall in the next step.

*Pushing the imported firewall configuration from Panorama to remove local firewall configuration updates **Policy** rule **Creation** and **Modified** dates to reflect the date you pushed to your newly managed firewalls when you monitor policy rule usage for a managed firewall. Additionally, a new universially unique identifier (UUID) for each policy rule is created.*

This step is required to successfully migrate firewall management to the Panorama management server. Failure to perform this step successfully causes configuration errors and commit failures.

1. Log in to the Panorama Web Interface.
2. Select **Panorama** > **Setup** > **Operations** and **Export or push device config bundle**.
3. Select the **Device** from which you imported the configuration and click **OK**.

   *If a master key is configured, **Use Master Key** and enter the master key before you click **OK**.*
4. Select **Push & Commit**. Panorama pushes the bundle and initiates a commit on the firewall.
5. Click **Close** after the push has committed successfully.
6. Launch the Web Interface of the firewall and ensure that the configuration has been successfully committed. If not, **Commit** the changes locally on the firewall.
7. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 6 | Push the device group and template configurations to complete the transition to centralized management.

This step overwrites any local **Network** and **Device** settings configured on the firewall.

If you are migrating multiple firewalls, perform all the preceding steps—including this one—for each firewall before continuing.

1. Select **Commit** > **Commit and Push** and **Edit Selections** in the Push Scope.
2. Select **Device Groups** and select the device groups that contain the imported firewall configurations.
3. Select **Merge with Device Candidate Config**, **Include Device and Network Templates**, and **Force Template Values**.
4. Click **OK** to save your changes to the Push Scope.
5. **Commit and Push** your changes.

STEP 7 | On the Panorama web interface, select **Panorama** > **Managed Devices** > **Summary** and verify that the device group and template stack are in sync for the firewall. On the firewall

 ©2025 Palo Alto Networks, Inc.

web interface, verify that configuration objects display a green cog, signifying that the configuration object is pushed from Panorama.

STEP 8 | 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.
2. Update the device group and template configurations as needed based on the configuration audit and any warnings that Panorama displayed after the import. For example:
   - Delete redundant objects and policy rules.
   - Move or Clone a Policy Rule or Object to a Different Device Group.
   - Move firewalls to different device groups or templates.
   - Move a device group that Panorama created during the import to a different parent device group: Select **Panorama** > **Device Groups**, select the device group you want to move, select a new **Parent Device Group**, and click **OK**.

STEP 9 | Consolidate all the imported firewall configurations.

This step is required if you are migrating multiple firewalls.

1. After importing all the firewall configurations, update the device groups and templates as needed to eliminate redundancy and streamline configuration management: see Fine-tune the imported configuration. (You don't need to push firewall configuration bundles again.)
2. Configure any firewall-specific settings.

   If the firewalls will have local zones, you must create them before performing a device group or template commit; Panorama can't poll the firewalls for zone name or zone configuration. If you will use local firewall rules, ensure their names are unique (not duplicated in Panorama). If necessary, you can Override a Template or Template Stack Value with a firewall-specific value.
3. Commit and push your changes:
   1. Select **Commit** > **Commit and Push** and **Edit Selections** in the Push Scope.
   2. Select **Device Groups**, select the device groups you changed, and **Include Device and Network Templates**.
   3. Click **OK** to save your changes to the Push Scope.
   4. **Commit and Push** your changes.

STEP 10 | Perform your post-migration test plan.

Perform the verification tasks that you devised during the migration planning to confirm that the firewalls work as efficiently with the Panorama-pushed configuration as they did with their original local configuration: see Create a post-migration test plan.

 ©2025 Palo Alto Networks, Inc.

Migrate a Firewall to Panorama Management and Push a New Configuration

This procedure overwrites the local firewall configuration with the configuration changes pushed from Panorama.

Migrate a firewall to Panorama management and create a new Panorama-managed configuration using device groups and template stacks.

When you perform the following steps, Panorama imports the entire firewall configuration. Alternatively, you can Load a Partial Firewall Configuration into Panorama.

To migrate a firewall to Panorama management and reuse the existing configuration, see Migrate a Firewall to Panorama Management and Reuse Existing Configuration. To migrate a firewall HA pair to Panorama management, see Migrate a Firewall HA Pair to Panorama Management and Reuse Existing Configuration.

Panorama can import configurations from firewalls that run PAN-OS 5.0 or later releases and can push configurations to those firewalls. The exception is that Panorama 6.1 and later releases cannot push configurations to firewalls running PAN-OS 6.0.0 through 6.0.3.

Panorama can import configurations from firewalls that are already managed devices but only if they are not already assigned to device groups or templates.

STEP 1 | Plan the migration.

See the checklist in Plan the Transition to Panorama Management.

STEP 2 | Add the firewall as a managed device.

See Add a Firewall as a Managed Device for more information on adding a firewall to Panorama management.

1. Log in to the Panorama Web Interface
2. Select **Panorama** > **Device Registration Auth Key** and **Add** a new authentication key.
   Copy Auth Key after you successfully create the device registration authentication key.
3. Select **Panorama** > **Managed Devices** > **Summary** to **Add** a firewall as a managed device.
4. Enter the serial number of the firewall and click **OK**.

   To add multiple firewalls at the same time, enter the serial number of each one on a separate line.
5. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Set up a connection from the firewall to Panorama.

1. Log in to the firewall web interface
2. Select **Device** > **Setup** > **Management** and edit the Panorama Settings.
3. In the **Panorama Servers** fields, enter the IP addresses of the Panorama management server.
4. Paste the **Auth Key** you copied in the previous step.
5. Click **OK** and **Commit**.

STEP 4 | On the Panorama web interface, select **Panorama** > **Managed Devices** > **Summary** and verify that the `Device State` is `Connected`.

STEP 5 | Add a Device Group.

Repeat this step to create as many device groups as needed to logically group your firewall configurations. Device groups are required to manage device group objects and policies. Learn more about how to manage your device groups.

STEP 6 | Create a template and template stack.

Templates and template stacks are used to configure the firewall **Network** and **Device** settings that enable firewall to operate on the network.

1. Add a Template.

   Repeat this step to create as many templates as needed to define your required networking configurations.
2. Configure a Template Stack.

   Repeat this step to create as many template stacks as needed to quickly apply your defined networking configurations. When you create a template stack, assign the relevant templates and managed firewalls.

STEP 7 | Configure the device groups, templates, and template stacks as needed.

STEP 8 | Push the device group and template configurations to complete the transition to centralized management.

1. Select **Commit** > **Commit and Push**.
2. (Optional) Click **Edit Selections** to modify the Push Scope.

   - **Merge with Device Candidate Config**—This setting is enabled by default and merges any pending local firewall configurations with the configuration push from Panorama. The local firewall configuration is merged and committed regardless of the admin

 ©2025 Palo Alto Networks, Inc.

pushing the changes from Panorama or the admin who made the local firewall configuration changes.

Disable this setting if you manage and commit local firewall configuration changes independently of the Panorama managed configuration.

- **Force Template Values**—Overwrites any local firewall configurations with those in the template stack configuration pushed from Panorama in the event of conflicting values.

  This setting is enabled by default. Enable this setting to overwrite any conflicting firewall configurations with those defined in the template or template stack. Before enabling this setting, review any overridden values to ensure an outage does not occur.

Click **OK** to save your changes to the Push Scope.

3. **Commit and Push** your changes.

STEP 9 | Select **Panorama** > **Managed Devices** > **Summary** and verify that the `Shared Policy` and `Template` status is `In Sync` for the newly added firewalls.

On the firewall web interface, verify that configuration objects display a green cog, signifying that the configuration object is pushed from Panorama.

STEP 10 | Perform your post-migration test plan.

Perform the verification tasks that you devised during the migration planning to confirm that the firewalls work as efficiently with the Panorama-pushed configuration as they did with their original local configuration: see Create a post-migration test plan.

Migrate a Firewall HA Pair to Panorama Management and Reuse Existing Configuration

If you have a pair of firewalls in an HA configuration that you want to manage using Panorama, you have the option to import the configuration local to your firewall HA pair to Panorama without needing to recreate any configurations or policies. This allows you to reuse the existing firewall configuration. You first import the firewall configurations to Panorama, which are used to create a new device group and template. You will perform a special configuration push of the device group and template to the firewalls to overwrite the local firewall configurations and synchronize the firewalls with Panorama.

To migrate a firewall HA pair to Panorama management and create a new configuration, see Migrate a Firewall HA Pair to Panorama Management and Push a New Configuration.

Panorama can import configurations from firewalls that run PAN-OS 5.0 or later releases and can push configurations to those firewalls. The exception is that Panorama 6.1 and later releases cannot push configurations to firewalls running PAN-OS 6.0.0 through 6.0.3.

Panorama can import configurations from firewalls that are already managed devices but only if they are not already assigned to device groups or templates.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Plan the migration.

See the checklist in Plan the Transition to Panorama Management.

STEP 2 | Disable configuration synchronization between the HA peers.

Repeat these steps for both firewalls in the HA pair.

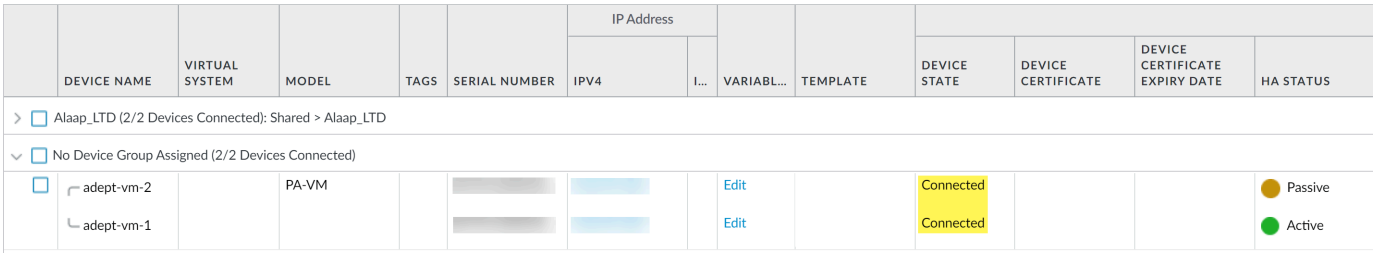
1. Log in to the web interface on each firewall, select **Device** > **High Availability** > **General** and edit the Setup section.
2. Clear **Enable Config Sync** and click **OK**.
3. **Commit** the configuration changes on each firewall.

STEP 3 | Add your HA firewalls to Panorama management.

Confirm that **Panorama Policy and Objects** and **Device and Network Template** are enabled.

If Panorama is already receiving logs from these firewalls, you do not need to perform this step. Continue to Step 7.

STEP 4 | Log in to the Panorama web interface and select **Panorama** > **Managed Devices** > **Summary** and verify the Device State for each firewall is `Connected`.



 ©2025 Palo Alto Networks, Inc.

STEP 5 | Import each firewall configuration into Panorama.

Do no push any device group or template stack configuration to your managed firewalls in this step. Pushing the device group and template stack configuration during this step wipes the local firewall HA configuration in the next steps.

*If you later decide to re-import a firewall configuration, first remove the firewall device groups and template to which it is a member. If the device group and template names are the same as the firewall hostname, then you can delete the device group and template before re-importing the firewall configuration or use the **Device Group Name Prefix** fields to enter a new name for the device group and template created by the re-import. Additionally, firewalls don't lose logs when you remove them from device groups or templates.*

1. From Panorama, select **Panorama** > **Setup** > **Operations**, click **Import device configuration to Panorama**, and select the **Device**.

   Panorama can't import a configuration from a firewall that is assigned to an existing device group or template stack.

2. (Optional) Edit the **Template Name**. The default value is the firewall name. You can't use the name of an existing template or template stack.
3. (Optional) Edit the **Device Group** names. For a multi-vsys firewall, each device group has a vsys name by default, so add a character string as a Device Group Name Prefix for each. Otherwise, the default value is the firewall name. You can't use the names of existing device groups.

   *The **Imported devices' shared objects into Panorama's shared context** check box is selected by default, which means Panorama compares imports objects that belong to the Shared location in the firewall to Shared in Panorama. If an imported object is not in the Shared context of the firewall, it is applied to each device group being imported. If you clear the check box, Panorama copies will not compare imported objects, and apply all shared firewall objects into device groups being imported instead of Shared. This could create duplicate objects, so selecting the check box is a best practice in most cases. To understand the consequences of importing shared or duplicate objects into Panorama, see* Plan how to manage shared settings.

4. **Commit to Panorama**.
5. Select **Panorama** > **Setup** > **Operations** and **Export or push device config bundle**. Select the **Device**, select **OK** and **Push & Commit** the configuration.

   The Enable Config Sync setting in Step 2 must be cleared on both firewalls before you push the device group and template stack.

6. Launch the Web Interface of the firewall HA peer and ensure that the configuration pushed in the previous step committed successfully. If not, **Commit** the changes locally on the firewall.
7. Repeat Step 1-6 above on the second firewall. The process creates a device group and template stack per each firewall.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Add the HA firewall pair into the same device group and template stack.

(Firewalls in active/active configuration) It is recommended to add HA peers to the same device group but not to the same template stack because firewalls in an active/active HA configuration typically need unique network configurations. This simplifies policy management for the HA peers while reducing the operational burden of managing the network configuration of each HA peer when their network configurations are independent of each other. For example, firewalls in an active/active HA configuration often times need unique network configurations, such as unique floating IP that are used as the default gateway for hosts.

Ultimately, deciding whether to add firewalls in an active/active HA configuration to the same device group and template stack is a design decision you must make when designing your configuration hierarchy.

1. Select **Panorama** > **Device Group**, select the device group of the second firewall, and remove the second firewall from the device group.
2. Select the device group from which you removed the second firewall and **Delete** it.
3. Select the device group for the first firewall, select the second firewall, click **OK** and **Commit to Panorama** to add it to the same device group as the HA peer.
4. Select **Panorama** > **Templates**, select the template stack of the second firewall, and remove the second firewall from the template stack.
5. Select the template stack from which you removed the second firewall and **Delete** it.
6. Select the template stack for the first firewall, add the second firewall, select **OK** and **Commit to Panorama** to add it to the same template stack as the HA peer.
7. (Optional) Remove the HA settings in the template associated with the newly migrated firewalls.

   If both members of the HA pair are in the same template stack, then you must remove the HA configuration from the template.

   You can manage the firewall HA configuration from Panorama or configure the HA settings locally on the managed firewalls.

   Skip this step if you want to manage the firewall HA settings from Panorama.

   1. Select **Device** > **High Availability** and select the **Template** containing the HA configuration.
   2. Select **Remove All**.
   3. **Commit to Panorama**.
8. Select **Panorama** > **Managed Devices** > **Summary**, and verify that the device group and template are in sync for the passive firewall. Verify policy rules, objects and network settings on the passive firewall match the active firewall.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Push the device group and template stack configuration changes to your managed firewalls.

You must first push the device group and template stack configuration to your `passive` or `Active-Secondary` HA peer first and then to the `active` or `Active-Primary` HA peer.

*Pushing the imported firewall configuration from Panorama to remove local firewall configuration updates **Policy** rule **Creation** and **Modified** dates to reflect the date you pushed to your newly managed firewalls when you monitor policy rule usage for a managed firewall. Additionally, a new universially unique identifier (UUID) for each policy rule is created.*

1. Log into the firewall web interface of the `Passive` or `Active-Secondary` HA peer and select **Device** > **High Availability** > **Operational Commands** to **Suspend local device for high availability**.
2. Push the Panorama managed configuration to the `suspended` HA firewall.
   1. Log in to the Panorama web interface.
   2. Select **Commit** > **Push to Devices** and **Edit Selections**.
   3. Enable (select) **Merge Device Candidate Config** and **Include Device and Network Templates**.

      (Panorama-managed HA configuration) Enable (select) **Force Template Values**.
   4. In **Device Groups** and **Templates**, select the suspended HA firewall.
   5. Click **OK** and **Push**.
3. In the firewall web interface of the suspended `passive` or `Active-Secondary` HA peer and select **Device** > **High Availability** > **Operational Commands** to **Make local device functional for high availability**.
4. Log into the firewall web interface of the `active` or `Active-Primary` HA peer and select **Device** > **High Availability** > **Operational Commands** to **Suspend local device for high availability**.
5. Repeat Step 2 to push the Panorama managed configuration to the `suspended` HA peer.
6. Log into the firewall web interface of the suspended `active` or `Active-Primary` HA peer and select **Device** > **High Availability** > **Operational Commands** to **Make local device functional for high availability**.
7. In the Panorama web interface, select **Panorama** > **Managed Devices** > **Summary**, and verify that the device group and template are in sync for HA firewalls. Verify policy rules, objects and network settings on the passive firewall match the active firewall.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | (Local firewall HA configuration only) Enable configuration synchronization between the HA peers.

Repeat these steps for both firewalls in the HA pair if you plan on maintaining a local configuration that needs to be synchronized.

Skip this step if managing the firewall HA configuration from Panorama. This setting is enabled by default.

1. Log in to the web interface on each firewall, select **Device** > **High Availability** > **General** and edit the Setup section.
2. Select **Enable Config Sync** and click **OK**.
3. **Commit** the configuration changes on each firewall.

Migrate a Firewall HA Pair to Panorama Management and Push a New Configuration

This procedure overwrites the local firewall configuration with the configuration pushed from Panorama.

Migrate a firewall high availability (HA) pair to Panorama management and create a new Panorama-managed configuration using device groups and template stacks.

To migrate a firewall HA pair to Panorama management and reuse the existing configuration, see Migrate a Firewall HA Pair to Panorama Management and Reuse Existing Configuration.

Panorama can import configurations from firewalls that run PAN-OS 5.0 or later releases and can push configurations to those firewalls. The exception is that Panorama 6.1 and later releases cannot push configurations to firewalls running PAN-OS 6.0.0 through 6.0.3.

Panorama can import configurations from firewalls that are already managed devices but only if they are not already assigned to device groups or templates.

STEP 1 | Plan the migration.

See the checklist in Plan the Transition to Panorama Management.

STEP 2 | Disable configuration synchronization between the HA peers.

Repeat these steps for both firewalls in the HA pair.

1. Log in to the web interface on each firewall, select **Device** > **High Availability** > **General** and edit the Setup section.
2. Clear **Enable Config Sync** and click **OK**.
3. **Commit** the configuration changes on each firewall.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Add the firewall as a managed device.

See Add a Firewall as a Managed Device for more information on adding a firewall to Panorama management.

1. Log in to the Panorama Web Interface
2. Select **Panorama** > **Device Registration Auth Key** and **Add** a new authentication key.

   Copy Auth Key after you successfully create the device registration authentication key.
3. Select **Panorama** > **Managed Devices** > **Summary** to **Add** a firewall as a managed device.
4. Enter the serial number of each firewall in the HA pair and click **OK**.

   To add multiple firewalls at the same time, enter the serial number of each one on a separate line.
5. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 4 | Set up a connection from the firewall to Panorama.

Repeat these steps for both firewalls in the HA pair.

1. Log in to the firewall web interface
2. Select **Device** > **Setup** > **Management** and edit the Panorama Settings.
3. In the **Panorama Servers** fields, enter the IP addresses of the Panorama management server.
4. Paste the **Auth Key** you copied in the previous step.
5. Click **OK** and **Commit**.

STEP 5 | On the Panorama web interface, select **Panorama** > **Managed Devices** > **Summary** and verify that the `Device State` is `Connected`.

STEP 6 | Add a Device Group.

Repeat this step to create as many device groups as needed to logically group your firewall configurations. Device groups are required to manage device group objects and policies. Learn more about how to manage your device groups.

You must add the HA peers to the same device group.

STEP 7 | Create a template and template stack.

Templates and template stacks are used to configure the firewall **Network** and **Device** settings that enable firewall to operate on the network.

1. Add a Template.

   Repeat this step to create as many templates as needed to define your required networking configurations.
2. Configure a Template Stack.

   Repeat this step to create as many template stacks as needed to quickly apply your defined networking configurations. When you create a template stack, assign the relevant templates and managed firewalls.

   You must add the HA peers to the same template stack.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Configure the device groups, templates, and template stacks as needed.

STEP 9 | Push the device group and template stack configuration changes to your managed firewalls.

You must first push the device group and template stack configuration to your `passive` or `Active-Secondary` HA peer first and then to the `active` or `Active-Primary` HA peer.

1. Log into the firewall web interface of the `Passive` or `Active-Secondary` HA peer and select **Device** > **High Availability** > **Operational Commands** to **Suspend local device for high availability**.
2. Push the Panorama managed configuration to the `suspended` HA firewall.
   1. Log in to the Panorama web interface.
   2. Select **Commit** > **Push and Push** and **Edit Selections** to modify the Push Scope.
      - **Merge with Device Candidate Config**—This setting is enabled by default and merges any pending local firewall configurations with the configuration push from Panorama. The local firewall configuration is merged and committed regardless of the admin pushing the changes from Panorama or the admin who made the local firewall configuration changes.

        Disable this setting if you manage and commit local firewall configuration changes independently of the Panorama managed configuration.
      - **Force Template Values**—Overwrites any local firewall configurations with those in the template stack configuration pushed from Panorama in the event of conflicting values.

        This setting is enabled by default. Enable this setting to overwrite any conflicting firewall configurations with those defined in the template or template stack. Before enabling this setting, review any overridden values to ensure an outage does not occur.
   3. In **Device Groups** and **Templates**, select the suspended HA firewall.
   4. Click **OK** and **Push**.
3. In the firewall web interface of the suspended `passive` or `Active-Secondary` HA peer and select **Device** > **High Availability** > **Operational Commands** to **Make local device functional for high availability**.
4. Log into the firewall web interface of the `active` or `Active-Primary` HA peer and select **Device** > **High Availability** > **Operational Commands** to **Suspend local device for high availability**.
5. Repeat Step 2 to push the Panorama managed configuration to the `suspended` HA peer.
6. Log into the firewall web interface of the suspended `active` or `Active-Primary` HA peer and select **Device** > **High Availability** > **Operational Commands** to **Make local device functional for high availability**.
7. In the Panorama web interface, select **Panorama** > **Managed Devices** > **Summary**, and verify that the device group and template are in sync for HA firewalls. Verify policy rules, objects and network settings on the passive firewall match the active firewall.

 ©2025 Palo Alto Networks, Inc.

STEP 10 | Select **Panorama** > **Managed Devices** > **Summary** and verify that the `Shared Policy` and `Template` status is `In Sync` for the newly added firewalls.

On the firewall web interface, verify that configuration objects display a green cog, signifying that the configuration object is pushed from Panorama.

STEP 11 | Perform your post-migration test plan.

Perform the verification tasks that you devised during the migration planning to confirm that the firewalls work as efficiently with the Panorama-pushed configuration as they did with their original local configuration: see Create a post-migration test plan.

Load a Partial Firewall Configuration into Panorama

If some configuration settings on a firewall are common to other firewalls, you can load those specific settings into Panorama and then push them to all the other firewalls or to the firewalls in particular device groups and templates.

Loading a configuration into a Panorama management server requires a full commit and must be performed by a superuser. Full commits are required when performing certain Panorama operations, such as reverting and loading a configuration snapshot, and are not supported for custom Admin Role profiles.

STEP 1 | Plan the transition to Panorama.

See the checklist in Plan the Transition to Panorama Management.

STEP 2 | Resolve how to manage duplicate settings, which are those that have the same names in Panorama as in a firewall.

Before you load a partial firewall configuration, Panorama and that firewall might already have duplicate settings. Loading a firewall configuration might also add settings to Panorama that are duplicates of settings in other managed firewalls.

If Panorama has policy rules or objects with the same names as those on a firewall, a commit failure will occur when you try to push device group settings to that firewall. If Panorama has template settings with the same names as those on a firewall, the template values will override the firewall values when you push the template.

1. On Panorama, perform a global find to determine if duplicate settings exist.
2. Delete or rename the duplicate settings on the firewall if you will use Panorama to manage them, or delete or rename the duplicate settings on Panorama if you will use the firewall to manage them. If you will use the firewall to manage device or network settings, instead of deleting or renaming the duplicates on Panorama, you can also push the settings from Panorama (Step 6) and then Override a Template or Template Stack Value on the firewall with firewall-specific values.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Export the entire firewall configuration to your local computer.

1. On the firewall, select **Device** > **Setup** > **Operations**.
2. Click **Save named configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**.
3. Click **Export named configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**. The firewall exports the configuration as an XML file.

STEP 4 | Import the firewall configuration snapshot into Panorama.

1. On Panorama, select **Panorama** > **Setup** > **Operations**.
2. Click **Import named Panorama configuration snapshot**, **Browse** to the firewall configuration file you exported to your computer, and click **OK**.

   After using this option to import a firewall configuration file, you can't use the Panorama web interface to load it. You must use the XML API or CLI, as described in the next step.

STEP 5 | Load the desired part of the firewall configuration into Panorama.

To specify a part of the configuration (for example, all application objects), you must identify the:

- Source xpath—The XML node in the firewall configuration file from which you are loading.
- Destination xpath—The node in the Panorama configuration to which you are loading.

Use the XML API or CLI to identify and load the partial configuration:

1. Use the firewall XML API or CLI to identify the source xpath.

   For example, the xpath for application objects in vsys1 of the firewall is:

   ```
   /config/devices/entry[@name='localhost.localdomain']/vsys/
   entry[@name='vsys1']/application
   ```

2. Use the Panorama XML API or CLI to identify the destination xpath.

   For example, to load application objects into a device group named US-West, the xpath is:

   ```
   /config/devices/entry[@name='localhost.localdomain']/device-
   group/entry[@name='US-West']/application
   ```

3. Use the Panorama CLI to load the configuration and commit the change:

   ```
   # load config partial mode [append|merge|replace]
    from-xpath <source-xpath> to-xpath <destination-xpath>
    from <filename>
   ```

 ©2025 Palo Alto Networks, Inc.

```
# commit
```

For example, enter the following to load the application objects from vsys1 on an imported firewall configuration named fw1-config.xml into a device group named US-West on Panorama:

```
# load config partial mode merge from-xpath /config/
devices/entry[@name='localhost.localdomain']/vsys/
entry[@name='vsys1']/application to-xpath /config/
devices/entry[@name='localhost.localdomain']/device-group/
entry[@name='US-West']/application from fw1-config.xml
# commit
```

STEP 6 | Push the partial configuration from Panorama to the firewall to complete the transition to centralized management.

1. On the firewall, delete any rules or objects that have the same names as those in Panorama. If the device group for that firewall has other firewalls with rules or objects that are duplicated in Panorama, perform this step on those firewalls also. For details, see Step 2.
2. On Panorama, push the partial configuration to the firewall.
   1. Select **Commit** > **Commit and Push** and **Edit Selections** in the Push Scope.
   2. Select **Device Groups** and select the device groups that contain the imported firewall configurations.
   3. Select **Merge with Device Candidate Config**, **Include Device and Network Templates**, and **Force Template Values**.
   4. Click **OK** to save your changes to the Push Scope.
   5. **Commit and Push** your changes.
3. If the firewall has a device or network setting that you won't use Panorama to manage, Override a Template or Template Stack Value on the firewall.

STEP 7 | Perform your post-migration test plan.

Perform the verification tasks that you devised during the migration planning to confirm that the firewall works as efficiently with the Panorama-pushed configuration as it did with its original local configuration: see Create a post-migration test plan.

Localize a Panorama Pushed Configuration on a Managed Firewall

You can localize the template and device group configurations pushed from the Panorama™ management server to:

- Remove the firewall from Panorama management.
- Migrate firewall management to a different Panorama.
- In the case of an emergency where Panorama isn't accessible, ensure that administrators can modify the managed firewall configuration locally.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Launch the web interface of the managed firewall as an administrator with the Superuser role. You can directly access the firewall by entering its IP address in the browser URL field or, in Panorama, select the firewall in the **Context** drop-down.

STEP 2 | (Best Practice) Select **Device** > **Setup** > **Operations** and **Export device state**.

Save a copy of the firewall system state, including device group and template settings pushed from Panorama, in the event you need to reload a known working configuration on the managed firewall.

STEP 3 | (Active/passive HA only) Disable configuration synchronization for firewalls in an active/passive high availability (HA) configuration.

Repeat this step on each firewall HA peer. This is required to prevent duplication of objects on the passive HA peer that results in local commit failures.

1. Log in to the firewall web interface of one of the HA peers.
2. Select **Device** > **High Availability** > **General** and edit the HA pair Settings Setup.
3. Disable (uncheck) **Enable Config Sync** and click **OK**.
4. Select **Commit** and **Commit** your changes.

STEP 4 | Disable the template configuration to stop using template and template stacks to manage the network configuration objects of the managed firewall.

1. Select **Device** > **Setup** > **Management** and edit the Panorama Settings.
2. Click **Disable Device and Network Template**.
3. (Optional) Select **Import Device and Network Template before disabling** to save the template configuration settings locally on the firewall. If you don't select this option, PAN-OS deletes all Panorama-pushed settings from the firewall.
4. Click **OK** twice to continue.

STEP 5 | Disable the device group configuration to stop using a device group to manage the policy rule and object configurations of the managed firewall.

1. Select **Device** > **Setup** > **Management** and edit the Panorama Settings.
2. (Optional) Select **Import Panorama Policy Objects before disabling** to save the policy rule and object configurations locally on the firewall. If you don't select this option, PAN-OS deletes all Panorama-pushed configurations from the firewall.
3. Click **OK** to continue.

Don't attempt to commit your configuration changes on the managed firewall yet as all commits fail until the following steps are successfully completed.

STEP 6 | Select **Device** > **Setup** > **Operations** and **Save named configuration snapshot**.

STEP 7 | **Load named configuration snapshot** and enable (check) **Regenerate Rule UUIDs for selected named configuration** to generate new policy rule UUIDs.

This step is required to successfully localize the Panorama-pushed policy rules on the managed firewalls.

STEP 8 | Click **OK** to load the named configuration snapshot.

 ©2025 Palo Alto Networks, Inc.

STEP 9 | **Commit** the named configuration snapshot load.

STEP 10 | (Active/passive HA only) Enable configuration synchronization for firewalls in an active/passive high HA configuration.

Repeat this step each firewall HA peer.

1. Log in to the firewall web interface of one of the HA peers.
2. Select **Device** > **High Availability** > **General** and edit the HA pair Settings Setup.
3. Enable (check) **Enable Config Sync** and click **OK**.
4. Select **Commit** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

Device Monitoring on Panorama

After adding your firewalls and configuring policy rules, you can monitor the health status to ensure that your firewalls are operating within healthy parameters. For policy rules, monitor rule traffic matches to identify which rules match your traffic enforcement needs.

- Monitor Device Health
- Monitor Policy Rule Usage

Monitor Device Health

Monitor the health information of your managed firewalls to identify and resolve hardware issues before they impact your network security. Both Panorama™ and the managed firewalls must be running PAN-OS® 8.1 or later releases but firewalls do not need to be part of a device group or template stack to monitor their summary session, logging, resource, and environmental performance. Panorama stores the last 90 days of health monitoring statistics of your managed firewalls so when you select a firewall, you can view the time-trended graphs and tables for sessions, environmentals, interfaces, logging, resources, and high availability performance.

Panorama calculates the Baseline performance of each metric using seven-day averages and standard deviation to determine a normal operating range for the specific firewall. This is the value displayed when you view the high-level overview of your managed device health data. You can click on a Device, CPS, Session, Data Plane, Management Plane, or Logging Rate health metric value to **View Snapshot**. This shows detailed health data for that specific metric, including the Baseline, 24-hour, 7-day, and 15-day averages. When you view the health metric Snapshot, In addition to tracking the baseline and comparing time-trended performance, you can view which firewalls have deviating metrics and isolate performance-related issues before they impact your network. When Panorama identifies that a metric is outside the normal operating range, it marks the metric and populates the Deviating Devices tab with the deviating firewall.

The health monitoring data is stored on Panorama, and is preserved in the event a firewall is removed. When a firewall is removed from Panorama management, the health monitoring data no longer display but are preserved for 90 days. After 90 days, all health monitoring data of the removed firewall are removed from Panorama. If a firewall is added back to Panorama management, the latest health monitoring data from when the firewall was removed is displayed.

STEP 1 | Log in to the Panorama Web Interface.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Select **Panorama** > **Managed Devices** > **Health** to monitor the health of managed firewalls.

View **All Devices** to see a list of all managed firewalls and the monitored health metrics. Select an individual firewall to view Detailed Device View with time-trended graphs and tables of monitored metrics.

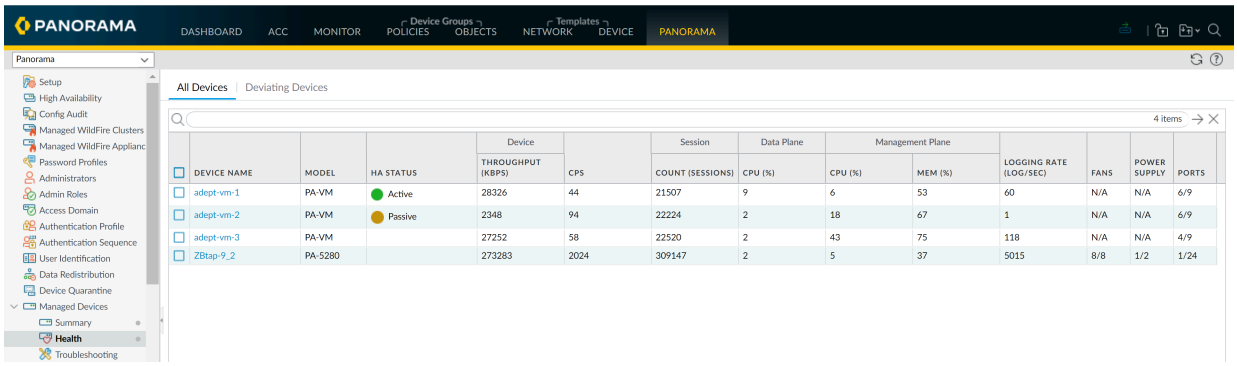


Figure 11: Managed Firewall Health Monitoring

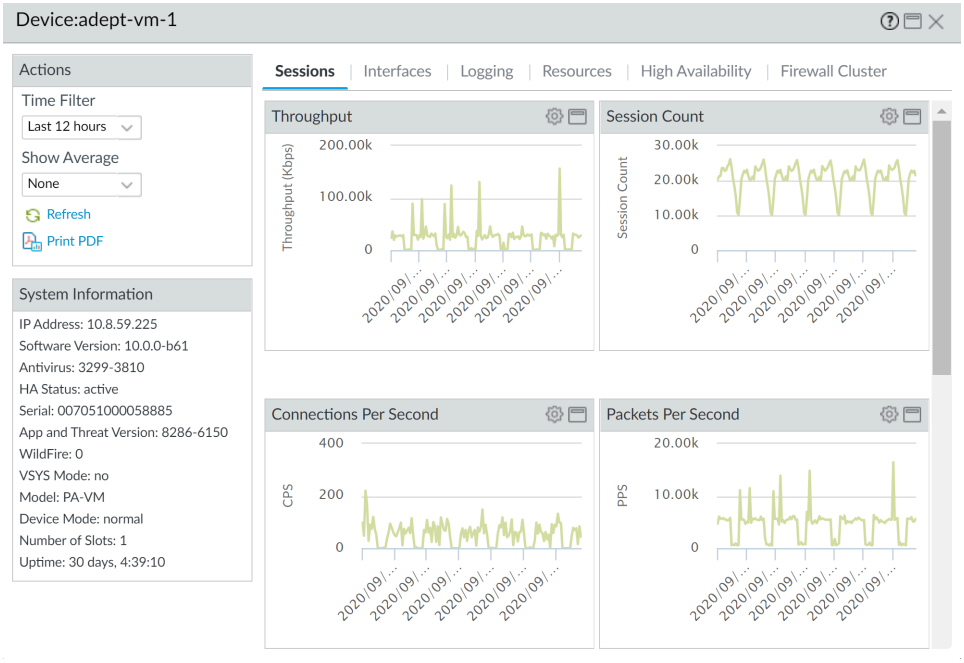


Figure 12: Detailed Device View

 ©2025 Palo Alto Networks, Inc.

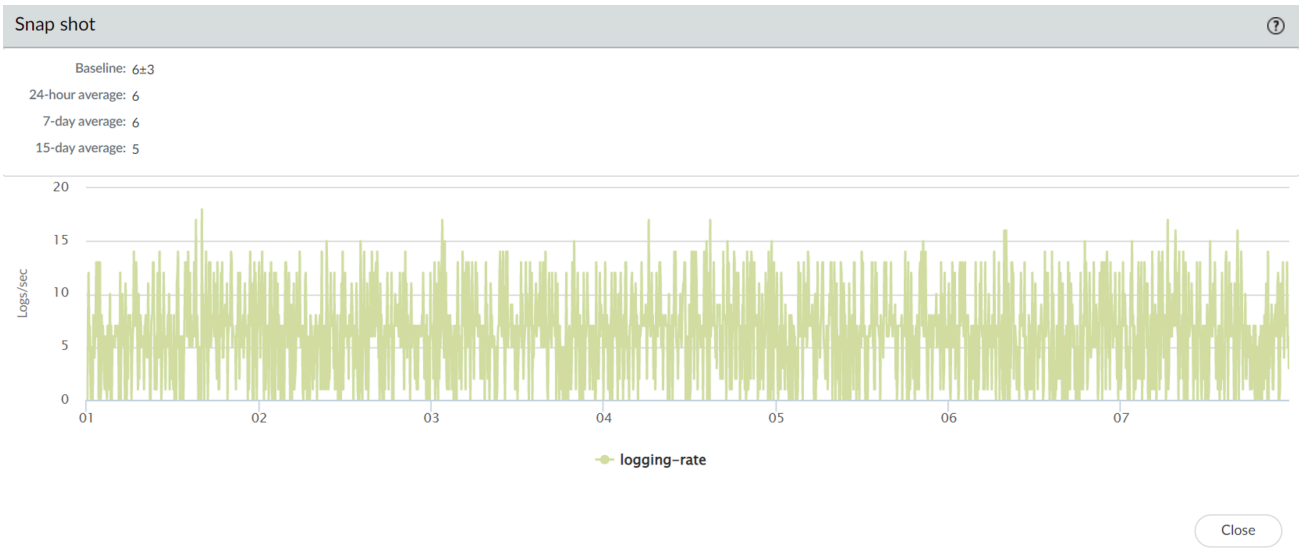
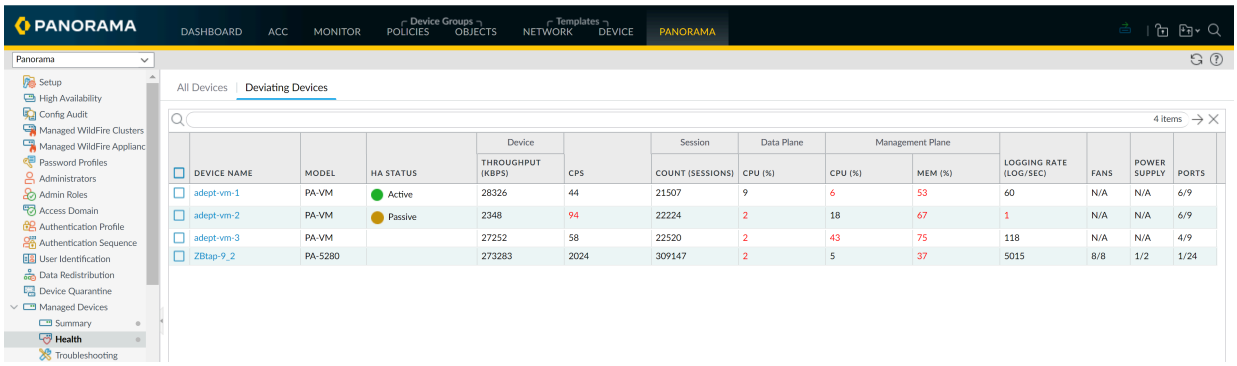


Figure 13: Metric Snap shot

STEP 3 | Select **Deviating Devices** to view firewalls with health metrics that deviated outside of the calculated baseline.

Panorama lists all firewalls that are reporting metrics that deviate from the calculated baseline and displays deviating metrics in red.



Monitor Policy Rule Usage

As your policies change, tracking rule usage on Panorama helps you evaluate whether your policy implementation continues to match your enforcement needs. This visibility helps you identify and remove unused rules to reduce security risks and keep your policy rule base organized. Additionally, rule usage tracking allows you to quickly validate new rule additions and rule changes and to monitor rule usage for operations and troubleshooting tasks. On Panorama, you can view the rule usage of firewalls in a device group—to which you pushed policies—to determine if all, some, or none of the firewalls have traffic matches instead of being able to monitor only the total number of hits across all firewalls in a device group. You can quickly filter rules using the rule usage data, such as Created and Modified dates, within a customizable time frame. The displayed rule usage information persists across reboot, dataplane restarts, and upgrades.

 ©2025 Palo Alto Networks, Inc.

On Panorama, you can view the rule usage details for managed firewalls that are running a PAN-OS 8.1 or later release, that have policy rule hit count enabled (default), and for which you have defined and pushed policy rules using device groups. Panorama cannot retrieve rule usage details for policy rules configured locally on the firewall so you must log in to the firewall to view rule usage information for locally configured rules.

After filtering your policy rulebase, administrators can take action to delete, disable, enable, and tag policy rules directly from the policy optimizer. For example, you can filter for unused rules and then tag them for review to determine whether they can be safely deleted or kept in the rulebase. By enabling administrators to take action directly from the policy optimizer, you reduce the management overhead required to further assist in simplifying your rule lifecycle management and ensure that your firewalls are not over-provisioned.

Policy rule usage data may also be useful when using Policy Optimizer to prioritize which rules to migrate or clean up first.

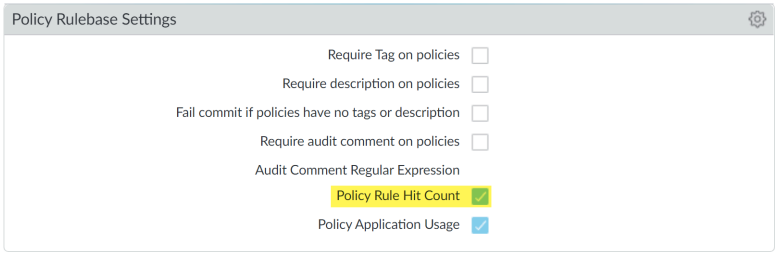
Policy rule usage data is not preserved when you transition to a different Panorama model. This means that all existing policy rule usage data from the old Panorama is no longer displayed after a successful migration to a new Panorama model. After a successful migration, Panorama begins tracking policy rule usage data based on the date the migration was completed. For example, the `Created` date displays the date the migration was completed.

To view the rule usage across any Shared rule or for a specific device group:

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Verify that the **Policy Rule Hit Count** is enabled.

1. Navigate to Policy Rulebase Settings (**Panorama** > **Setup** > **Management**.
2. Verify that **Policy Rule Hit Count** is enabled.



STEP 3 | Select **Policies** > ***<policy rule>*** to view a rule.

STEP 4 | Change the Device Group context to **Shared** or to the specific device group you want to view.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Determine whether the rule is being used (Rule Usage). The policy rule usage status is one of the following:

Firewalls must run PAN-OS 8.1 or later release with Policy Rule Hit Count enabled for Panorama to determine rule usage.

- **Used**—When all firewalls in the device group—to which you pushed the policy rule—have traffic matches for the policy rule.
- **Partially Used**—When some of the firewalls in the device group—to which you pushed the policy rule—have traffic matches for the policy rule.
- **Unused**—When no firewalls in the device group—to which you pushed the policy rule—have traffic matches for the policy rule.
- **Em-dash (—)**—When no firewalls in the device group—to which you pushed the policy rule—have Policy Rule Hit Count enabled or available for Panorama to determine the rule usage.
- **Modified**—The date and time the policy rule was last modified.
- **Created**—The date and time the policy rule was created.

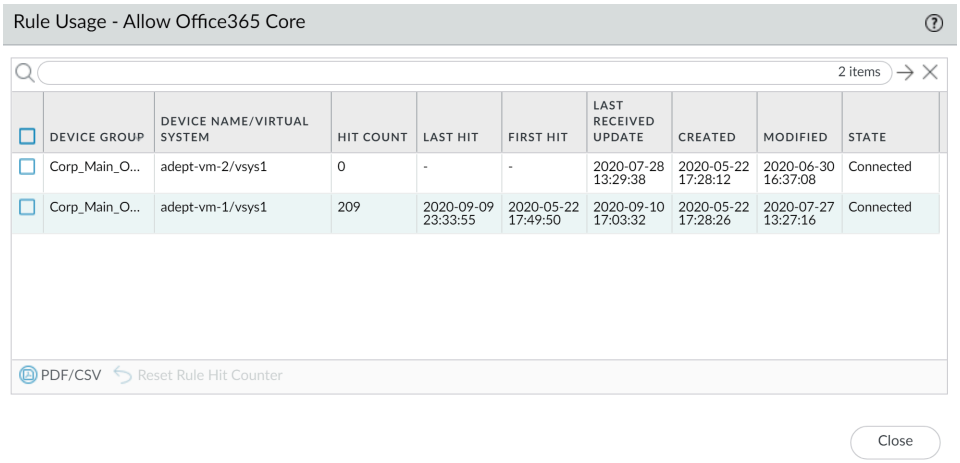


If the rule was created when Panorama was running PAN-OS 8.1 and the Policy Rule Hit Count setting is enabled, the First Hit date and time is used as the Created date and time on upgrade to PAN-OS 9.0 or later releases. If the rule was created in PAN-OS 8.1 when the Policy Rule Hit Count setting was disabled or if the rule was created when Panorama was running PAN-OS 8.0 or an earlier release, the Created date for the rule will be the date and time you successfully upgraded Panorama to PAN-OS 9.0 or later releases.

Rule Usage		DAYS WITH NO NEW APPS	MODIFIED	CREATED
RULE USAGE	APPS SEEN			
Used	6	150	2020-06-24 10:34:...	2020-04-09 11:34:03
Unused	0	-	2020-06-24 10:34:...	2020-04-16 11:42:46
Used	11	57	2020-06-24 10:34:...	2020-04-16 11:42:46
Partially Used	3	111	2020-06-24 10:34:...	2020-05-22 17:26:44
Unused	0	-	2020-06-24 10:34:...	2020-05-22 22:45:53

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Click the Rule Usage status to view the list of firewalls using the rule and the hit-count data for traffic that matches that rule on each firewall.



STEP 7 | (Optional) View the policy rule hit-count data for individual firewalls in the device group.

1. Click **Preview Rules**.
2. From the Device context, select the firewall for which you want to view the policy rule usage data.

STEP 8 | Select **Policies** and, in the Policy Optimizer dialog, view the **Rule Usage** filter.

STEP 9 | Filter rules in the selected rulebase.

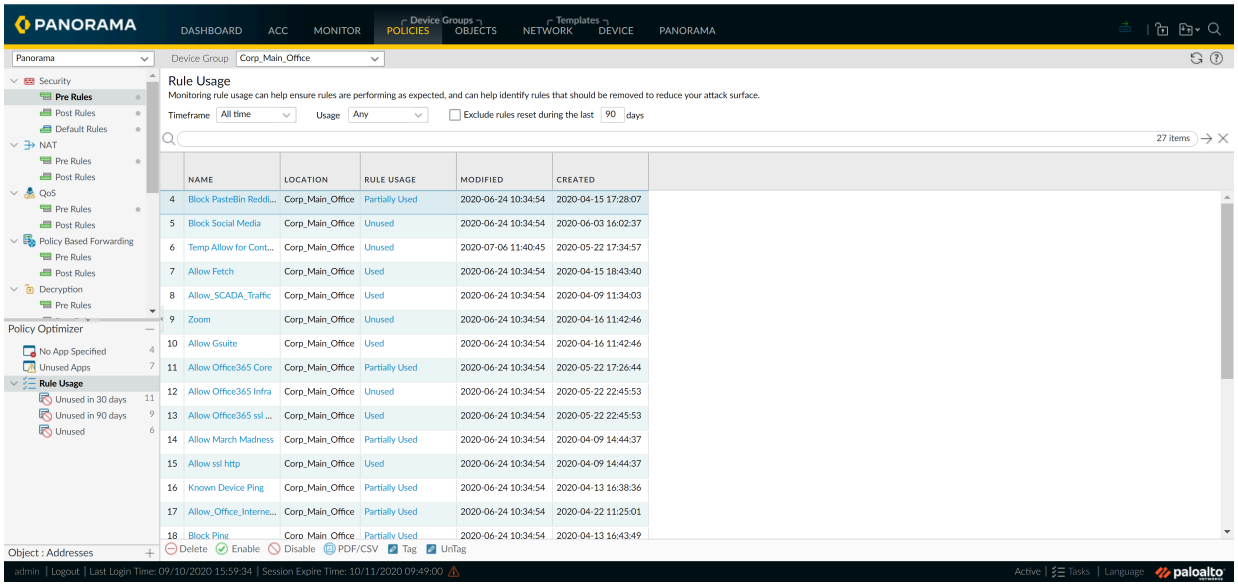
You can filter the rule usage for rules pushed to firewalls from Panorama. Panorama cannot filter rule usage for rules configured locally on the firewall.

Use the rule usage filter to evaluate the rule usage within a specified period of time. For example, filter the selected rulebase for Unused rules within the last 30 days. You can also evaluate rule usage with other rule attributes, such as the Created and Modified dates, which enables you to filter for the correct set of rules to review. You can use this data to help manage your rule lifecycle and to determine if a rule needs to be removed to reduce your network attack surface.

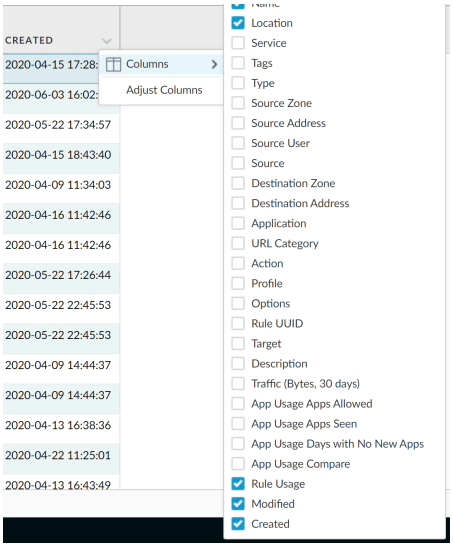
1. Select the **Timeframe** you want to filter on, or specify a **Custom** time frame.
2. Select the rule **Usage** on which you want to filter.
3. (Optional) If you have reset the rule usage data for any rules, check for **Exclude rules reset during the last <*number of days*> days** and decide when to exclude a rule based

 ©2025 Palo Alto Networks, Inc.

on the number of days you specify since the rule was reset. Only rules that were reset before your specified number of days are included in the filtered results.

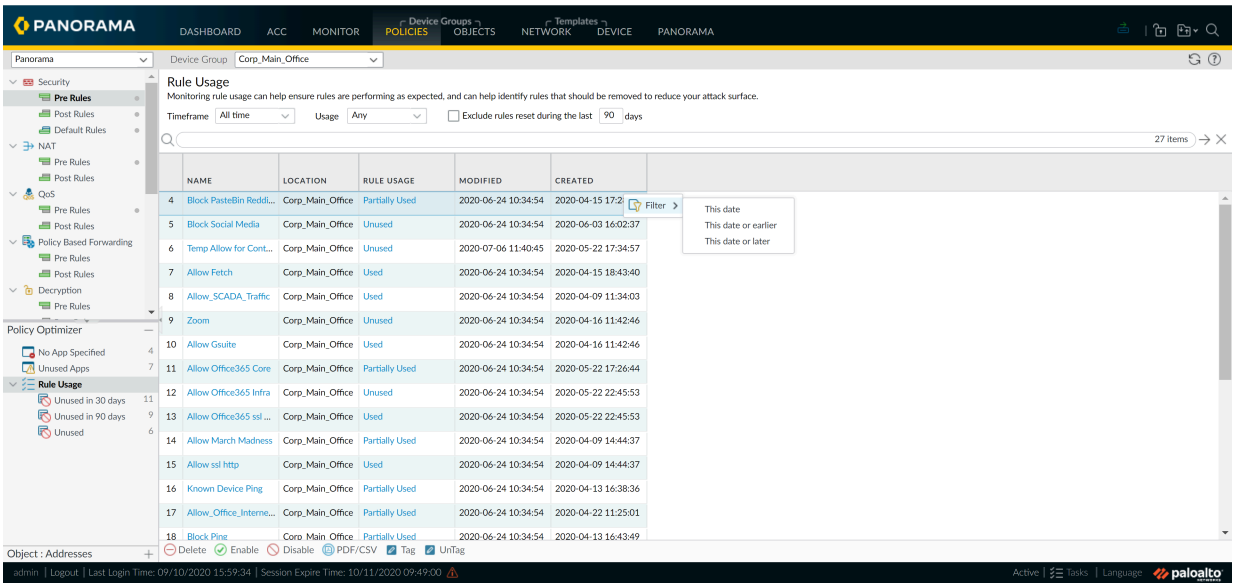


4. (Optional) Specify search filters based on additional rule data, other than the rule usage.
   1. Hover your mouse over the column header, and from the drop-down select **Columns**.
   2. Add any additional columns you want to filter with or to display.



   3. Hover your mouse over the column data that you would like to filter, and select **Filter** from the drop-down. For data that contain dates, select whether to filter using **This date**, **This date or earlier**, or **This date or later**.
   4. Click **Apply Filter** (→).

 ©2025 Palo Alto Networks, Inc.



STEP 10 | Take action on one or more unused policy rules.

1. Select one or more unused policy rules.
2. Perform one of the following actions:
   - **Delete**—Delete one or more selected policy rules.
   - **Enable**—Enable one or more selected policy rules when disabled.
   - **Disable**—Disable one or more selected policy rules.
   - **Tag**—Apply one or more group tags to one or more selected policy rules. The group tag must already exist in order to tag policy rule.
   - **Untag**—Remove one or more group tags from one or more selected policy rules.
3. Select **Commit** and **Commit and Push** your changes.

 ©2025 Palo Alto Networks, Inc.

Use Case: Configure Firewalls Using Panorama

Let's say that you want to use Panorama in a high availability configuration to manage a dozen firewalls on your network: you have six firewalls deployed across six branch offices, a pair of firewalls in a high availability configuration at each of two data centers, and a firewall in each of the two regional head offices.

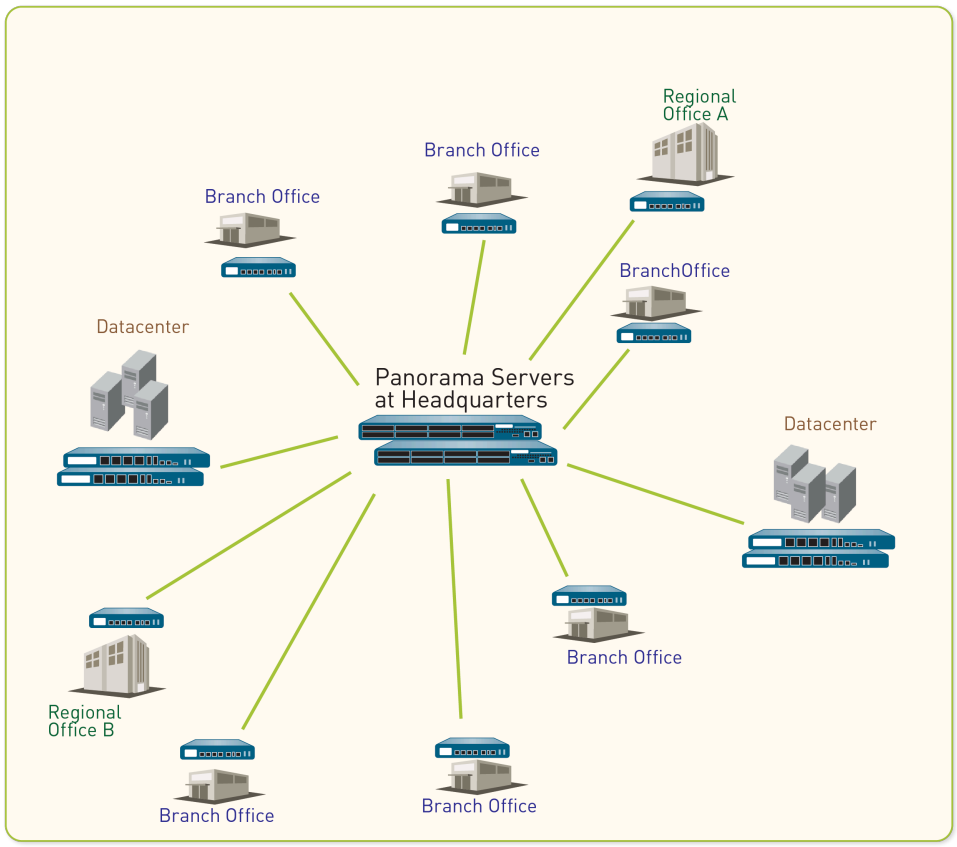


Figure 14: Firewall Distribution Example

The first step in creating your central management strategy is to determine how to group the firewalls into device groups and templates to efficiently push configurations from Panorama. You can base the grouping on the business functions, geographic locations, or administrative domains of the firewalls. In this example, you create two device groups and three templates to administer the firewalls using Panorama:

- Device Groups in this Use Case
- Templates in this Use Case
- Set Up Your Centralized Configuration and Policies

Device Groups in this Use Case

In Use Case: Configure Firewalls Using Panorama, we need to define two device groups based on the functions the firewalls will perform:

 ©2025 Palo Alto Networks, Inc.

- DG_BranchAndRegional for grouping firewalls that serve as the security gateways at the branch offices and at the regional head offices. We placed the branch office firewalls and the regional office firewalls in the same device group because firewalls with similar functions will require similar policy rulebases.
- DG_DataCenter for grouping the firewalls that secure the servers at the data centers.

We can then administer shared policy rules across both device groups as well as administer distinct device group rules for the regional office and branch office groups. Then for added flexibility, the local administrator at a regional or branch office can create local rules that match specific source, destination, and service flows for accessing applications and services that are required for that office. In this example, we create the following hierarchy for security rules. you can use a similar approach for any of the other rulebases.

Device Groups	DG_BranchAndRegional		DG_DataCenter
Rules	**Regional**	**Branch**	**Datacenter**
Shared pre-rule	Allow DNS and SNMP services.		
	Acceptable use policy that denies access to specified URL categories and peer-to-peer traffic that is of risk level 3, 4, and 5.		
Device Group pre-rule	Allow Facebook to all users in the marketing group in the regional offices only.		Allow access to the Amazon cloud application for the specified hosts/servers in the datacenter.
Local rules on a device	None		
Device Group post-rule	None		
Shared post-rule	To enable logging for all Internet-bound traffic on your network, create a rule that allows or denies all traffic from the trust zone to the untrust zone.		

Figure 15: Security Rules Hierarchy

Templates in this Use Case

When grouping firewalls for templates, we must take into account the differences in the networking configuration. For example, if the interface configuration is not the same—the interfaces are unlike in type, or the interfaces used are not alike in the numbering scheme and link capacity, or the zone to interface mappings are different—the firewalls must be in separate templates. Further, the way the firewalls are configured to access network resources might be different because the firewalls are spread geographically; for example, the DNS server, syslog servers and gateways that they access might be different. So, to allow for an optimal base configuration, in Use Case: Configure Firewalls Using Panorama we must place the firewalls in separate templates as follows:

- T_Branch for the branch office firewalls
- T_Regional for the regional office firewalls

 ©2025 Palo Alto Networks, Inc.

- T_DataCenter for the data center firewalls

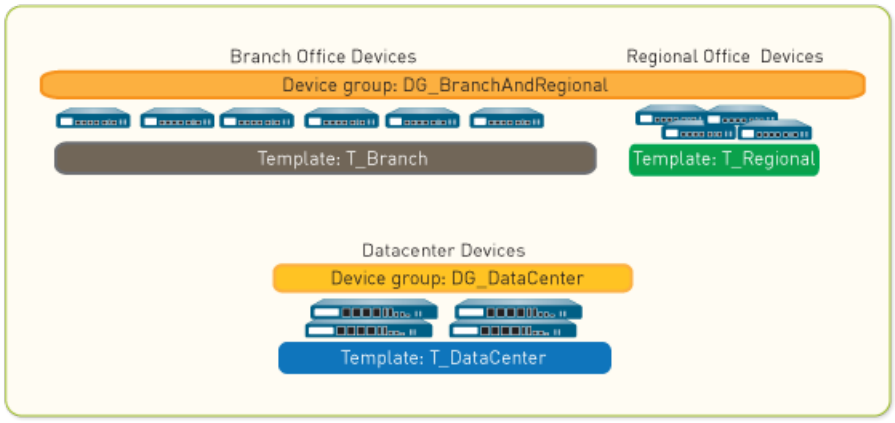


Figure 16: Device Group Example

If you plan to deploy your firewalls in an active/active HA configuration, assign each firewall in the HA pair to a separate template. Doing so gives you the flexibility to set up separate networking configurations for each peer. For example, you can manage the networking configurations in a separate template for each peer so that each can connect to different northbound and southbound routers, and can have different OSPF or BGP peering configurations.

Set Up Your Centralized Configuration and Policies

In Use Case: Configure Firewalls Using Panorama, we would need to perform the following tasks to centrally deploy and administer firewalls:

- Add the Managed Firewalls and Deploy Updates
- Use Templates to Administer a Base Configuration
- Use Device Groups to Push Policy Rules
- Preview the Rules and Commit Changes

Add the Managed Firewalls and Deploy Updates

The first task in Use Case: Configure Firewalls Using Panorama is to add the firewalls as managed devices and deploy content updates and PAN-OS software updates to those firewalls.

STEP 1 | For each firewall that Panorama will manage, Add a Firewall as a Managed Device.

In this example, add 12 firewalls.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Deploy the content updates to the firewalls. If you purchased a Threat Prevention subscription, the content and antivirus databases are available to you. First install the **Applications** or **Applications and Threats** database, then the **Antivirus**.

To review the status or progress for all tasks performed on Panorama, see Use the Panorama Task Manager.

1. Select **Panorama** > **Device Deployment** > **Dynamic Updates**.
2. Click **Check Now** to check for the latest updates. If the value in the Action column is **Download**, this indicates an update is available.
3. Click **Download**. When the download completes, the value in the Action column changes to **Install**.
4. In the **Action** column, click **Install**. Use the filters or user-defined tags to select the managed firewalls on which you would like to install this update.
5. Click **OK**, then monitor the status, progress, and result of the content update for each firewall. The **Result** column displays the success or failure of the installation.

STEP 3 | Deploy the software updates to the firewalls.

1. Select **Panorama** > **Device Deployment** > **Software**.
2. Click **Check Now** to check for the latest updates. If the value in the Action column is **Download**, this indicates an update is available.
3. Locate the version that you need for each hardware model and click **Download**. When the download completes, the value in the Action column changes to **Install**.
4. In the Action column, click the **Install** link. Use the filters or user-defined tags to select the managed firewalls on which to install this version.
5. Enable the check box for **Reboot device after install** or **Upload only to device (do not install)** and click **OK**. The **Results** column displays the success or failure of the installation.

Use Templates to Administer a Base Configuration

The second task in Use Case: Configure Firewalls Using Panorama is to create the templates you will need to push the base configuration to the firewalls.

STEP 1 | For each template you will use, Add a Template and assign the appropriate firewalls to each.

In this example, create templates named T_Branch, T_Regional, and T_DataCenter.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Define a DNS server, NTP server, syslog server, and login banner. Repeat this step for each template.

1. In the **Device** tab, select the **Template** from the drop-down.
2. Define the DNS and NTP servers:
   1. Select **Device** > **Setup** > **Services** > **Global** and edit the Services.
   2. In the **Services** tab, enter an IP address for the **Primary DNS Server**.

      *For any firewall that has more than one virtual system (vsys), for each vsys, add a DNS server profile to the template (**Device** > **Server Profiles** > **DNS**).*
   3. In the **NTP** tab, enter an IP address for the **Primary NTP Server**.
   4. Click **OK** to save your changes.
3. Add a login banner: select **Device** > **Setup** > **Management**, edit the General Settings, enter text for the **Login Banner** and click **OK**.
4. Configure a Syslog server profile (**Device** > **Server Profiles** > **Syslog**).

STEP 3 | Enable HTTPS, SSH, and SNMP access to the management interface of the managed firewalls. Repeat this step for each template.

1. In the **Device** tab, select the **Template** from the drop-down.
2. Select **Setup** > **Management**, and edit the Management Interface Settings.
3. Under Services, select the **HTTPS**, **SSH**, and **SNMP** check boxes, and click **OK**.

STEP 4 | Create a Zone Protection profile for the firewalls in the data center template (T_DataCenter).

1. Select the **Network** tab and, in the **Template** drop-down, select T_DataCenter.
2. Select **Network Profiles** > **Zone Protection** and click **Add**.
3. For this example, enable protection against a SYN flood—In the **Flood Protection** tab, select the **SYN** check box, set the **Action** to **SYN Cookies** as, set the **Alert** packets/second to **100**, set the **Activate** packets/second to **1000**, and set the **Maximum** packets/second to **10000**.
4. For this example, enable alerts—In the **Reconnaissance Protection** tab, select the **Enable** check boxes for **TCP Port Scan**, **Host Sweep**, and **UDP Port Scan**. Ensure the Action values are set to **alert** (the default value).
5. Click **OK** to save the Zone Protection profile.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Configure the interface and zone settings in the data center template (T_DataCenter), and then attach the Zone Protection profile you just created.

> *Before performing this step, you must have configured the interfaces locally on the firewalls. As a minimum, for each interface, you must have defined the interface type, assigned it to a virtual router (if needed), and attached a security zone.*

1. Select the **Network** tab and, in the **Template** drop-down, select T_DataCenter.
2. Select **Network** > **Interface** and, in the Interface column, click the interface name.
3. Select the **Interface Type** from the drop-down.
4. In the **Virtual Router** drop-down, click **New Virtual Router**. When defining the router, ensure the **Name** matches what is defined on the firewall.
5. In the **Security Zone** drop-down, click **New Zone**. When defining the zone, ensure that the **Name** matches what is defined on the firewall.
6. Click **OK** to save your changes to the interface.
7. Select **Network** > **Zones**, and select the zone you just created. Verify that the correct interface is attached to the zone.
8. In the **Zone Protection Profile** drop-down, select the profile you created, and click **OK**.

STEP 6 | Push your template changes.

1. Select **Commit** > **Commit and Push** and **Edit Selections** in the Push Scope.
2. Select **Templates** and select the firewalls assigned to the templates where you made changes.
3. **Commit and Push** your changes to the Panorama configuration and to the template.

Use Device Groups to Push Policy Rules

The third task in Use Case: Configure Firewalls Using Panorama is to create the device groups to manage policy rules on the firewalls.

STEP 1 | Create device groups and assign the appropriate firewalls to each device group: see Add a Device Group.

In this example, create device groups named DG_BranchAndRegional and DG_DataCenter.

When configuring the DG_BranchAndRegional device group, you must assign a **Master** firewall. This is the only firewall in the device group that gathers user and group mapping information for policy evaluation.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Create a shared pre-rule to allow DNS and SNMP services.

1. Create a shared application group for the DNS and SNMP services.
   1. Select **Objects** > **Application Group** and click **Add**.
   2. Enter a **Name** and select the **Shared** check box to create a shared application group object.
   3. Click **Add**, type **DNS**, and select **dns** from the list. Repeat for SNMP and select **snmp**, **snmp-trap**.
   4. Click **OK** to create the application group.
2. Create the shared rule.
   1. Select the **Policies** tab and, in the **Device Group** drop-down, select **Shared**.
   2. Select the **Security** > **Pre-Rules** rulebase.
   3. Click **Add** and enter a **Name** for the security rule.
   4. In the **Source** and **Destination** tabs for the rule, click **Add** and enter a **Source Zone** and a **Destination Zone** for the traffic.
   5. In the **Applications** tab, click **Add**, type the name of the applications group object you just created, and select it from the drop-down.
   6. In the **Actions** tab, set the **Action** to **Allow**, and click **OK**.

STEP 3 | Define the corporate acceptable use policy for all offices. In this example, create a shared rule that restricts access to some URL categories and denies access to peer-to-peer traffic that is of risk level 3, 4, or 5.

1. Select the **Policies** tab and, in the **Device Group** drop-down, select **Shared**.
2. Select **Security** > **Pre-Rules** and click **Add**.
3. In the **General** tab, enter a **Name** for the security rule.
4. In the **Source** and **Destination** tabs, click **Add** and select **any** for the traffic **Source Zone** and **Destination Zone**.
5. In the **Application** tab, define the application filter:
   1. Click **Add** and click **New Application Filter** in the footer of the drop-down.
   2. Enter a **Name**, and select the **Shared** check box.
   3. In the Risk column, select levels **3**, **4**, and **5**.
   4. In the Technology column, select **peer-to-peer**.
   5. Click **OK** to save the new filter.
6. In the **Service/URL Category** tab, URL Category section, click **Add** and select the categories you want to block (for example, **streaming-media**, **dating**, and **online-personal-storage**).
7. You can also attach the default URL Filtering profile—In the **Actions** tab, Profile Setting section, select the **Profile Type** option **Profiles**, and select the **URL Filtering** option **default**.
8. Click **OK** to save the security pre-rule.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Allow Facebook for all users in the Marketing group in the regional offices only.

Enabling a security rule based on user and group has the following prerequisite tasks:

- Set up User-ID on the firewalls.
- Enable User-ID for each zone that contains the users you want to identify.
- Define a master firewall for the DG_BranchAndRegional device group (see step 1).

1. Select the **Policies** tab and, in the **Device Group** drop-down, select DG_BranchAndRegional.
2. Select the **Security** > **Pre-Rules** rulebase.
3. Click **Add** and enter a **Name** for the security rule.
4. In the **Source** tab, **Add** the Source Zone that contains the Marketing group users.
5. In the **Destination** tab, **Add** the Destination Zone.
6. In the **User** tab, **Add** the Marketing user group to the Source User list.
7. In the **Application** tab, click **Add**, type **Facebook**, and then select it from the drop-down.
8. In the **Action** tab, set the **Action** to **Allow**.
9. In the **Target** tab, select the regional office firewalls and click **OK**.

STEP 5 | Allow access to the Amazon cloud application for the specified hosts/servers in the data center.

1. Create an address object for the servers/hosts in the data center that need access to the Amazon cloud application.
    1. Select **Objects** > **Addresses** and, in the **Device Group** drop-down, select DG_DataCenter.
    2. Click **Add** and enter a **Name** for the address object.
    3. Select the **Type**, and specify an IP address and netmask (**IP Netmask**), range of IP addresses (**IP Range**), or **FQDN**.
    4. Click **OK** to save the object.
2. Create a security rule that allows access to the Amazon cloud application.
    1. Select **Policies** > **Security** > **Pre-Rules** and, in the **Device Group** drop-down, select DG_DataCenter.
    2. Click **Add** and enter a **Name** for the security rule.
    3. Select the **Source** tab, **Add** the Source Zone for the data center, and **Add** the address object (Source Address) you just defined.
    4. Select the **Destination** tab and **Add** the Destination Zone.
    5. Select the **Application** tab, click **Add**, type **amazon**, and select the Amazon applications from the list.
    6. Select the **Action** tab and set the **Action** to **Allow**.
    7. Click **OK** to save the rule.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | To enable logging for all internet-bound traffic on your network, create a rule that matches trust zone to untrust zone.

1. Select the **Policies** tab and, in the **Device Group** drop-down, select **Shared**.
2. Select the **Security** > **Pre-Rules** rulebase.
3. Click **Add** and enter a **Name** for the security rule.
4. In the **Source** and **Destination** tabs for the rule, **Add** **`trust_zone`** as the Source Zone and **`untrust_zone`** as the Destination Zone.
5. In the **Action** tab, set the **Action** to **Deny**, set the **Log Setting** to **Log at Session end**, and click **OK**.

Preview the Rules and Commit Changes

The final task in Use Case: Configure Firewalls Using Panorama is to review the rules and commit the changes you have made to Panorama, device groups, and templates.

STEP 1 | Preview the rules.

This preview enables you to visually evaluate how rules are layered for a particular rulebase.

1. Select **Policies** and **Preview Rules**.
2. Select a **Rulebase**, **Device Group**, and **Device**.
3. Close the preview dialog when you finish.

STEP 2 | Commit and push your configuration changes.

1. Select **Commit** > **Commit and Push** and **Edit Selections** in the Push Scope.
2. Select **Device Groups**, select the device groups you added, and **Include Device and Network Templates**.
3. (Optional) Disable **Merge with Device Candidate Config** if you manage local firewall configuration changes independently of configuration changes from Panorama.

   This setting is enabled by default and merges any pending local firewall configurations with the configuration push from Panorama. The local firewall configuration is merged and committed regardless of the admin pushing the changes from Panorama or the admin who made the local firewall configuration changes.
4. Click **OK** to save your changes to the Push Scope.
5. **Commit and Push** your changes.

STEP 3 | Verify that Panorama applied the template and policy configurations.

1. In the Panorama header, set the **Context** to the firewall to access its web interface.
2. Review the template and policy configurations to ensure your changes are there.

 ©2025 Palo Alto Networks, Inc.

 ©2025 Palo Alto Networks, Inc.

Manage Log Collection

All Palo Alto Networks firewalls can generate logs that provide an audit trail of firewall activities. For Centralized Logging and Reporting, you must forward the logs generated on the firewalls to your on-premise infrastructure that includes the Panorama™ management server or Log Collectors or send the logs to the cloud-based Strata Logging Service. Optionally, you can then configure Panorama to forward the logs to external logging destinations (such as syslog servers).

If you forward logs to a Panorama virtual appliance in Legacy mode, you don't need to perform any additional tasks to enable logging. If you forward logs to Log Collectors, you must configure them as managed collectors and assign them to Collector Groups. A managed collector can be local to an M-Series appliance, or Panorama virtual appliance in Panorama mode. Additionally, an M-Series appliance, or Panorama virtual appliance in Log Collector mode can be Dedicated Log Collectors. To determine whether to deploy either or both types of managed collectors, see Local and Distributed Log Collection.

To manage the System and Config logs that Panorama generates locally, see Monitor Panorama.

- Configure a Managed Collector
- Monitor Managed Collector Health Status
- Configure Authentication for a Dedicated Log Collector
- Manage Collector Groups
- Configure Log Forwarding to Panorama
- Configure Syslog Forwarding to External Destinations
- Forward Logs to Strata Logging Service
- Verify Log Forwarding to Panorama
- Modify Log Forwarding and Buffering Defaults
- Configure Log Forwarding from Panorama to External Destinations
- Log Collection Deployments

Configure a Managed Collector

To enable the Panorama management server to manage a Log Collector, you must add it as a managed collector. Log Collectors support communication using a public or private IPv4 or IPv6 address only, including when you configure custom certificates for mutual authentication.

You can add two types of managed collectors:

- **Dedicated Log Collector**—To set up a new M-700, M-600, M-500, M-300, or M-200 appliance or a Panorama virtual appliance as a Log Collector to switch an existing M-Series appliance or Panorama virtual appliance from Panorama mode to Log Collector mode, you must Set Up the M-Series Appliance as a Log Collector. Keep in mind that switching from Panorama Mode to Log Collector Mode removes the local Log Collector that is predefined on the M-Series appliance in Panorama mode.
- **Local Log Collector**—A Log Collector can run locally on a M-700, M-600, M-500, M-300, or M-200 appliance or a Panorama virtual appliance in Panorama mode. On the M-Series appliances, the Log Collector is predefined; on the virtual appliance, you must add the Log Collector. When the Panorama management server has a high availability (HA) configuration, each HA peer can have a local Log Collector. However, relative to the primary Panorama, the Log Collector on the secondary Panorama is remote, not local. Therefore, to use the Log Collector on the secondary Panorama, you must manually add it to the primary Panorama (for details, see Deploy Panorama M-Series Appliances with Local Log Collectors or Deploy Panorama Virtual Appliances with Local Log Collectors). If you delete a local Log Collector, you can later add it back. The following steps describe how to add a local Log Collector.

If the Panorama virtual appliance is in Legacy mode, you must switch to Panorama mode to create a Log Collector. For details, see Set Up the Panorama Virtual Appliance with Local Log Collector.

A device registration authentication key is used to securely authenticate and connect the Panorama management server and the managed collector on first connect. To configure the device registration authentication key, specify the key lifetime and the number of times you can use the authentication key to onboard new Log Collectors. Additionally, you can specify one or more Log Collector serial numbers for which the authentication key is valid.

The authentication key expires 90 days after the key lifetime expires. After 90 days, you are prompted to re-certify the authentication key to maintain its validity. If you do not re-certify, then the authentication key becomes invalid. A system log is generated each time a Log Collector uses the Panorama-generated authentication key. The Log Collector uses the authentication key to authenticate Panorama when it delivers the device certificate that is used for all subsequent communications.

As a best practice, retain a local Log Collector and Collector Group on the Panorama management server, regardless whether it manages Dedicated Log Collectors.

(Panorama evaluation only) If you are evaluating a Panorama virtual appliance with a local Log Collector, Configure Log Forwarding from Panorama to External Destinations to preserve logs generated during your evaluation period.

Logs stored on the local Log Collector cannot be preserved when you Convert Your Evaluation Panorama Instance to a Production Panorama Instance with a Local Log Collector.

 ©2025 Palo Alto Networks, Inc.

For Dedicated Log Collectors running a PAN-OS 10.1 release, Panorama running PAN-OS 10.2 or later release supports onboarding Dedicated Log Collectors running PAN-OS 10.1.3 or later release only. You cannot add a Dedicated Log Collector running PAN-OS 10.1.2 or earlier PAN-OS 10.1 release to Panorama management if Panorama is running PAN-OS 10.2 or later release.

Panorama supports onboarding Dedicated Log Collectors running the following releases:

- ***Panorama running PAN-OS 10.2 or later release—** Dedicated Log Collectors running PAN-OS 10.1.3 or later release, and Dedicated Log Collectors running PAN-OS 10.0 or earlier PAN-OS release.*

There is no impact to Dedicated Log Collectors already managed by Panorama on upgrade to PAN-OS 10.2.

If you are experiencing issues adding a Dedicated Log Collector to Panorama management, you may need to recover managed device connectivity to Panorama.

STEP 1 | Record the serial number of the Log Collector.

You will need the serial number when you add the Log Collector as a managed collector.

1. Access the Panorama web interface.
2. Select **Dashboard** and record the **Serial #** in the General Information section.

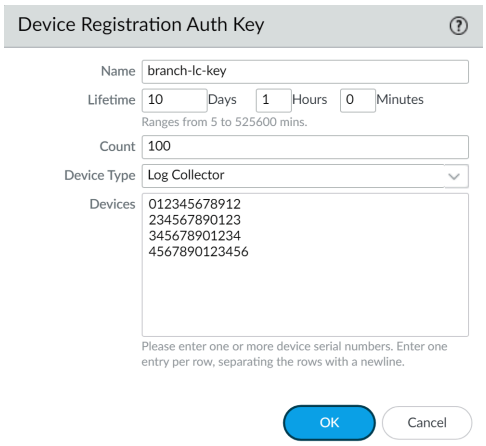
STEP 2 | Log in to the Panorama Web Interface.

 ©2025 Palo Alto Networks, Inc.

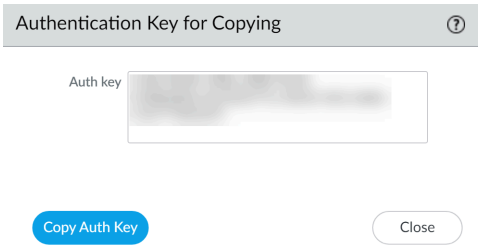
STEP 3 | Create a device registration authentication key.

1. Select **Panorama** > **Device Registration Auth Key** and **Add** a new authentication key.
2. Configure the authentication key.
   - **Name**—Add a descriptive name for the authentication key.
   - **Lifetime**—Specify the key lifetime for how long you can use the authentication key to onboard new Log Collectors.
   - **Count**—Specify how many times you can use the authentication key to onboard new Log Collectors.
   - **Device Type**—Specify that this authentication key is used to authenticate only a **Log Collector**.

     *You can select **Any** to use the device registration authentication key to onboard firewalls, Log Collectors, and WildFire appliances.*
   - (Optional) **Devices**—Enter one or more device serial numbers to specify for which Log Collectors the authentication key is valid.
3. Click **OK**.



4. **Copy Auth Key** and **Close**.



 ©2025 Palo Alto Networks, Inc.

STEP 4 | (Dedicated Log Collector only) Add the device registration authentication key to the Log Collector.

Add the device registration authentication key only to a Dedicated Log Collector. A Panorama in Panorama mode does not need to authenticate its own local Log Collector.

1. Log in to the Log Collector CLI.
2. Add the device registration authentication key.

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



STEP 5 | Add the Log Collector as a managed collector.

1. In the Panorama web interface, select **Panorama** > **Managed Collectors** and **Add** a new Log Collector.
2. In the **General** settings, enter the serial number (**Collector S/N**) you recorded for the Log Collector.
3. Click **OK** to save your changes.
4. Select **Commit** > **Commit to Panorama**.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | (Optional) Configure the Log Collector admin authentication.

*Palo Alto Networks recommends adding at least one Local Administrator with Superuser privileges for **Authentication** if you configure an authorization list for your managed collector.*

If you added any Imported Panorama Admin Users, then you are required to add at least one Local Administrator with Superuser privileges.

1. Select **Panorama** > **Managed Collectors** and edit the Log Collector by clicking its name.
2. Configure the Log Collector admin password:
   1. Select the password **Mode**.
   2. If you selected **Password** mode, enter a plaintext **Password** and **Confirm Password**. If you selected **Password Hash** mode, enter a hashed password string of up to 63 characters.
3. Configure the admin login security requirements:

   *If you set the **Failed Attempts** to a value other than 0 but leave the **Lockout Time** at 0, then the admin user is indefinitely locked out until another administrator manually unlocks the locked out admin. If no other administrator has been created, you must reconfigure the **Failed Attempts** and **Lockout Time** settings on Panorama and push the configuration change to the Log Collector. To ensure that an admin is never locked out, use the default 0 value for both **Failed Attempts** and **Lockout Time**.*

   1. Enter the number of login **Failed Attempts** value. The range is between the default value **0** to the maximum of **10** where the value **0** specifies unlimited login attempts.
   2. Enter the **Lockout Time** value between the default value **0** to the maximum of **60** minutes.
4. Click **OK** to save your changes.

STEP 7 | Enable the logging disks.

1. Select **Panorama** > **Managed Collectors** and edit the Log Collector by clicking its name.

   The Log Collector name has the same value as the hostname of the Panorama management server.
2. Select **Disks** and **Add** each disk pair.
3. Click **OK** to save your changes.
4. Select **Commit** > **Commit to Panorama**.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | (Optional) If your deployment is using custom certificates for authentication between Panorama and managed devices, deploy the custom client device certificate. For more information, see Set Up Authentication Using Custom Certificates.

1. Select **Panorama** > **Certificate Management** > **Certificate Profile** and choose the certificate profile from the drop-down or click **New Certificate Profile** to create one.
2. Select **Panorama** > **Managed Collectors** and **Add** a new Log Collector or select an existing one. Select **Communication**.
3. Select the type of device certificate the Type drop-down.
   - If you are using a local device certificate, select the **Certificate** and **Certificate Profile** from the respective drop-downs.
   - If you are using SCEP as the device certificate, select the **SCEP Profile** and **Certificate Profile** from the respective drop-downs.
4. Click **OK**.

 ©2025 Palo Alto Networks, Inc.

STEP 9 | (Optional) Configure **Secure Server Communication** on a Log Collector. For more information, see Set Up Authentication Using Custom Certificates.

1. Select **Panorama** > **Managed Collectors** and click **Add**. Select **Communication**.
2. Verify that the **Custom Certificate Only** check box is not selected. This allows you to continue managing all devices while migrating to custom certificates.

   When the Custom Certificate Only check box is selected, the Log Collector does not authenticate and cannot receive logs from devices using predefined certificates.
3. Select the SSL/TLS service profile from the **SSL/TLS Service Profile** drop-down. This SSL/TLS service profile applies to all SSL connections between the Log Collector and devices sending it logs.
4. Select the certificate profile from the **Certificate Profile** drop-down.
5. Select **Authorize Client Based on Serial Number** to have the server check clients against the serial numbers of managed devices. The client certificate must have the special keyword $UDID set as the CN to authorize based on serial numbers.
6. In **Disconnect Wait Time (min)**, enter the number of minutes Panorama should before breaking and reestablishing the connection with its managed devices. This field is blank by default and the range is 0 to 44,640 minutes.

   The disconnect wait time does not begin counting down until you commit the new configuration.
7. (Optional) Configure an authorization list.
   1. **Add** an Authorization List.
   2. Select the **Subject** or **Subject Alt Name** as the Identifier type.
   3. Specify an identifier of the selected type.
   4. Click **OK**.
   5. Enable the Log Collector to **Check Authorization List** to enforce the authorization list.
8. Click **OK**.
9. Select **Commit** > **Commit to Panorama**.

STEP 10 | Verify your changes.

1. Verify that the **Panorama** > **Managed Collectors** page lists the Log Collector you added. The Connected column displays a check mark to indicate that the Log Collector is connected to Panorama. You might have to wait a few minutes before the page displays the updated connection status.

   Until you Configure a Collector Group and push configuration changes to the Collector Group, the Configuration Status column displays Out of Sync, the Run Time Status column displays disconnected, and the CLI command ***`show interface all`*** *displays the interfaces as down.*
2. Click **Statistics** in the last column to verify that the logging disks are enabled.

 ©2025 Palo Alto Networks, Inc.

STEP 11 | Next steps...

Before a Log Collector can receive firewall logs, you must:

1. Configure Log Forwarding to Panorama.
2. Configure a Collector Group—On the M-Series appliances, a default Collector Group is predefined and already contains the local Log Collector as a member. On the Panorama virtual appliance, you must add the Collector Group and add the local Log Collector as a member. On both models, assign firewalls to the local Log Collector for log forwarding.

   You must add the Log Collector to a Collector Group before it can start ingesting firewall logs. The ElasticSearch health status displays as degraded and the Log Collector cannot ingest logs until added to a Collector Group.
3. Monitor Managed Collector Health Status to identify and resolve issues impacting log collection should they arise.

 ©2025 Palo Alto Networks, Inc.

Monitor Managed Collector Health Status

Monitor the health status of your managed Log Collector to identify and resolve issues impacting log collection. The Log Collector health status is based on the health status of vital Log Collector processes and you can view both the overall health status and the health status of each log collection process.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Configure a Managed Collector.

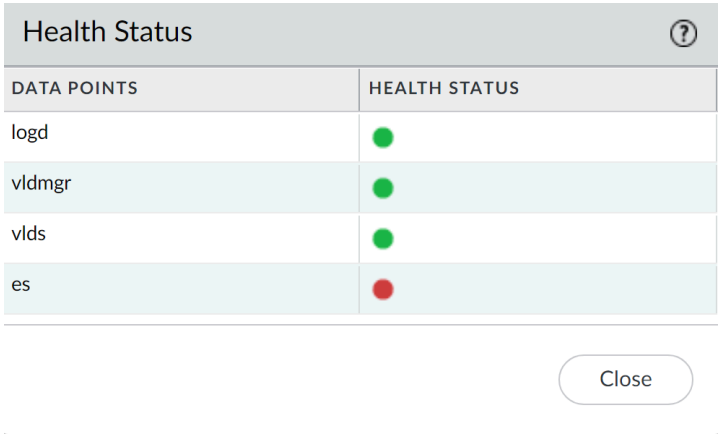
STEP 3 | Select **Panorama** > **Managed Collectors** and navigate to the Health column.

STEP 4 | Review the overall health status of the Log Collector.

A green circle (●) indicates that the Log Collector is healthy and red circle (●) indicates that one or more log collection processes are experiencing degraded health.

STEP 5 | View the **Health Status** details to view the health status of each log collection process.

- **logd—** Process responsible for ingesting logs received from the managed firewall and for transferring ingested logs to the vldmgr.
- **vldmgr—**Process responsible for managing the vld processes.
- **vlds—**Process responsible for managing individual logging disks, writing logs to the logging disks, and ingesting logs into ElasticSearch.
- **es—**ElasticSearch process running on the Log Collector.



 ©2025 Palo Alto Networks, Inc.

Configure Authentication for a Dedicated Log Collector

Create and configure enhanced authentication for your Dedicated Log Collector by configuring local administrative users with granular authentication parameters, as well as leveraging RADIUS, TACAS+, or LDAP for authorization and authentication.

When you Configure and push administrators from Panorama, you overwrite the existing administrators on the Dedicated Log Collectors with those you configure on Panorama.

Only Superuser administrators are supported when configuring an administrative account for a Dedicated Log Collector. Local or Panorama Administrators with any other admin role type is not supported.

(RADIUS and TACAS+) Only Superuser administrators are supported when configuring an administrative account for a Dedicated Log Collector. Remote, Local, or Panorama Administrators with any other admin role type is not supported.

- Configure an Administrative Account for a Dedicated Log Collector
- Configure RADIUS Authentication for a Dedicated Log Collector
- Configure TACACS+ Authentication for a Dedicated Log Collector
- Configure LDAP Authentication for a Dedicated Log Collector

Configure an Administrative Account for a Dedicated Log Collector

Create one or more administrators with granular authentication parameters for your Dedicated Log Collector to manage from the Panorama™ management server. Additionally, you can configure local administrators from Panorama that can be configured directly on the CLI of the Dedicated Log Collector. However, pushing new configuration changes to a Dedicated Log Collector overwrites existing local administrators with the administrators configured for the Dedicated Log Collector.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Configure a Managed Collector.

STEP 3 | (Optional) Configure an authentication profile to define the authentication service that validates the login credentials of the administrators who access the Dedicated Log Collector CLI.

STEP 4 | Configure one or more administrator accounts as needed.

The administrator accounts created on Panorama are later imported to the Dedicated Log Collector and managed from Panorama.

Only Superuser administrators are supported when configuring an administrative account for a Dedicated Log Collector. Local or Panorama Administrators with any other admin role type is not supported.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Configure the authentication for the Dedicated Log Collector.

1. Select **Panorama** > **Managed Collectors** and select the Dedicated Log Collector you previously added.
2. (Optional) Select the **Authentication Profile** you configured in the previous step.
3. Configure the authentication **Timeout Configuration** for the Dedicated Log Collector.
   1. Enter the number of **Failed Attempt**s before a user is locked out of the Dedicated Log Collector CLI.
   2. Enter the **Lockout Time**, in minutes, for which the Dedicated Log Collector locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access the Dedicated Log Collector.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
4. Add the Dedicated Log Collector administrators.

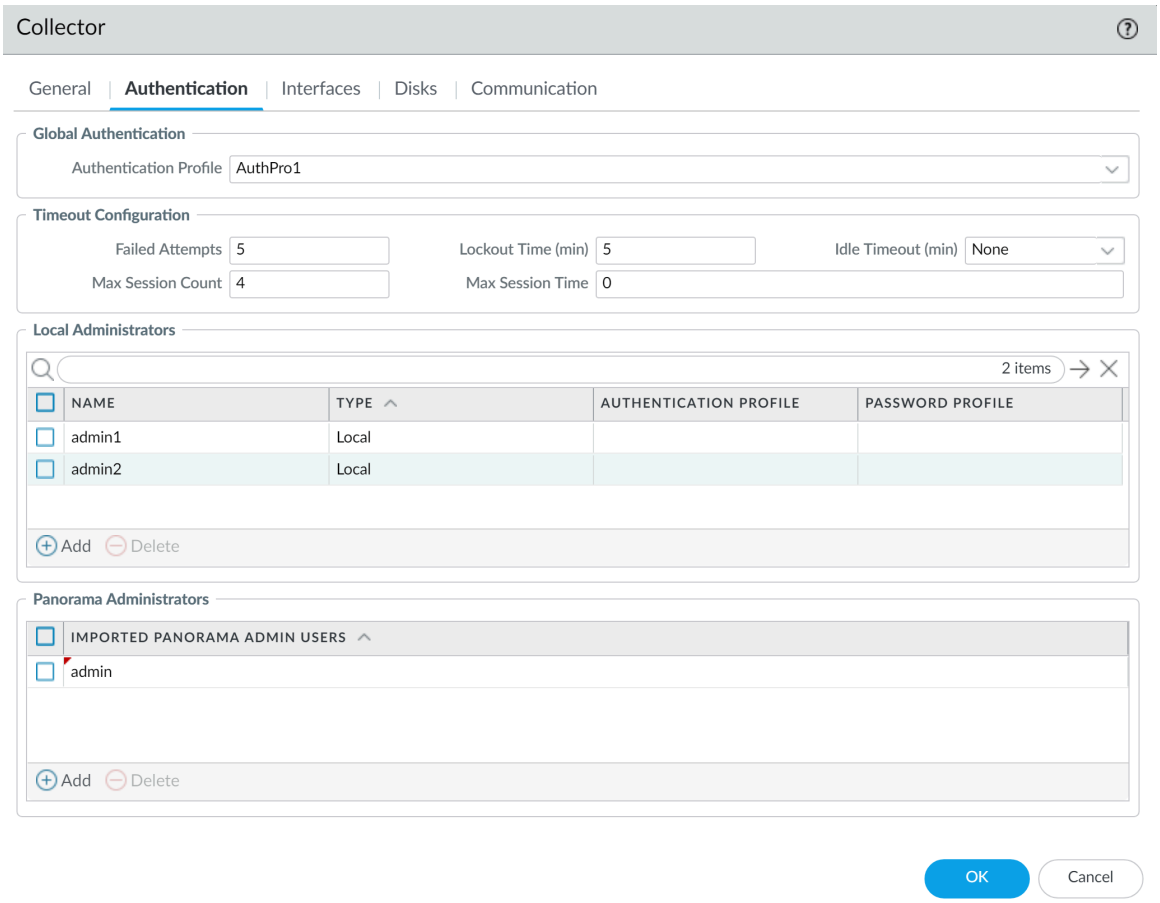
   Administrators may either be added as a local administrator or as an imported Panorama administrator—but not both. Adding the same administrator as both a local administrator and as an imported Panorama administrator is not supported and causes the Panorama

 ©2025 Palo Alto Networks, Inc.

commit to fail. For example, the commit to Panorama fails if you add **admin1** as both a local and Panorama administrator.

1. **Add** and configure new administrators unique to the Dedicated Log Collector. These administrators are specific to the Dedicated Log Collector for which they are created and you manage these administrators from this table.
2. **Add** any administrators configured on Panorama. These administrators are created on Panorama and imported to the Dedicated Log Collector.

5. Click **OK** to save the Dedicated Log Collector authentication configuration.



STEP 6 | **Commit** and then **Commit and Push** your configuration changes.

STEP 7 | Log in to the Panorama CLI of the Dedicated Log Collector to verify you can successfully access the Dedicated Log Collector using the local admin user.

Configure RADIUS Authentication for a Dedicated Log Collector

Use a RADIUS server to authenticate administrative access to the Dedicated Log Collector CLI. You can also define Vendor-Specific Attributes (VSAs) on the RADIUS server to manage administrator authorization. Using VSAs enables you to quickly change the roles, access domains, and user groups of administrators through your directory service, which is often easier than reconfiguring settings on the Panorama™ management server.

 ©2025 Palo Alto Networks, Inc.

You can Import the Palo Alto Networks RADIUS dictionary into RADIUS server to define the authentication attributes needed for communication between Panorama and the RADIUS server.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Configure a Managed Collector.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Configure RADIUS authentication.

Only Superuser administrators are supported when configuring an administrative account for a Dedicated Log Collector. Remote, Local, or Panorama Administrators with any other admin role type is not supported.

1. Add a RADIUS server profile.

   The profile defines how the Dedicated Log Collector connects to the RADIUS server.

   1. Select **Panorama** > **Server Profiles** > **RADIUS** and **Add** a profile.
   2. Enter a **Profile Name** to identify the server profile.
   3. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–20).
   4. Select the **Authentication Protocol** (default is **CHAP**) that the Dedicated Log Collector uses to authenticate to the RADIUS server.

      *Select **CHAP** if the RADIUS server supports that protocol; it is more secure than **PAP**.*

   5. **Add** each RADIUS server and enter the following:
      1. **Name** to identify the server.
      2. **RADIUS Server** IP address or FQDN.
      3. **Secret**/**Confirm Secret** (a key to encrypt usernames and passwords).
      4. Server **Port** for authentication requests (default is 1812).
   6. Click **OK** to save the server profile.
2. Assign the RADIUS server profile to an authentication profile.

   The authentication profile defines authentication settings that are common to a set of administrators.

   1. Select **Panorama** > **Authentication Profile** and **Add** a profile.
   2. Enter a **Name** to identify the authentication profile.
   3. Set the **Type** to **RADIUS**.
   4. Select the **Server Profile** you configured.
   5. Select **Retrieve user group from RADIUS** to collect user group information from VSAs defined on the RADIUS server.

      Panorama matches the group information against the groups you specify in the Allow List of the authentication profile.
   6. Select **Advanced** and, in the Allow List, **Add** the administrators that are allowed to authenticate with this authentication profile.
   7. Click **OK** to save the authentication profile.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the authentication for the Dedicated Log Collector.

1. Select **Panorama** > **Managed Collectors** and select the Dedicated Log Collector you previously added.
2. Select the **Authentication Profile** you configured in the previous step.

   If a global authentication profile is not assigned you must assign an authentication profile to each individual local administrator to leverage remote authentication.
3. Configure the authentication **Timeout Configuration** for the Dedicated Log Collector.
   1. Enter the number of **Failed Attempt**s before a user is locked out of the Dedicated Log Collector CLI.
   2. Enter the **Lockout Time**, in minutes, for which the Dedicated Log Collector locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access the Dedicated Log Collector.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
4. Add the Dedicated Log Collector administrators.

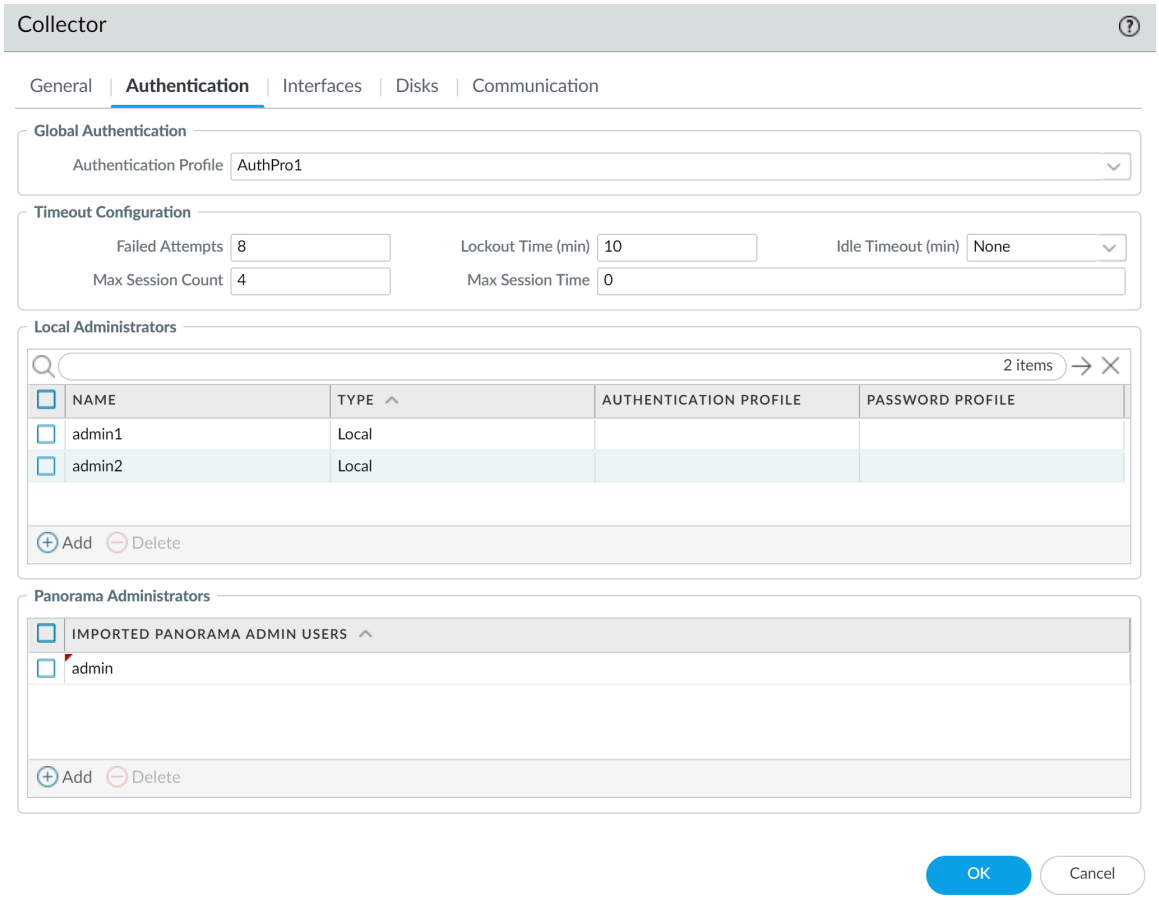
   Administrators may either be added as a local administrator or as an imported Panorama administrator—but not both. Adding the same administrator as both a local administrator and as an imported Panorama administrator is not supported and causes the Panorama

 ©2025 Palo Alto Networks, Inc.

commit to fail. For example, the commit to Panorama fails if you add **admin1** as both a local and Panorama administrator.

1. **Add** and configure new administrators unique to the Dedicated Log Collector. These administrators are specific to the Dedicated Log Collector for which they are created and you manage these administrators from this table.
2. **Add** any administrators configured on Panorama. These administrators are created on Panorama and imported to the Dedicated Log Collector.

5. Click **OK** to save the Dedicated Log Collector authentication configuration.



STEP 5 | **Commit** and then **Commit and Push** your configuration changes.

STEP 6 | Log in to the Panorama CLI of the Dedicated Log Collector to verify you can successfully access the Dedicated Log Collector using the local admin user.

Configure TACACS+ Authentication for a Dedicated Log Collector

You can use a TACACS+ server to authenticate administrative access to the Dedicated Log Collector CLI. You can also define Vendor-Specific Attributes (VSAs) on the TACACS+ server to manage administrator authorization. Using VSAs enables you to quickly change the roles, access domains, and user groups of administrators through your directory service, which is often easier than reconfiguring settings on Panorama.

STEP 1 | Log in to the Panorama Web Interface.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Configure a Managed Collector.

STEP 3 | Configure TACACS+ authentication.

> *Only Superuser administrators are supported when configuring an administrative account for a Dedicated Log Collector. Remote, Local, or Panorama Administrators with any other admin role type is not supported.*

1. Add a TACACS+ server profile.

   The profile defines how the Dedicated Log Collector connects to the TACACS+ server.

   1. Select **Panorama** > **Server Profiles** > **TACACS+** and **Add** a profile.
   2. Enter a **Profile Name** to identify the server profile.
   3. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–20).
   4. Select the **Authentication Protocol** (default is **CHAP**) that Panorama uses to authenticate to the TACACS+ server.
   5. Select **CHAP** if the TACACS+ server supports that protocol; it is more secure than **PAP**.
   6. **Add** each TACACS+ server and enter the following:
      1. **Name** to identify the server.
      2. **TACACS+ Server** IP address or FQDN.
      3. **Secret/Confirm Secret** (a key to encrypt usernames and passwords).
      4. Server **Port** for authentication requests (default is 49).
   7. Click **OK** to save the server profile.
2. Assign the TACACS+ server profile to an authentication profile.

   The authentication profile defines authentication settings that are common to a set of administrators.

   1. Select **Panorama** > **Authentication Profile** and **Add** a profile.
   2. Enter a **Name** to identify the profile.
   3. Set the **Type** to **TACACS+**.
   4. Select the **Server Profile** you configured.
   5. Select **Retrieve user group from TACACS+** to collect user group information from VSAs defined on the TACACS+ server.

      Panorama matches the group information against the groups you specify in the Allow List of the authentication profile.
   6. Select **Advanced** and, in the Allow List, **Add** the administrators that are allowed to authenticate with this authentication profile.
   7. Click **OK** to save the authentication profile.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the authentication for the Dedicated Log Collector.

1. Select **Panorama** > **Managed Collectors** and select the Dedicated Log Collector you previously added.
2. Select the **Authentication Profile** you configured in the previous step.

   If a global authentication profile is not assigned you must assign an authentication profile to each individual local administrator to leverage remote authentication.
3. Configure the authentication **Timeout Configuration** for the Dedicated Log Collector.
   1. Enter the number of **Failed Attempt**s before a user is locked out of the Dedicated Log Collector CLI.
   2. Enter the **Lockout Time**, in minutes, for which the Dedicated Log Collector locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access the Dedicated Log Collector.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
4. Add the Dedicated Log Collector administrators.

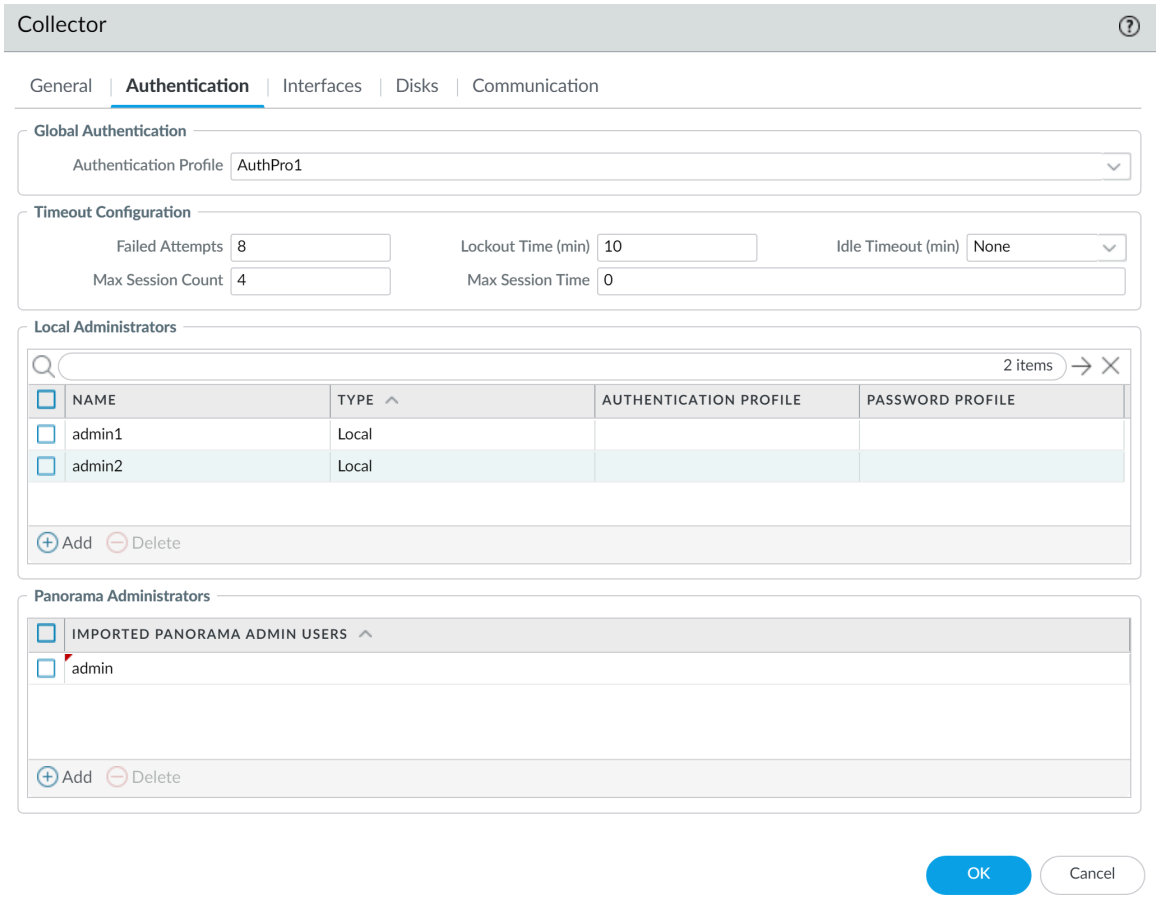
   Administrators may either be added as a local administrator or as an imported Panorama administrator—but not both. Adding the same administrator as both a local administrator and as an imported Panorama administrator is not supported and causes the Panorama

 ©2025 Palo Alto Networks, Inc.

commit to fail. For example, the commit to Panorama fails if you add **admin1** as both a local and Panorama administrator.

1. **Add** and configure new administrators unique to the Dedicated Log Collector. These administrators are specific to the Dedicated Log Collector for which they are created and you manage these administrators from this table.
2. **Add** any administrators configured on Panorama. These administrators are created on Panorama and imported to the Dedicated Log Collector.

5. Click **OK** to save the Dedicated Log Collector authentication configuration.



STEP 5 | **Commit** and then **Commit and Push** your configuration changes.

STEP 6 | Log in to the Panorama CLI of the Dedicated Log Collector to verify you can successfully access the Dedicated Log Collector using the local admin user.

Configure LDAP Authentication for a Dedicated Log Collector

You can use LDAP to authenticate end users who access Dedicated Log Collector web interface.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Configure a Managed Collector.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Add an LDAP server profile.

The profile defines how the Dedicated Log Collector connects to the LDAP server.

Only Superuser administrators are supported when configuring an administrative account for a Dedicated Log Collector. Local or Panorama Administrators with any other admin role type is not supported.

1. Select **Panorama** > **Server Profiles** > **LDAP** and **Add** a server profile.
2. Enter a **Profile Name** to identify the server profile.
3. **Add** the LDAP servers (up to four). For each server, enter a **Name** (to identify the server), **LDAP Server** IP address or FQDN, and server **Port** (default 389).

   If you use an FQDN address object to identify the server and you subsequently change the address, you must commit the change for the new server address to take effect.
4. Select the server **Type**.
5. Select the **Base DN**.

   To identify the Base DN of your directory, open the **Active Directory Domains and Trusts** Microsoft Management Console snap-in and use the name of the top-level domain.
6. Enter the **Bind DN** and **Password** to enable the authentication service to authenticate the firewall.

   The Bind DN account must have permission to read the LDAP directory.
7. Enter the **Bind Timeout** and **Search Timeout** in seconds (default is 30 for both).
8. Enter the **Retry Interval** in seconds (default is 60).
9. (Optional) If you want the endpoint to use SSL or TLS for a more secure connection with the directory server, enable the option to **Require SSL/TLS secured connection** (enabled by default). The protocol that the endpoint uses depends on the server port:
   - 389 (default)—TLS (Specifically, the Dedicated Log Collector uses the StartTLS operation, which upgrades the initial plaintext connection to TLS.)
   - 636—SSL
   - Any other port—The Dedicated Log Collector first attempts to use TLS. If the directory server doesn't support TLS, the Dedicated Log Collector falls back to SSL.
10. (Optional) For additional security, enable to the option to **Verify Server Certificate for SSL sessions** so that the endpoint verifies the certificate that the directory server presents for SSL/TLS connections. To enable verification, you must also enable the

 ©2025 Palo Alto Networks, Inc.

option to **Require SSL/TLS secured connection**. For verification to succeed, the certificate must meet one of the following conditions:

- It is in the list of Panorama certificates: **Panorama** > **Certificate Management** > **Certificates** > **Device Certificates.** If necessary, import the certificate into Panorama.
- The certificate signer is in the list of trusted certificate authorities: **Panorama** > **Certificate Management** > **Certificates**.

11. Click **OK** to save the server profile.

STEP 4 | Configure the authentication for the Dedicated Log Collector.

1. Select **Panorama** > **Managed Collectors** and select the Dedicated Log Collector you previously added.
2. Configure the authentication **Timeout Configuration** for the Dedicated Log Collector.
   1. Enter the number of **Failed Attempt**s before a user is locked out of the Dedicated Log Collector CLI.
   2. Enter the **Lockout Time**, in minutes, for which the Dedicated Log Collector locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access the Dedicated Log Collector.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
3. Add the Dedicated Log Collector administrators.

   Administrators may either be added as a local administrator or as an imported Panorama administrator—but not both. Adding the same administrator as both a local administrator and as an imported Panorama administrator is not supported and causes the Panorama

 ©2025 Palo Alto Networks, Inc.

commit to fail. For example, the commit to Panorama fails if you add **admin1** as both a local and Panorama administrator.

- Configure the local administrators.

  Configure new administrators unique to the Dedicated Log Collector. These administrators are specific to the Dedicated Log Collector for which they are created and you manage these administrators from this table.

  1. **Add** one or more new local administrator.
  2. Enter a **Name** for the local administrator.
  3. Assign an **Authentication Profile** you previously created.

     LDAP authentication profiles are supported only for individual local administrators.
  4. Enable (check) **Use Public Key Authentication (SSH)** to import a public key file for authentication.
  5. Select a **Password Profile** to set the expiration parameters.
- Import existing Panorama administrators

  Import existing administrators configured on Panorama. These administrators are configured and managed on Panorama and imported to Dedicated Log Collector.

  1. **Add** an existing Panorama administrator

4. Click **OK** to save the Dedicated Log Collector authentication configuration.

STEP 5 | Configure the authentication for the Dedicated Log Collector.

1. Select **Panorama** > **Managed Collectors** and select the Dedicated Log Collector you previously added.
2. Select the **Authentication Profile** you configured in the previous step.
3. Configure the authentication **Timeout Configuration** for the Dedicated Log Collector.
   1. Enter the number of **Failed Attempt**s before a user is locked out of the Dedicated Log Collector CLI.
   2. Enter the **Lockout Time**, in minutes, for which the Dedicated Log Collector locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access the Dedicated Log Collector.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
4. Add the Dedicated Log Collector administrators.

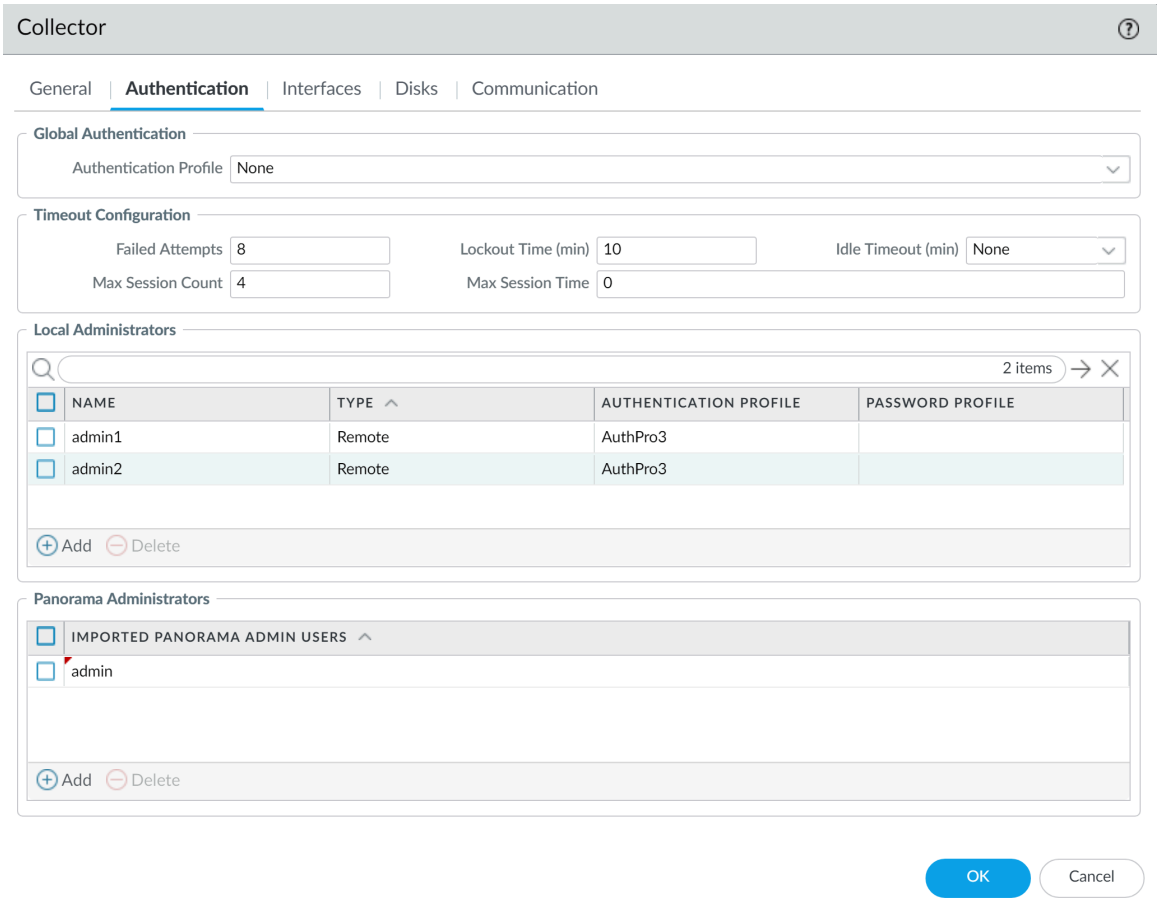
   You must add the administrator (**admin**) as either a local administrator or as an imported Panorama administrator—but not both. The push to managed collectors fails if an

 ©2025 Palo Alto Networks, Inc.

administrator is not added or if the administrator is added as both a local administrator and as an imported Panorama administrator.

1. **Add** and configure new administrators unique to the Dedicated Log Collector. These administrators are specific to the Dedicated Log Collector for which they are created and you manage these administrators from this table.
2. **Add** any administrators configured on Panorama. These administrators are created on Panorama and imported to the Dedicated Log Collector.

5. Click **OK** to save the Dedicated Log Collector authentication configuration.



STEP 6 | **Commit** and then **Commit and Push** your configuration changes.

STEP 7 | Log in to the Panorama CLI of the Dedicated Log Collector to verify you can successfully access the Dedicated Log Collector using the local admin user.

 ©2025 Palo Alto Networks, Inc.

Manage Collector Groups

A Collector Group is 1 to 16 Log Collectors that operate as a single logical unit for collecting firewall logs. You must assign at least one Log Collector to a Collector Group for firewalls to successfully send logs to a Log Collector. Firewall logs are dropped if there is no Collector Group configured or none of the Log Collectors are assigned to a Collector Group. You can configure a Collector Group with multiple Log Collectors to ensure log redundancy or to accommodate logging rates that exceed the capacity of a single Log Collector (see Panorama Models). To understand the risks and recommended mitigations, see Caveats for a Collector Group with Multiple Log Collectors.

The M-700, M-600, M-500, M-300, and M-200 appliances in Panorama mode have a predefined Collector Group that contains a predefined local Log Collector. You can edit all the settings of the predefined Collector Group except its name (default).

If you delete a Collector Group, you will lose logs.

Palo Alto Networks recommends preserving the predefined Log Collector and Collector Group on the Panorama management server, regardless of whether Panorama also manages Dedicated Log Collectors.

If you switch an M-Series appliance from Panorama mode to Log Collector mode, the appliance will lose its predefined Collector Group and Log Collector. You would then have to Set Up the M-Series Appliance as a Log Collector, add it as a managed collector to Panorama, and configure a Collector Group to contain the managed collector.

- Configure a Collector Group
- Configure Authentication with Custom Certificates Between Log Collectors
- Move a Log Collector to a Different Collector Group
- Remove a Firewall from a Collector Group

Configure a Collector Group

Before configuring Collector Groups, decide whether each one will have a single Log Collector or multiple Log Collectors (up to 16). A Collector Group with multiple Log Collectors supports higher logging rates and log redundancy but has the following requirements:

- In any single Collector Group, all the Log Collectors must run on the same Panorama model: all M-700 appliances, all M-600 appliances, all M-500 appliances, all M-300 appliances, all M-200 appliances, or all Panorama virtual appliances.
- Log redundancy is available only if each Log Collector has the same number of logging disks. To add disks to a Log Collector, see Increase Storage on the M-Series Appliance.
- (Best Practice) All Log Collectors in the same Collector Group should be in the same local area network (LAN). Avoid adding Log Collectors in the same or different wide area networks (WAN) to the same Collector Group as network disruption are much more common and may result in log data loss. Additionally, it is recommended that Log Collectors in the same Collector Group be in close physical proximity to each other to allow Panorama to quickly query the Log Collectors when needed.

 ©2025 Palo Alto Networks, Inc.

You must add the Log Collector to a Collector Group and push the Collector Group configuration to the Log Collector before it can start ingesting firewall logs. Otherwise, the ElasticSearch health status displays as degraded and the Log Collector cannot ingest logs until added to a Collector Group.

STEP 1 | Perform the following tasks before configuring the Collector Group.

1. Add a Firewall as a Managed Device for each firewall that you will assign to the Collector Group.
2. Configure a Managed Collector for each Log Collector that you will assign to the Collector Group.

STEP 2 | Add the Collector Group.

1. Access the Panorama web interface, select **Panorama** > **Collector Groups**, and **Add** a Collector Group or edit an existing one.
2. Enter a **Name** for the Collector Group if you are adding one.

   You cannot rename an existing Collector Group.
3. Enter the **Minimum Retention Period** in days (1 to 2,000) for which the Collector Group will retain firewall logs.

   By default, the field is blank, which means the Collector Group retains logs indefinitely.
4. **Add** Log Collectors (1 to 16) to the Collector Group Members list.
5. (Recommended) **Enable log redundancy across collectors** if you are adding multiple Log Collectors to a single Collector group.

   Redundancy ensures that no logs are lost if any one Log Collector becomes unavailable. Each log will have two copies and each copy will reside on a different Log Collector. For example, if you have two Log Collectors in the collector group the log is written to both Log Collectors.

   Enabling redundancy creates more logs and therefore requires more storage capacity, reducing storage capability in half. When a Collector Group runs out of space, it deletes older logs. Redundancy also doubles the log processing traffic in a Collector Group, which reduces its maximum logging rate by half, as each Log Collector must distribute a copy of each log it receives.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Assign Log Collectors and firewalls to the Collector Group.

1. Select **Device Log Forwarding** and **Add** *log forwarding preference lists* for the firewalls.

   Log data is forwarded over a separate TCP channel. By adding a log forwarding preference, list you enable the creation of separate TCP connections for forwarding log data.

   A preference list determines the order in which Log Collectors receive logs from a firewall. If a log forwarding preference list is not assigned, you may encounter one of the following scenarios:

   - *If Panorama is in Management Only mode, Panorama drops all incoming logs.*
   - *If the local Log Collector is not configured as a managed collector when Panorama is in Panorama mode, Panorama drops all incoming logs.*
   - *If the local Log Collector is configured as a managed collector when Panorama is in Panorama mode, incoming logs are received but the Panorama may act as a bottleneck because all managed firewalls are forwarding logs to the local Log Collector first before being redistributed to other available Log Collectors.*

   1. In the Devices section, **Modify** the list of firewalls and click **OK**.
   2. In the Collectors section, **Add** Log Collectors to the preference list.

      If you enabled redundancy in Step 2, it is recommended to add at least two Log Collectors. If you assign multiple Log Collectors, the first one will be the primary; if the primary becomes unavailable, the firewalls send logs to the next Log Collector in the list. To change the priority of a Log Collector, select it and **Move Up** (higher priority) or **Move Down** (lower priority).
   3. Click **OK**.

STEP 4 | Define the storage capacity (log quotas) and expiration period for each log type.

1. Return to the **General** tab and click the **Log Storage** value.

   *If the field displays 0MB, verify that you enabled the disk pairs for logging and committed the changes (see Configure a Managed Collector, **Disks** tab).*
2. Enter the log storage **Quota(%)** for each log type.
3. Enter the **Max Days** (expiration period) for each log type (1 to 2,000).

   By default, the fields are blank, which means the logs never expire.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Commit and verify your changes.

1. Select **Commit** > **Commit and Push** and then **Commit and Push** your changes to Panorama and the Collector Group you configured.
2. Select **Panorama** > **Managed Collectors** to verify the Log Collectors in the Collector Group are:
   - **Connected to Panorama**—The Connected column displays a check mark icon to indicate that a Log Collector is connected to Panorama.
   - **Synchronized with Panorama**—The Configuration Status column indicates whether a Log Collector is `In Sync` (green icon) or `Out of Sync` (red icon) with Panorama.

STEP 6 | Troubleshoot Connectivity to Network Resources to verify your firewalls successfully connected to the Log Collector.

STEP 7 | Next steps...

1. Configure Log Forwarding to Panorama.

   The Collector Group won't receive firewall logs until you configure the firewalls to forward to Panorama.
2. (Optional) Configure Log Forwarding from Panorama to External Destinations.

   You can configure each Collector Group to forward logs to separate destinations (such as a syslog server).

Configure Authentication with Custom Certificates Between Log Collectors

Complete the following procedure to configure custom certificates for communication between Log Collectors. You must configure secure server communication and secure client communication on each Log Collector in a Collector Group because the server and client roles are chosen dynamically. Use custom certificates to create a unique chain of trust that ensures mutual authentication between the members of your Log Collector Group.

For more information about using custom certificates, see How Are SSL/TLS Connections Mutually Authenticated?

STEP 1 | Obtain key pairs and certificate authority (CA) certificates for each Log Collector.

STEP 2 | Import the CA certificate to validate the identity of the client Log Collector, the server key pair, and the client key pair for each Log Collector in the Collector Group.

1. Select **Panorama** > **Certificate Management** > **Certificates** > **Import**.
2. Import the CA certificate, server key pair, and client key pair.
3. Repeat th step for the each Log Collector.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Configure a certificate profile that includes the root CA and intermediate CA for secure server communication. This certificate profile defines the authentication between Log Collectors.

1. Select **Panorama** > **Certificate Management** > **Certificate Profile**.
2. Configure a certificate profile.

   If you configure an intermediate CA as part of the certificate profile, you must also include the root CA.

STEP 4 | Configure the certificate profile for secure client communication. You can configure this profile on each client Log Collector individually or you can push the configuration from Panorama™ to managed Log Collectors.

If you are using SCEP for the client certificate, configure a SCEP profile instead of a certificate profile.

1. Select **Panorama** > **Certificate Management** > **Certificate Profile**.
2. Configure a Certificate Profile.

STEP 5 | Configure an SSL/TLS service profile.

1. Select **Panorama** > **Certificate Management** > **SSL/TLS Service Profile**.
2. Configure an SSL/TLS service profile to define the certificate and protocol that the Log Collectors use for SSL/TLS services.

STEP 6 | After deploying custom certificates on all Log Collectors, enforce custom-certificate authentication.

1. Select **Panorama** > **Collector Groups** and select the Collector Group.
2. On the General tab, **Enable secure inter LC Communication**.

   If you enable secure inter LC communication and your Collector Group includes a local Log Collector, a link should appear that stating that the **Log Collector on local Panorama is using the secure client configuration from Panorama** > **Secure Communication Settings**. You can click this link to open the Secure Communication Settings dialog and configure the secure server and secure client settings for the Local Log Collector from there.
3. Click **OK**.
4. **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Configure secure server communication on each Log Collector.

1. Select **Panorama** > **Managed Collectors** for Dedicated Log Collectors or **Panorama** > **Setup** > **Management** and **Edit** the Secure Communication Settings for a Local Log Collector.
2. For Dedicated Log Collectors, click the Log Collector and select **Communications**.
3. Enable the **Customize Secure Server Communication** feature.
4. Select the SSL/TLS service profile from the **SSL/TLS Service Profile** drop-down. This SSL/TLS service profile applies to all SSL connections between Log Collectors.
5. Select the **Certificate Profile** from the drop-down.
6. Verify that the **Custom Certificates Only** is disabled (cleared). This allows the inter Log Collector communication to continue with the predefined certificate while configuring to custom certificates.
7. Set the disconnect wait time—the number of minutes Log Collectors wait before breaking and reestablishing the connection with other Log Collectors. This field is empty by default (range is 0 to 44,640).
8. (Optional) Configure an authorization list. The authorization list adds an additional layer of security beyond certificate authentication. The authorization list checks the client certificate Subject or Subject Alt Name. If the Subject or Subject Alt Name presented with the client certificate does not match an identifier in the authorization list, authentication is denied.
    1. **Add** an Authorization List.
    2. Select the **Subject** or **Subject Alt Name** configured in the certificate profile as the Identifier type.
    3. Enter the Common Name if the identifier is `Subject` or an IP address, hostname, or email if the identifier is `Subject Alt Name`.
    4. Click **OK**.
    5. Enable the **Check Authorization List** option to configure Panorama to enforce the authorization list.
9. Click **OK**.
10. **Commit** your changes.

After committing these changes, the disconnect wait time countdown begins. When the wait time ends, Log Collectors in the Collector Group cannot connect without the configured certificates.

STEP 8 | Configure secure client communication on each Log Collector.

1. Select **Panorama** > **Managed Collectors** for Dedicated Log Collectors or **Panorama** > **Setup** > **Management** and **Edit** the Secure Communication Settings for a Local Log Collector.
2. For Dedicated Log Collectors, click the Log Collector and select **Communications**.
3. Under Secure Client Communications, select the **Certificate Type**, **Certificate**, and **Certificate Profile** from the respective drop-downs.
4. Click **OK**.
5. **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

Move a Log Collector to a Different Collector Group

M-700, M-600, M-500, M-300, M-200, and Panorama virtual appliances can have one or more Log Collectors in each Collector Group. You assign Log Collectors to a Collector Group based on the logging rate and log storage requirements of that Collector Group. If the rates and required storage increase in a Collector Group, the best practice is to Increase Storage on the M-Series Appliance or Configure a Collector Group with additional Log Collectors. However, in some deployments, it might be more economical to move Log Collectors between Collector Groups.

When a Log Collector is local to an M-700, M-600, M-500, M-300, or M-200 in Panorama mode, move it only if the appliance is the passive peer in a high availability (HA) configuration. HA synchronization applies the configurations associated with the new Collector Group. Never move a Log Collector that is local to the active HA peer.

In any single Collector Group, all the Log Collectors must run on the same Panorama model: all M-700 appliances, all M-600 appliances, all M-500 appliances, all M-300 appliances, all M-200 appliances, or all Panorama virtual appliances.

Log redundancy is available only if each Log Collector has the same number of logging disks. To add disks to a Log Collector, see Increase Storage on the M-Series Appliance.

STEP 1 | Remove the Log Collector from Panorama management.

1. Select **Panorama** > **Collector Groups** and edit the Collector Group that contains the Log Collector you will move.
2. In the Collector Group Members list, select and **Delete** the Log Collector.
3. Select **Device Log Forwarding** and, in the Log Forwarding Preferences list, perform the following steps for each set of firewalls assigned to the Log Collector you will move:
   1. In the Devices column, click the link for the firewalls assigned to the Log Collector.
   2. In the Collectors column, select and **Delete** the Log Collector.

      *To reassign the firewalls, **Add** the new Log Collector to which they will forward logs.*
   3. Click **OK** twice to save your changes.
4. Select **Panorama** > **Managed Collectors** and then select and **Delete** the Log Collector you will move.

STEP 2 | Configure a Collector Group.

Add the Log Collector to its new Collector Group and assign firewalls to the Log Collector.

*When you push changes to the Collector Group configuration, Panorama starts redistributing logs across the Log Collectors. This process can take hours for each terabyte of logs. During the redistribution process, the maximum logging rate is reduced. In the **Panorama** > **Collector Groups** page, the Log Redistribution State column indicates the completion status of the process as a percentage.*

STEP 3 | Configure Log Forwarding to Panorama for the new Collector Group you configured.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Select **Commit** > **Commit and Push** to commit your changes to Panorama and push the changes to device groups, templates, and Collector Groups if you have not already done so.

Remove a Firewall from a Collector Group

If you use a Panorama virtual appliance in Legacy mode to manage Dedicated Log Collectors, you have the option to forward firewall logs to Panorama instead of forwarding to the Log Collectors. For such cases, you must remove the firewall from the Collector Group; the firewall will then automatically forward its logs to Panorama.

To temporarily remove the log forwarding preference list on the firewall, you can delete it using the CLI on the firewall. You must however, remove the assigned firewalls in the Collector Group configuration on Panorama. Otherwise, the next time you push changes to the Collector Group, the firewall will be reconfigured to send logs to the assigned Log Collector.

STEP 1 | Select **Panorama** > **Collector Groups** and edit the Collector Group.

STEP 2 | Select **Device Log Forwarding**, click the firewall in the Devices list, **Modify** the Devices list, clear the check box of the firewall, and click **OK** three times.

STEP 3 | Select **Commit** > **Commit and Push** and then **Commit and Push** your changes to Panorama and the Collector Group from which you removed the firewall.

 ©2025 Palo Alto Networks, Inc.

Configure Log Forwarding to Panorama

Each firewall stores its log files locally by default and cannot display the logs that reside on other firewalls. Therefore, to achieve global visibility into the network activity that all your firewalls monitor, you must forward all firewall logs to Panorama and Use Panorama for Visibility. In cases where some teams in your organization can achieve greater efficiency by monitoring only the logs that are relevant to their operations, you can create forwarding filters based on any log attributes (such as threat type or source user). For example, a security operations analyst who investigates malware attacks might be interested only in Threat logs with the type attribute set to wildfire-virus.

The following steps describe how to use Panorama templates and device groups for configuring multiple firewalls to forward logs.

If Panorama manages firewalls running software versions earlier than PAN-OS 7.0, specify a WildFire® server from which Panorama can gather analysis information for WildFire samples that those firewalls submit. Panorama uses the information to complete WildFire Submissions logs that are missing field values introduced in PAN-OS 7.0. Firewalls running earlier releases won't populate those fields. To specify the server, select ***Panorama*** *>* ***Setup*** *>* ***WildFire****, edit the General Settings, and enter the* ***WildFire Private Cloud*** *name. The default is* ***wildfire-public-cloud****, which is the WildFire cloud hosted in the United States.*

You can also forward firewall logs to external services (such as a syslog server). For details, see Log Forwarding Options.

STEP 1 | Add a Device Group for the firewalls that will forward logs.

Panorama requires a device group to push a Log Forwarding profile to firewalls. Create a new device group or assign the firewalls to an existing device group.

STEP 2 | Add a Template for the firewalls that will forward logs.

Panorama requires a template to push log settings to firewalls. Create a new template or assign the firewalls to an existing template.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Create a Log Forwarding profile.

The profile defines the destinations for Traffic, Threat, WildFire Submission, URL Filtering, Data Filtering, Tunnel and Authentication logs.

1. Select **Objects** > **Log Forwarding**, select the **Device Group** of the firewalls that will forward logs, and **Add** a profile.
2. Enter a **Name** to identify the Log Forwarding profile.
3. **Add** one or more *match list profiles*.

   The profiles specify log query filters, forwarding destinations, and automatic actions such as tagging. For each match list profile:

   1. Enter a **Name** to identify the profile.
   2. Select the **Log Type**.
   3. In the **Filter** drop-down, select **Filter Builder**. Specify the following and then **Add** each query:

      Connector logic (and/or)

      Log **Attribute**

      Operator to define inclusion or exclusion logic

      Attribute **Value** for the query to match
   4. Select **Panorama/Strata Logging Service**.
4. Click **OK** to save the Log Forwarding profile.

STEP 4 | Assign the Log Forwarding profile to policy rules and network zones.

Security, Authentication, and DoS Protection rules support log forwarding. In this example, you assign the profile to a Security rule.

Perform the following steps for each rule that will trigger log forwarding:

1. Select the rulebase (for example, **Policies** > **Security** > **Pre Rules**), select the **Device Group** of the firewalls that will forward logs, and edit the rule.
2. Select **Actions** and select the **Log Forwarding** profile you created.
3. Set the **Profile Type** to **Profiles** or **Group**, and then select the security profiles or **Group Profile** required to trigger log generation and forwarding for:
   - Threat logs—Traffic must match any security profile assigned to the rule.
   - WildFire logs—Traffic must match a WildFire Analysis profile assigned to the rule.
4. For Traffic logs, select **Log At Session Start** and/or **Log At Session End**.

   Log At Session Start consumes more resources than logging only at the session end. In most cases, you only **Log At Session End**. Enable both **Log At Session Start** and **Log At Session End** only for troubleshooting, for long-lived tunnel sessions such as GRE tunnels (you can't see these sessions in the ACC unless you log at the start of the session), and to gain visibility into Operational Technology/Industrial Control Systems (OT/ICS) sessions, which are also long-lived sessions.
5. Click **OK** to save the rule.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Configure the destinations for System logs, Configuration logs, User-ID™ logs, and HIP Match logs.

Panorama generates Correlation logs based on the firewall logs it receives, rather than aggregating Correlation logs from firewalls.

1. Select **Device** > **Log Settings** and select the **Template** of the firewalls that will forward logs.
2. For each log type that the firewall will forward, see step Add one or more match list profiles.

STEP 6 | (PA-7000 Series firewalls only) Configure a log card interface to perform log forwarding.

When you configure a data port on one of the PA-7000 Series Network Processing Cards (NPCs) as a Log Card interface, the firewall will automatically begin using this interface to forward logs to the logging destinations you configure and forward files for WildFire analysis.

 ©2025 Palo Alto Networks, Inc.

Make sure that the interface you configure can reach the log forwarding destinations and the WildFire cloud, WildFire appliance, or both.

Because PA-7000 Series firewall can now forward logs to Panorama, Panorama no longer treats the PA-7000 Series firewalls it manages as Log Collectors. If you have not configured the PA-7000 Series firewalls to forward logs to Panorama, all logs a managed PA-7000 Series firewall generates are only viewable from the local firewall and not from Panorama. If you do not yet have a log forwarding infrastructure that is capable of handling the logging rate and volume from the PA-7000 Series firewalls, starting with PAN-OS 8.0.8 you can enable Panorama to directly query PA-7000 Series firewalls when monitoring logs. To use this functionality, both Panorama and the PA-7000 Series firewalls must be running PAN-OS 8.0.8 or later. Enable Panorama to directly query PA-7000 Series firewalls by entering the following command from the Panorama CLI:

```
> debug reportd send-request-to-7k yes
```

After running this command, you will be able to view logs for managed PA-7000 Series firewalls on the Panorama ***Monitor*** *tab. Additionally, as with all managed devices, you can also generate reports that include PA-7000 Series log data by selecting* ***Remote Device Data*** *as the* ***Data Source****. If you later decide to enable the PA-7000 Series firewalls to forward logs to Panorama, you must first disable this option using the* ***debug reportd send-request-to-7k no*** *command.*

1. Select **Network** > **Interfaces** > **Ethernet**, select the **Template** of the firewalls that will forward logs, and **Add Interface**.
2. Select the **Slot** and **Interface Name**.
3. Set the **Interface Type** to **Log Card**.
4. Enter the **IP Address**, **Default Gateway**, and (for IPv4 only) **Netmask**.
5. Select **Advanced** and specify the **Link Speed**, **Link Duplex**, and **Link State**.

   These fields default to ***auto****, which specifies that the firewall automatically determines the values based on the connection. However, the minimum recommended* ***Link Speed*** *for any connection is* ***1000*** *(Mbps).*
6. Click **OK** to save your changes.

STEP 7 | Configure Panorama to receive the logs.

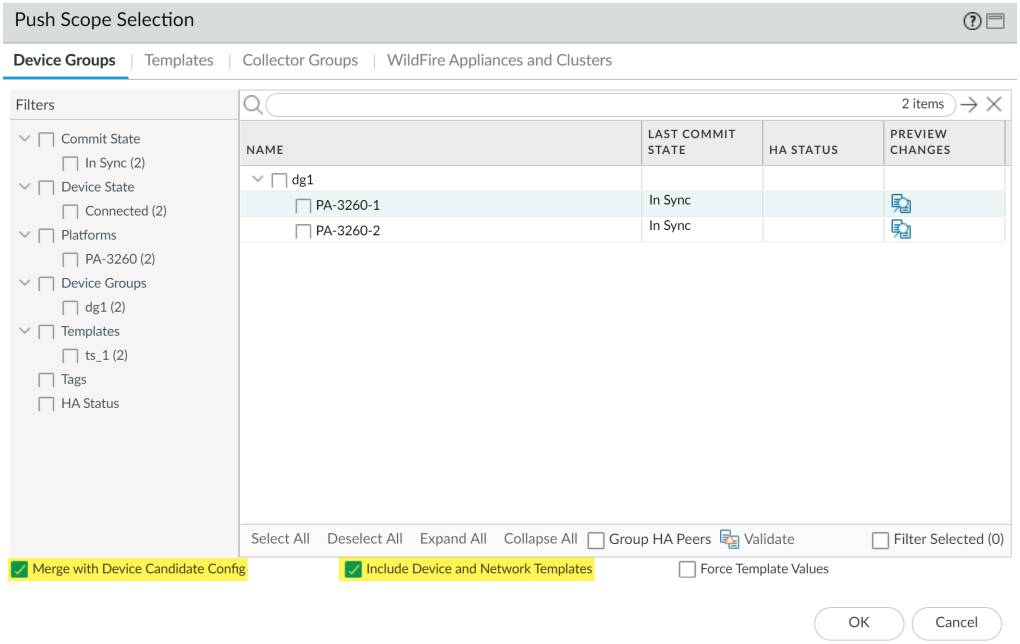
If you will forward logs to a Panorama virtual appliance in Legacy mode, you can skip this step.

1. For each Log Collector that will receive logs, Configure a Managed Collector.
2. Configure a Collector Group to assign firewalls to specific Log Collectors for log forwarding.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Commit your configuration changes.

1. Select **Commit** > **Commit and Push** and **Edit Selections.**
2. Select **Merge with Device Candidate Config** and **Include Device and Network Templates**.



3. Click **Collector Groups** to verify your target Collector Group is selected, and click **OK**.
4. **Commit and Push** your changes to Panorama and push the changes to the device groups, templates, and Collector Groups.
5. Verify Log Forwarding to Panorama to confirm that your configuration is successful.

> *To change the log forwarding mode that the firewalls use to send logs to Panorama, you can* Modify Log Forwarding and Buffering Defaults. *You can also* Manage Storage Quotas and Expiration Periods for Logs and Reports.

 ©2025 Palo Alto Networks, Inc.

Configure Syslog Forwarding to External Destinations

In the case of a deployment with a high rate of log generation, you can forward syslogs over an Ethernet interface to prevent loss of logs and reduce the load on the management interface, which optimizes management operations.

Syslog forwarding using an Ethernet interface is supported only for a Panorama™ management server in Panorama mode or in Log Collector mode. Additionally, you can enable syslog forwarding on only a single interface regardless whether Panorama is in Panorama mode or Log Collector mode.

STEP 1 | Log in to the Panorama web interface.

STEP 2 | Configure a Managed Collector.

STEP 3 | Configure a Collector Group.

On the M-Series appliance, a default Collector Group is predefined and already contains the local Log Collector as a member. However on the Panorama virtual appliance, you must add the Collector Group and add the local Log Collector as a member. For both configurations, you need to assign firewalls to a Log Collector for log forwarding.

STEP 4 | Configure a Syslog server profile.

1. Select **Panorama** > **Server Profiles** > **Syslog** and **Add** a new syslog server profile.
2. Enter a **Name** for the syslog server profile.
3. For each syslog server, **Add** the information that Panorama or the Dedicated Log Collector requires to connect to it:
   - **Name**—Unique name for the syslog server.
   - **Syslog Server**—IP address or fully qualified domain name (FQDN) of the syslog server.
   - **Transport**—Select **UDP**, **TCP**, or **SSL** as the method of communication with the syslog server.
   - **Port**—The port number to use when sending syslog messages (default is UDP on port 514); you must use the same port number on Panorama and on the Dedicated Log Collector.
   - **Format**—Select the syslog message format to use: **BSD** (default) or **IETF**. Traditionally, **BSD** format is over UDP and **IETF** format is over TCP or SSL.
   - **Facility**—Select the syslog standard value (default is **LOG_USER**) to calculate the priority (PRI) field in your syslog server implementation. Select the value that maps to how you use the PRI field to manager your syslogs.
4. (Optional) To customize which format of syslog messages that Panorama or the Dedicated Log Collector sends, select **Custom Log Format**. For details about how to create custom formats for the various log types, refer to the Common Event Format Configuratiuon Guide.
5. Click **OK** to save the syslog server profile.

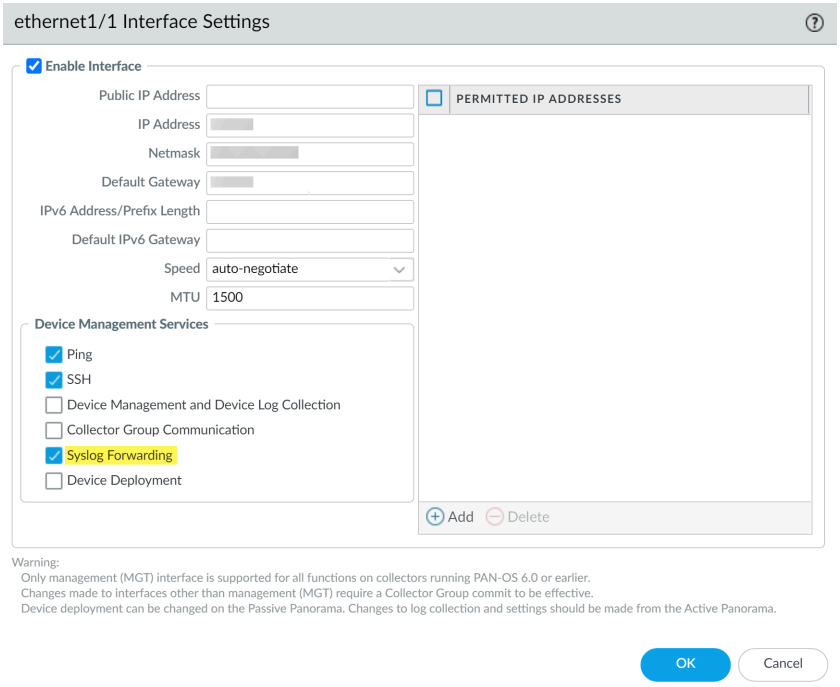
 ©2025 Palo Alto Networks, Inc.

STEP 5 | Configure an Ethernet interface for forwarding syslogs.

By default, syslog forwarding is enabled on the management interface and is supported on only one interface at a time.

- Configure an Ethernet interface on the local Log Collector from the Panorama web interface.
  1. Select **Panorama** > **Setup** > **Interfaces** and select an Ethernet interface.
  2. **Enable Interface**.
  3. Configure the Ethernet interface as appropriate.
  4. In the Device Management Services section, enable **Syslog Forwarding**.
  5. Select **Yes** to confirm your syslog forwarding change.

     You can only on a single Ethernet interface on the local Log Collector.
  6. Click **OK** to save your changes.
  7. **Commit** and then **Commit and Push** your configuration changes.

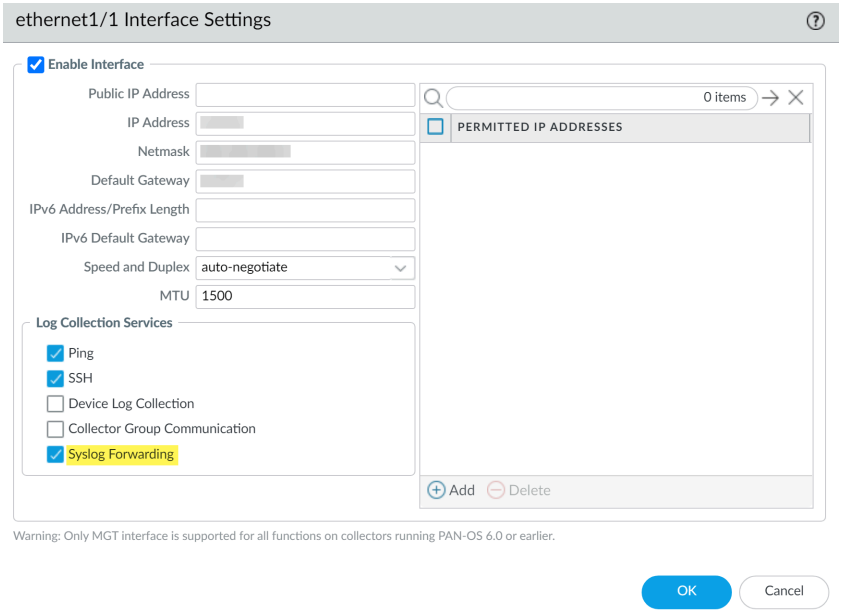


- Configure an Ethernet interface on a Dedicated Log Collector.
  1. Select **Panorama** > **Managed Collectors** and select a Dedicated Log Collector.
  2. **Enable Interface**.
  3. Configure the Ethernet interface as appropriate.
  4. In the Log Collection Services section, enable **Syslog Forwarding**.
  5. Select **Yes** to confirm your syslog forwarding change.

     You can only on a single Ethernet interface on the Dedicated Log Collector.

 ©2025 Palo Alto Networks, Inc.

6. Click **OK** to save your changes.
7. **Commit** and then **Commit and Push** your configuration changes.



- Configure an Ethernet interface on the local Log Collector or Dedicated Log Collector from the Panorama CLI.

  To successfully configure syslog forwarding over an Ethernet interface from the CLI, you must first disable syslog forwarding on the management interface and then enable syslog forwarding on the Ethernet interface from the CLI; Panorama does not automatically disable syslog forwarding over the management interface you enable syslog forwarding on an

 ©2025 Palo Alto Networks, Inc.

Ethernet interface from the CLI so syslog forwarding continues over the management interface if you enable it on both the management and Ethernet interfaces.

1. Log in to the Panorama CLI
2. Disable syslog forwarding on the management interface:

```
admin@Panorama> configure
```

```
admin@Panorama> set log-collector <Log Collector Serial Number>
 deviceconfig system service disable-syslog-forwarding yes
```

3. Enable syslog forwarding on the Ethernet interface:

```
admin@Panorama> configure
```

```
admin@Panorama> set log-collector <Log Collector Serial Number>
 deviceconfig system eth<Interface Number>  service disable-
syslog-forwarding no
```

```
admin@Panorama> commit
```

4. Commit your configuration changes:

```
admin@Panorama> run commit-all log-collector-config log-
collector-group <Collector Group name>
```

STEP 6 | Configure Log Forwarding to Panorama.

STEP 7 | Configure syslog forwarding from Panorama to a syslog server.

 ©2025 Palo Alto Networks, Inc.

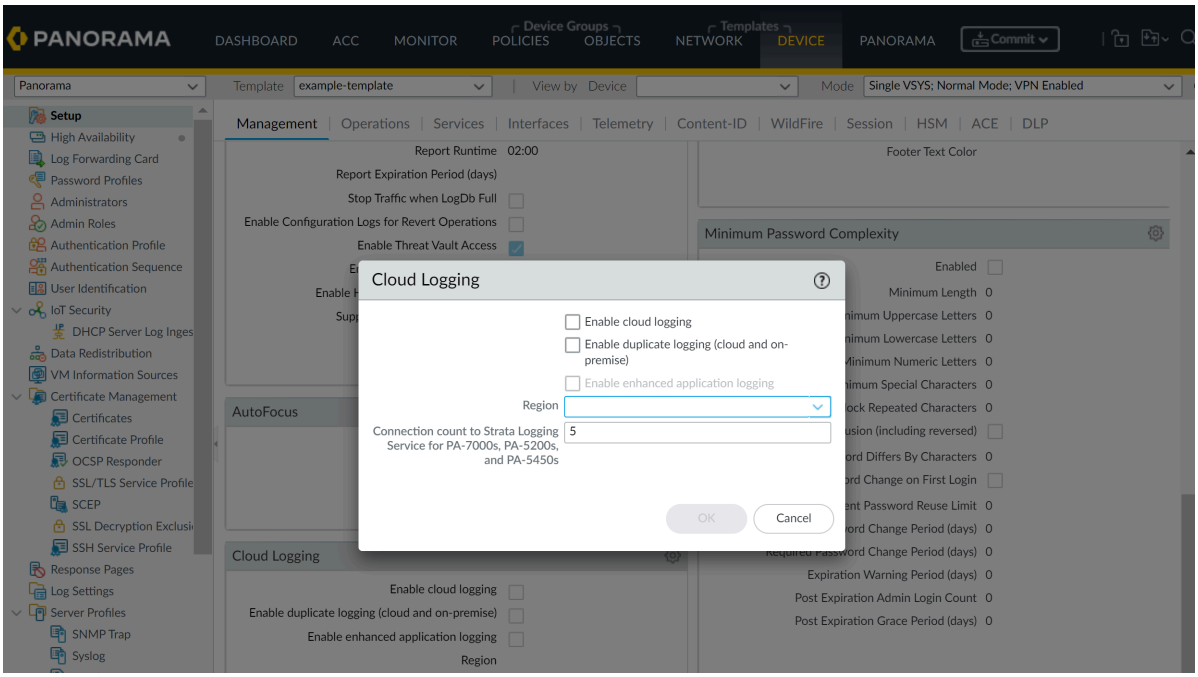
Forward Logs to Strata Logging Service

Strata Logging Service is Palo Alto Networks' cloud-based logging infrastructure. Before you can configure your managed firewalls to send logs to Strata Logging Service (previously called the Logging Service), you need to purchase a license for the volume of logs in your deployment, and install the cloud services plugin. If you already have on premise Log Collectors, you can use Strata Logging Service to complement and augment your existing setup.

STEP 1 | Install Panorama Plugins.

STEP 2 | Configure the firewalls to send logs to Strata Logging Service.

For firewalls running PAN-OS 8.1 or later releases, you can opt to send logs to both the Strata Logging Service and to your Panorama and on premise log collection setup when you select **Enable Duplicate Logging (Cloud and On-Premise)**. When enabled, the firewalls that belong to the selected Template will save a copy of the logs to both locations. You may select either **Enable Duplicate Logging (Cloud and On-Premise)** or **Enable Strata Logging Service**, but not both.



 ©2025 Palo Alto Networks, Inc.

Verify Log Forwarding to Panorama

Verify log forwarding to Panorama once you Configure Log Forwarding to Panorama or to the Strata Logging Service to test that your configuration succeeded.

After you configure log forwarding to Log Collectors, managed firewalls open a TCP connection to all configured Log Collectors. These connections timeout every sixty (60) seconds and do not indicate that the firewall has lost connection to the Log Collectors. When you configure log forwarding to a local or Dedicated Log Collector over a supported ethernet interface, the firewall traffic logs show `incomplete` sessions despite the firewall being able to successfully connect to the Log Collectors. If you configure log forwarding over the management port, no traffic logs showing `incomplete` sessions are generated. Traffic logs showing `incomplete` sessions are generated by all firewalls except for the PA-5200 and PA-7000 series firewalls.

STEP 1 | Access the firewall CLI.

STEP 2 | If you configured Log Collectors, verify that each firewall has a log forwarding preference list.

```
> show log-collector preference-list
```

If the Collector Group has only one Log Collector, the output will look something like this:

```
Forward to all: No
Log collector Preference List
Serial Number: 003001000024
IP Address: 10.2.133.48
IPV6 Address: unknown
```

STEP 3 | Verify that each firewall is forwarding logs.

```
> show logging-status
```

For successful forwarding, the output indicates that the log forwarding agent is active.

- For a Panorama virtual appliance, the agent is `Panorama`.
- For an M-Series appliance, the agent is a `LogCollector`.
- For the Strata Logging Service, the agent is `Log CollectionService.`. And the

```
‘Log Collection log forwarding agent’ is active and connected
 to <IP_address>.
```

 ©2025 Palo Alto Networks, Inc.

STEP 4 | View the average logging rate. The displayed rate will be the average logs/second for the last five minutes.

- If Log Collectors receive the logs, access the Panorama web interface, select **Panorama** > **Managed Collectors** and click the **Statistics** link in the far-right column.
- If a Panorama virtual appliance in Legacy mode receives the logs, access the Panorama CLI and run the following command: **`debug log-collector log-collection-stats show incoming-logs`**



This command also works on an M-Series appliance.

 ©2025 Palo Alto Networks, Inc.

Modify Log Forwarding and Buffering Defaults

You can define the log forwarding mode that the firewalls use to send logs to Panorama and, when configured in a high availability (HA) configuration, specify which Panorama peer can receive logs. To access these options, select **Panorama** > **Setup** > **Management**, edit the Logging and Reporting Settings, and select **Log Export and Reporting**.

- Define the log forwarding mode on the firewall: The firewalls can forward logs to Panorama (pertains to both the M-Series appliance and the Panorama virtual appliance) in either Buffered Log Forwarding mode or in the Live Mode Log Forwarding mode.

Logging Options	Description
(Best Practice) **Buffered Log Forwarding from Device** Default: Enabled	Allows each managed firewall to buffer logs and send the logs at 30-second intervals to Panorama (not user configurable). Buffered log forwarding is very valuable when the firewall loses connectivity to Panorama. The firewall buffers log entries to its local hard disk and keeps a pointer to record the last log entry that was sent to Panorama. When connectivity is restored the firewall resumes forwarding logs from where it left off. The disk space available for buffering depends on the log storage quota for the firewall model and the volume of logs that are pending roll over. If the firewall was disconnected for a long time and the last log forwarded was rolled over, all the logs from its local hard disk will be forwarded to Panorama on reconnection. If the available space on the local hard disk of the firewall is consumed, the oldest entries are deleted to allow logging of new events.
Live Mode Log Forwarding from Device This option is enabled when the check box for **Buffered Log Forwarding from Device** is cleared.	In live mode, the managed firewall sends every log transaction to Panorama at the same time as it records it on the firewall.

 ©2025 Palo Alto Networks, Inc.

- Define log forwarding preference on a Panorama virtual appliance in Legacy mode that is deployed in a high availability (HA) configuration:
  - When logging to a virtual disk, enable logging to the local disk on the primary Panorama peer only. By default, both Panorama peers in the HA configuration receive logs.

    For the 5200 and 7000 series firewalls, only the active peer receive logs.
  - When logging to an NFS (ESXi server only), enable the firewalls to send only newly generated logs to a secondary Panorama peer, which is promoted to primary, after a failover.

Logging Options	Pertains to	Description
Only Active Primary Logs to Local Disk Default: Disabled	Panorama virtual appliance in Legacy mode that is logging to a virtual disk and is deployed in an HA configuration.	Allows you to configure only the primary Panorama peer to save logs to the local disk.
Get Only New Logs on Convert to Primary Default: Disabled	Panorama virtual appliance in Legacy mode that is mounted to a Network File System (NFS) datastore, runs on a VMware ESXi server, and is deployed in an HA configuration	With NFS logging, when you have a pair of Panorama servers configured in a high availability configuration, only the primary Panorama peer mounts the NFS datastore. Therefore, the firewalls can only send logs to the primary Panorama peer, which can write to the NFS datastore. When an HA failover occurs, the **Get Only New Logs on Convert to Primary** option allows an administrator to configure the managed firewalls to send only newly generated logs to Panorama. This event is triggered when the priority of the active-secondary Panorama is promoted to primary and it can begin logging to the NFS. This behavior is typically enabled to prevent the firewalls from sending a large volume of buffered logs when connectivity to Panorama is restored after a significant period of time.

 ©2025 Palo Alto Networks, Inc.

Configure Log Forwarding from Panorama to External Destinations

Panorama enables you to forward logs to external services, including syslog, email, SNMP trap, and HTTP-based services. Using an external service enables you to receive alerts for important events, archive monitored information on systems with dedicated long-term storage, and integrate with third-party security monitoring tools. In addition to forwarding firewall logs, you can forward the logs that the Panorama management server and Log Collectors generate. The Panorama management server or Log Collector that forwards the logs converts them to a format that is appropriate for the destination (syslog message, email notification, SNMP trap, or HTTP payload). Forwarded logs have a maximum log record size of 4,096 bytes. A forwarded log with a log record size larger than the maximum is truncated at 4,096 bytes while logs that do not exceed the maximum log record size are not.



Log forwarding is supported only for supported log fields. Forwarding logs that contain unsupported log fields or pseudo-fields causes the firewall to crash.



If your Panorama management server is a Panorama virtual appliance in Legacy mode, it converts and forwards logs to external services without using Log Collectors.

You can also forward logs directly from firewalls to external services: see Log Forwarding Options.

On a Panorama virtual appliance running Panorama 5.1 or earlier releases, you can use Secure Copy (SCP) commands from the CLI to export the entire log database to an SCP server and import it to another Panorama virtual appliance. A Panorama virtual appliance running Panorama 6.0 or later releases, and M-Series appliances running any release, do not support these options because the log database on those models is too large for an export or import to be practical.

To forward logs to external services, start by configuring the firewalls to forward logs to Panorama. Then you must configure the server profiles that define how Panorama and Log Collectors connect to the services. Lastly, you assign the server profiles to the log settings of Panorama and to Collector Groups.

STEP 1 | Configure the firewalls to forward logs to Panorama.

Configure Log Forwarding to Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Configure a server profile for each external service that will receive log information.

1. Select **Panorama** > **Server Profiles** and select the type of server that will receive the log data: **SNMP Trap**, **Syslog**, **Email**, or **HTTP**.
2. Configure the server profile:
   - Configure an SNMP Trap server profile. For details on how SNMP works for Panorama and Log Collectors, refer to SNMP Support.
   - Configure a Syslog server profile. If the syslog server requires client authentication, use the **Panorama** > **Certificate Management** > **Certificates** page to create a certificate for securing syslog communication over SSL.
   - Configure an Email server profile.
   - Configure an HTTP server profile.

   Log forwarding to an HTTP server is designed for log forwarding at low frequencies and is not recommend for deployments with a high volume of log forwarding. You may experience log loss when forwarding to an HTTP server if your deployment generate a high volume of logs that need to be forwarded.

STEP 3 | Configure destinations for:

- Logs that the Panorama management server and Log Collectors generate.
- Firewall logs that a Panorama virtual appliance in Legacy mode collects.

1. Select **Panorama** > **Log Settings**.
2. **Add** one or more *match list profiles* for each log type.

   The profiles specify log query filters, forwarding destinations, and automatic actions such as tagging. For each match list profile:

   1. Enter a **Name** to identify the profile.
   2. Select the **Log Type**.
   3. In the **Filter** drop-down, select **Filter Builder**. Specify the following and then **Add** each query:

      Connector logic (and/or)

      Log **Attribute**

      Operator to define inclusion or exclusion logic

      Attribute **Value** for the query to match
   4. **Add** the server profiles you configured for each external service.
   5. Click **OK** to save the profile.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure destinations for firewall logs that Log Collectors receive.

Each Collector Group can forward logs to different destinations. If the Log Collectors are local to a high availability (HA) pair of Panorama management servers, you must log into each HA peer to configure log forwarding for its Collector Group.

1. Select **Panorama** > **Collector Groups** and edit the Collector Group that receives the firewall logs.
2. (Optional, SNMP trap forwarding only) Select **Monitoring** and configure the SNMP settings.
3. Select **Collector Log Forwarding** and **Add** configured match list profiles as necessary.
4. Click **OK** to save your changes to the Collector Group.

STEP 5 | (Syslog forwarding only) If the syslog server requires client authentication and the firewalls forward logs to Dedicated Log Collectors, assign a certificate that secures syslog communication over SSL.

Perform the following steps for each Dedicated Log Collector:

1. Select **Panorama** > **Managed Collectors** and edit the Log Collector.
2. Select the **Certificate for Secure Syslog** and click **OK**.

STEP 6 | (SNMP trap forwarding only) Enable your SNMP manager to interpret traps.

Load the Supported MIBs and, if necessary, compile them. For the specific steps, refer to the documentation of your SNMP manager.

STEP 7 | Commit and verify your configuration changes.

1. Select **Commit** > **Commit and Push** to commit your changes to Panorama and push the changes to device groups, templates, and Collector Groups.
2. Verify that the external services are receiving the log information:
   - **Email server**—Verify that the specified recipients are receiving logs as email notifications.
   - **Syslog server**—Refer to the documentation for your syslog server to verify it's receiving logs as syslog messages.
   - **SNMP manager**—Refer to the documentation for your SNMP trap server to verify it's receiving logs as SNMP traps.
   - **HTTP server**—Verify that the HTTP-based server is receiving logs in the correct payload format.

 ©2025 Palo Alto Networks, Inc.

Log Collection Deployments

The following topics describe how to configure log collection in the most typical deployments. Before starting, Plan Your Panorama Deployment according to your current and future logging needs.

The deployments in these topics all describe Panorama in a high availability (HA) configuration. Palo Alto Networks recommends HA because it enables automatic recovery (in case of server failure) of components that are not saved as part of configuration backups. In HA deployments, the Panorama management server only supports an active/passive configuration.

- Deploy Panorama with Dedicated Log Collectors
- Deploy Panorama M-Series Appliances with Local Log Collectors
- Deploy Panorama Virtual Appliances with Local Log Collectors
- Deploy Panorama Virtual Appliances in Legacy Mode with Local Log Collection

Deploy Panorama with Dedicated Log Collectors

The following figures illustrate Panorama in a distributed log collection deployment. In these examples, the Panorama management server comprises two M-Series or Panorama virtual appliances in Panorama mode that are deployed in an active/passive high availability (HA) configuration. The firewalls send logs to Dedicated Log Collectors (M-Series or Panorama virtual appliances in Log Collector mode). This is the recommended configuration if the firewalls generate over 10,000 logs/second.

If you will assign more than one Log Collector to a Collector Group, see Caveats for a Collector Group with Multiple Log Collectors to understand the requirements, risks, and recommended mitigations.

 ©2025 Palo Alto Networks, Inc.

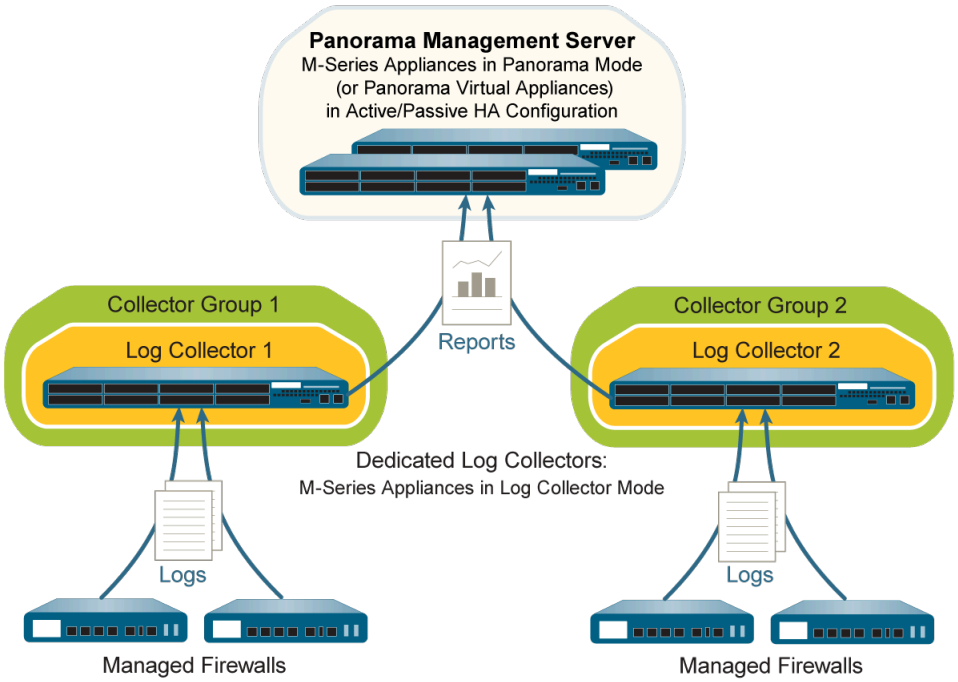


Figure 17: Single Dedicated Log Collector Per Collector Group

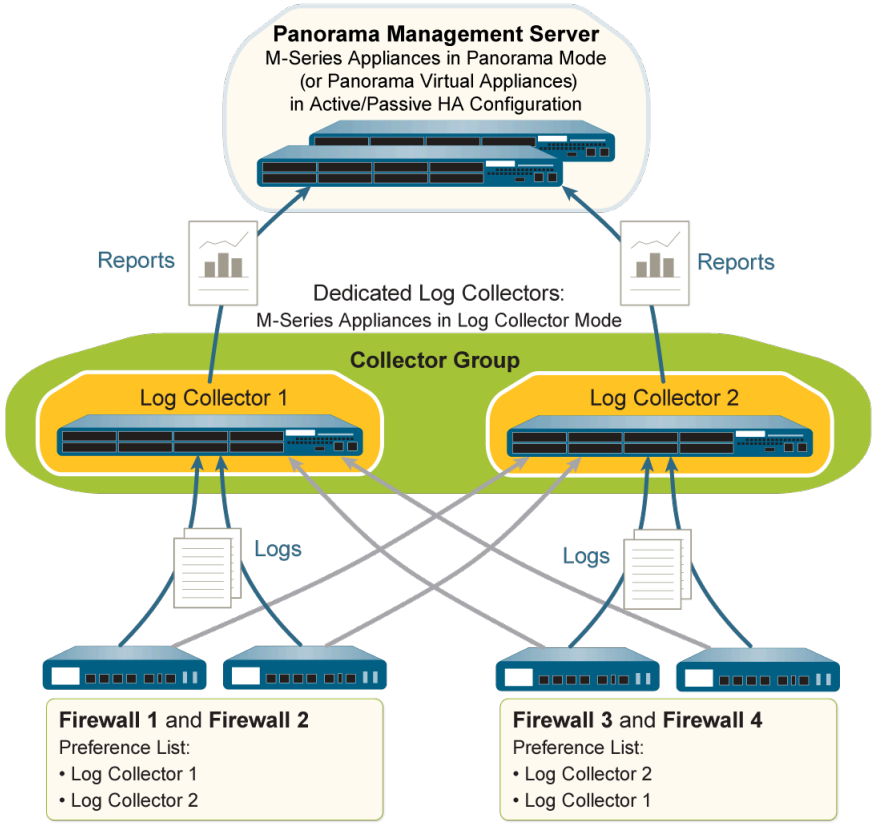


Figure 18: Multiple Dedicated Log Collectors Per Collector Group

Perform the following steps to deploy Panorama with Dedicated Log Collectors. Skip any steps you have already performed (for example, the initial setup).

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Perform the initial setup of the Panorama management server (virtual appliances or M-Series appliances) and the Dedicated Log Collectors.

For each M-Series appliance:

1. Rack mount the M-Series appliance. Refer to the M-Series Hardware Reference Guide for instructions.
2. Perform Initial Configuration of the M-Series Appliance.

   Palo Alto Networks recommends reserving the management (MGT) interface for administrative access to Panorama and dedicating separate M-Series Appliance Interfaces to other Panorama services.
3. Configure each array. This task is required to make the RAID disks available for logging. Optionally, you can add disks to Increase Storage on the M-Series Appliance.
4. Register Panorama and Install Licenses.
5. Install Content and Software Updates for Panorama.

For each virtual appliance (if any):

1. Install the Panorama Virtual Appliance.
2. Perform Initial Configuration of the Panorama Virtual Appliance.
3. Register Panorama and Install Licenses.
4. Install Content and Software Updates for Panorama.

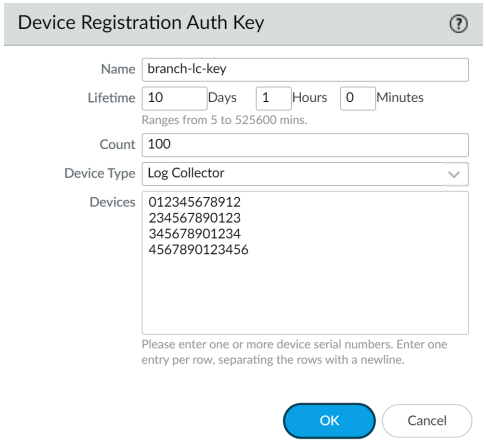
For the Panorama management server (virtual appliance or M-Series appliance), you must also Set Up HA on Panorama.

 ©2025 Palo Alto Networks, Inc.

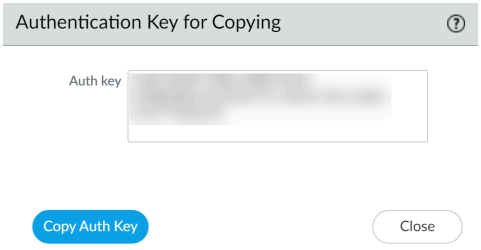
STEP 2 | On the Panorama management server, create a device registration authentication key to securely add the Dedicated Log Collector to Panorama management.

1. Log in to the Panorama Web Interface.
2. Select **Panorama** > **Device Registration Auth Key** and **Add** a new authentication key.
3. Configure the authentication key.
   - **Name**—Add a descriptive name for the authentication key.
   - **Lifetime**—Specify the key lifetime for how long you can use the authentication key to onboard new Log Collectors.
   - **Count**—Specify how many times you can use the authentication key to onboard new Log Collectors.
   - **Device Type**—Specify that this authentication key is used to authenticate only a **Log Collector**.

     *You can select **Any** to use the device registration authentication key to onboard firewalls, Log Collectors, and WildFire appliances.*
   - (Optional) **Devices**—Enter one or more device serial numbers to specify for which Log Collectors the authentication key is valid.
4. Click **OK**.



5. **Copy Auth Key** and **Close**.



 ©2025 Palo Alto Networks, Inc.

STEP 3 | Switch from Panorama mode to Log Collector mode on each Panorama management server that will be a Dedicated Log Collector.

Switching the mode of an M-Series or Panorama virtual appliance deletes any existing log data and deletes all configurations except the management access settings. After the switch, the M-Series or Panorama virtual appliance retains CLI access but loses web interface access.

1. Connect to Panorama in one of the following ways:
   - (M-Series appliances only) Attach a serial cable from your computer to the Console port on the M-Series appliance. Then use terminal emulation software (9600-8-N-1) to connect.
   - Use terminal emulation software such as PuTTY to open an SSH session to the IP address that you specified for the MGT interface of the Panorama management server during initial configuration.
2. Log in to the CLI when prompted. Use the default admin account and the password that you specified during initial configuration.
3. Switch to Log Collector mode by entering the following command:

   ```
   > request system system-mode logger
   ```
4. Enter **Y** to confirm the mode change. The Panorama management server reboots. If the reboot process terminates your terminal emulation software session, reconnect to Panorama to see the Panorama login prompt.

   *If you see a **CMS Login** prompt, this means the Log Collector has not finished rebooting. Press Enter at the prompt without typing a username or password.*
5. Log back in to the CLI.
6. Verify that the switch to Log Collector mode succeeded:

   ```
   > show system info | match system-mode
   ```

   If the mode change succeeded, the output displays:

   ```
   system-mode: logger
   ```

 ©2025 Palo Alto Networks, Inc.

STEP 4 | From the Dedicated Log Collector CLI, reset the secure connection state.

1. Reset the secure connection state.

   This command resets the managed device connection and is irreversible.

   ```
   admin> request sc3 reset
   ```

2. Restart the management server on the managed device.

   ```
   admin> debug software restart process management-server
   ```

STEP 5 | Add the device registration authentication key to the Dedicated Log Collector.

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



STEP 6 | Enable connectivity between each Log Collector and the Panorama management server.

This step is required before you can enable logging disks on the Log Collectors.

Enter the following commands at the CLI of each Log Collector. *<IPaddress1>* is for the MGT interface of the active Panorama and *<IPaddress2>* is for the MGT interface of the passive Panorama.

```
> configure
# set deviceconfig system panorama-server <IPaddress1> panorama-
server-2 <IPaddress2>
# commit
# exit
```

STEP 7 | Record the serial number of each Log Collector.

You need the serial numbers to add the Log Collectors as managed collectors on the Panorama management server.

1. At the CLI of each Log Collector, enter the following command to display its serial number.

   ```
   > show system info | match serial
   ```

2. Record the serial number.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Add each Log Collector as a managed collector.

Use the web interface of the primary Panorama management server peer to Configure a Managed Collector:

1. Select **Panorama** > **Managed Collectors** and **Add** the managed collector.
2. In the **General** tab, enter the serial number (**Collector S/N**) you recorded for the Log Collector.
3. Enter the IP address or FQDN of the active and passive Panorama HA peers in the **Panorama Server IP** field and **Panorama Server IP 2** field respectively. These fields are required.
4. Select **Interfaces**, click **Management**, and configure one or both of the following field sets for the MGT interface based on the IP protocols of your network.

   *If you configure a **Public IP Address** for the interface, Log Collectors in the Collector Group always use the public IP address for communication within the Collector Group. To ensure Log Collectors in a Collector use the private IP address to communicate, do not configure a public IP address.*

   - IPv4—**IP Address**, **Netmask**, and **Default Gateway**
   - IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
5. (Optional) Select **SNMP** if you will use an SNMP manager to monitor Log Collector statistics.

   Using SNMP requires additional steps besides configuring the Log Collector (see Monitor Panorama and Log Collector Statistics Using SNMP).
6. Click **OK** to save your changes.
7. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

   This step is required before you can enable logging disks on the Log Collectors.
8. Verify that the **Panorama** > **Managed Collectors** page lists the Log Collector you added. The Connected column displays a check mark to indicate that the Log Collector is connected to Panorama. You might have to wait a few minutes before the page displays the updated connection status.

   At this point, the Configuration Status column displays Out of Sync and the Run Time Status column displays disconnected. The status will change to In Sync and connected after you configure a Collector Group (Step 9).

STEP 9 | Enable the logging disks on each Log Collector.

Use the web interface of the primary Panorama management server peer to perform these steps:

1. Select **Panorama** > **Managed Collectors** and edit the Log Collector.
2. Select **Disks**, **Add** each disk pair, and click **OK**.
3. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

STEP 10 | (Recommended) Configure the **Ethernet1**, **Ethernet2**, **Ethernet3**, **Ethernet4**, and **Ethernet5** interfaces if the Log Collector will use them for **Device Log Collection** (receiving logs from firewalls) and **Collector Group Communication**.

By default, the Log Collector uses the MGT interface for log collection and Collector Group communication. Assigning other interfaces to these functions enables you to reserve the MGT interface for management traffic. In an environment with heavy log traffic, consider using the 10Gbps interfaces (**Ethernet4** and **Ethernet5**) on the M-500 appliance for log collection and Collector Group communication. To load balance the logging traffic across interfaces, you can enable **Device Log Collection** on multiple interfaces.

Use the web interface of the primary Panorama management server peer to perform these steps for each Log Collector:

1. Select **Panorama** > **Managed Collectors**, edit the Log Collector, and select **Interfaces**.
2. Perform the following steps for each interface:
   1. Click the name of the interface to edit it.
   2. Select **<interface-name>** to enable the interface.
   3. Complete one or both of the following field sets based on the IP protocols of your network:

      IPv4—**IP Address**, **Netmask**, and **Default Gateway**

      IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
   4. Select the Device Management Services that the interface supports:

      Device Log Collection—You can assign one or more interfaces.

      Collector Group Communication—You can assign only one interface.
   5. Click **OK** to save your changes to the interface.
3. Click **OK** to save your changes to the Log Collector.
4. Select **Commit** > **Commit to Panorama** and **Commit** your changes to the Panorama configuration.

STEP 11 | Add a Firewall as a Managed Device.

Use the web interface of the primary Panorama management server peer to perform this task for each firewall that will forward logs to Log Collectors.

 ©2025 Palo Alto Networks, Inc.

STEP 12 | Configure the Collector Group.

If each Collector Group will have one Log Collector, repeat this step for each Collector Group before continuing.

If you will assign all the Log Collectors to one Collector Group, perform this step only once.

Use the web interface of the primary Panorama management server peer to Configure a Collector Group:

1. Select **Panorama** > **Collector Groups** and **Add** the Collector Group.
2. Enter a **Name** to identify the Collector Group.
3. **Add** one or more Log Collectors to the Collector Group Members list.

   In any single Collector Group, all the Log Collectors must run on the same Panorama model: all M-700 appliances, all M-600 appliances, all M-500 appliances, all M-300 appliances, all M-200 appliances, or all Panorama virtual appliances.
4. (Best Practice) **Enable log redundancy across collectors** if you add multiple Log Collectors to a single Collector group. This option requires each Log Collector to have the same number of logging disks.
5. (Optional) Select **Monitoring** and configure the settings if you will use SNMP to monitor Log Collector statistics and traps.
6. Select **Device Log Forwarding** and configure the Log Forwarding Preferences list. This list defines which firewalls forward logs to which Log Collectors. Assign firewalls according to the number of Log Collectors in this Collector Group:
   - **Single**—Assign the firewalls that will forward logs to that Log Collector, as illustrated in Single Dedicated Log Collector Per Collector Group.
   - **Multiple**—Assign each firewall to both Log Collectors for redundancy. When you configure the preferences, make Log Collector 1 the first priority for half the firewalls and make Log Collector 2 the first priority for the other half, as illustrated in Multiple Dedicated Log Collectors Per Collector Group.
7. Click **OK** to save your changes to the Collector Group.
8. Select **Commit** > **Commit and Push** and then **Commit and Push** your changes to Panorama and to the Collector Groups you added.
9. Select **Panorama** > **Managed Collectors** to verify that the Log Collector configuration is synchronized with Panorama.

   The Configuration Status column should display In Sync and the Run Time Status column should display connected.

STEP 13 | Configure log forwarding from firewalls to Panorama.

Use the web interface of the primary Panorama management server peer to:

1. Configure Log Forwarding to Panorama.
2. Verify Log Forwarding to Panorama.
3. (Optional) Configure Log Forwarding from Panorama to External Destinations.

 ©2025 Palo Alto Networks, Inc.

Deploy Panorama M-Series Appliances with Local Log Collectors

The following figures illustrate Panorama in a centralized log collection deployment. In these examples, the Panorama management server comprises two M-Series appliances in Panorama mode that are deployed in an active/passive high availability (HA) configuration. The firewalls send logs to the predefined (default) local Log Collector on each Panorama M-Series appliance.

If you assign more than one Log Collector to a Collector Group, see Caveats for a Collector Group with Multiple Log Collectors to understand the requirements, risks, and recommended mitigations.

After implementing this deployment, if the logging rate exceeds the qualified maximum, Palo Alto Networks recommends that you add Dedicated Log Collectors (M-Series appliances in Log Collector mode) as described in the Deploy Panorama with Dedicated Log Collectors documentation. For more information about the qualified maximum logging rates, see Logging Rates for M-Series appliances and System Requirements for Panorama Virtual Appliances. Such an expansion might require reassigning firewalls from the local Log Collectors to Dedicated Log Collectors.

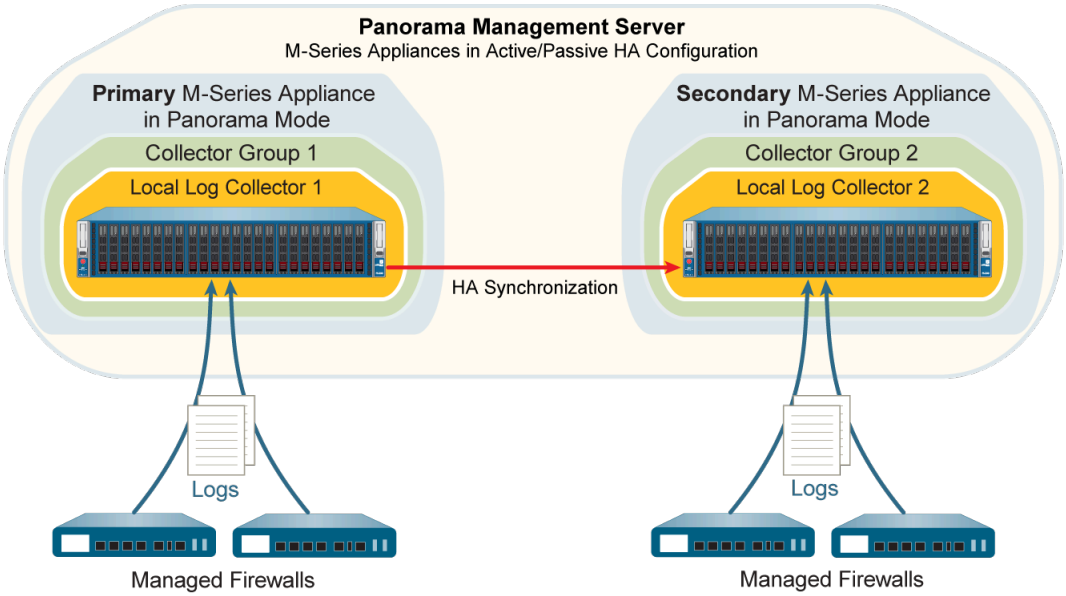


Figure 19: Single Local Log Collector Per Collector Group

©2025 Palo Alto Networks, Inc.

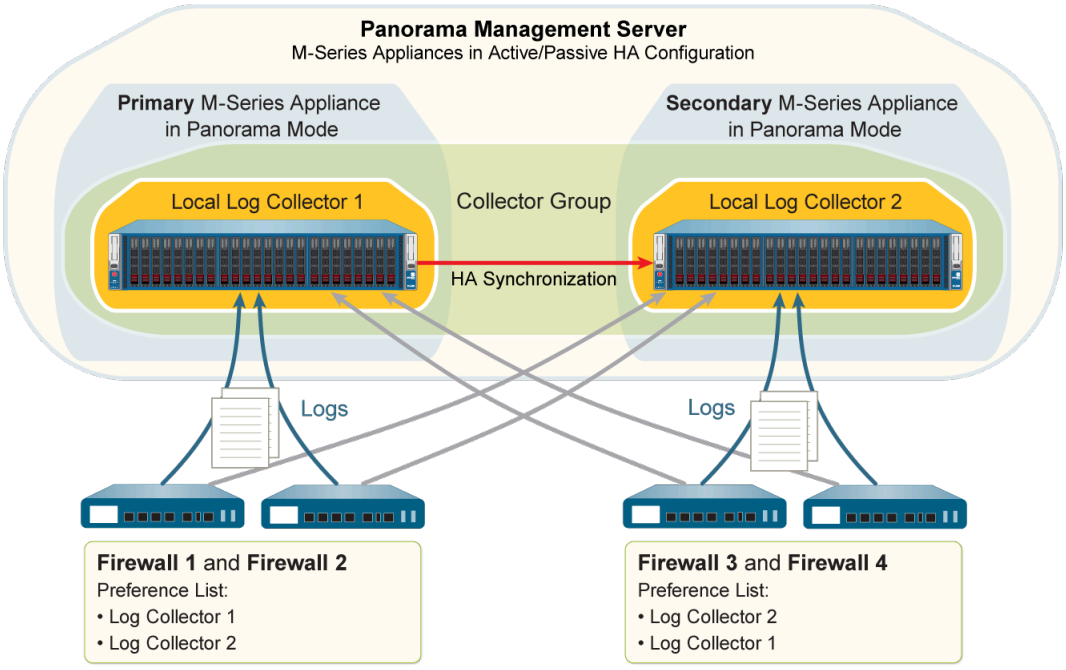


Figure 20: Multiple Local Log Collectors Per Collector Group

Perform the following steps to deploy Panorama with local Log Collectors. Skip any steps you have already performed (for example, the initial setup).

STEP 1 | Perform the initial setup of each M-Series appliance.

1. Rack mount the M-Series appliance. Refer to the M-Series Hardware Reference Guides for instructions.
2. Perform Initial Configuration of the M-Series Appliance.

   Palo Alto Networks recommends reserving the management (MGT) interface for administrative access to Panorama and dedicating separate M-Series Appliance Interfaces to other Panorama services.
3. Configure each array. This task is required to make the RAID disks available for logging. Optionally, you can add disks to Increase Storage on the M-Series Appliance.
4. Register Panorama and Install Licenses.
5. Install Content and Software Updates for Panorama.
6. Set Up HA on Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Perform the following steps to prepare Panorama for log collection.

1. Connect to the primary Panorama in one of the following ways:
   - Attach a serial cable from your computer to the Console port on the primary Panorama. Then use terminal emulation software (9600-8-N-1) to connect.
   - Use terminal emulation software such as PuTTY to open an SSH session to the IP address that you specified for the MGT interface of the primary Panorama during initial configuration.
2. Log in to the CLI when prompted. Use the default admin account and the password that you specified during initial configuration.
3. Enable the primary Panorama to connect to the secondary Panorama by entering the following command, where *<IPaddress2>* represents the MGT interface of the secondary Panorama:

   ```
   > configure
   # set deviceconfig system panorama-server <IPaddress2>
   # commit
   ```
4. Log in to the CLI of the secondary Panorama.
5. Enable the secondary Panorama to connect to the primary Panorama by entering the following command, where *<IPaddress1>* represents the MGT interface of the primary Panorama:

   ```
   > configure
   # set deviceconfig system panorama-server <IPaddress1>
   # commit
   # exit
   ```
6. In the CLI of the secondary Panorama, enter the following command to display the serial number, and then record it:

   ```
   > show system info | match serial
   ```

   You need the serial number to add the Log Collector of the secondary Panorama as a managed collector to the primary Panorama.

STEP 3 | Edit the Log Collector that is local to the primary Panorama.

Use the web interface of the primary Panorama to perform these steps:

1. Select **Panorama** > **Managed Collectors** and select the default (local) Log Collector.
2. Select **Disks** and **Add** each logging disk pair.
3. Click **OK** to save your changes.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the Log Collector that is local to the secondary Panorama.

Panorama treats this Log Collector as remote because it's not local to the primary Panorama. Therefore you must manually add it on the primary Panorama.

Use the web interface of the primary Panorama to Configure a Managed Collector:

1. Select **Panorama** > **Managed Collectors** and **Add** the Log Collector.
2. Enter the serial number (**Collector S/N**) you recorded for the Log Collector of the secondary Panorama.
3. Enter the IP address or FQDN of the primary and secondary Panorama HA peers in the **Panorama Server IP** field and **Panorama Server IP 2** field respectively.

   Both of these fields are required.
4. Select **Interfaces** and configure each interface that the Log Collector will use. The **Management** interface is required. Perform the following steps for each interface:
   1. Click the interface name.
   2. Configure one or both of the following field sets based on the IP protocols of your network.

      IPv4—**IP Address**, **Netmask**, and **Default Gateway**

      IPv6—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
   3. (Management interface only) Select **SNMP** if you will use an SNMP manager to monitor Log Collector statistics.

      Using SNMP requires additional steps besides configuring the Log Collector (see Monitor Panorama and Log Collector Statistics Using SNMP).
   4. Click **OK** to save your changes to the interface.
5. Click **OK** to save your changes to the Log Collector.
6. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

   This step is required before you can enable logging disks.
7. Edit the Log Collector by clicking its name.
8. Select **Disks**, **Add** each RAID disk pair, and click **OK**.
9. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 5 | Add a Firewall as a Managed Device.

Use the web interface of the primary Panorama to perform this task for each firewall that will forward logs to the Log Collectors.

STEP 6 | Edit the default Collector Group that is predefined on the primary Panorama.

Use the web interface of the primary Panorama to Configure a Collector Group:

1. Select **Panorama** > **Collector Groups** and edit the **default** Collector Group.
2. **Add** the local Log Collector of the secondary Panorama to the Collector Group Members list if you are adding multiple Log Collectors to a single Collector group. By default, the

 ©2025 Palo Alto Networks, Inc.

list displays the local Log Collector of the primary Panorama because it is pre-assigned to the default Collector Group.

> *In any single Collector Group, all the Log Collectors must run on the same Panorama model: all M-700 appliances, all M-600 appliances, all M-500 appliances, all M-300 appliances, all M-200 appliances, or all Panorama virtual appliances.*

3. (Best Practice) **Enable log redundancy across collectors** if you add multiple Log Collectors to a single Collector group. This option requires each Log Collector to have the same number of logging disks.
4. (Optional) Select **Monitoring** and configure the settings if you will use SNMP to monitor Log Collector statistics and traps.
5. Select **Device Log Forwarding** and configure the Log Forwarding Preferences list. This list defines which firewalls forward logs to which Log Collectors. Assign firewalls according to the number of Log Collectors in this Collector Group:
   - **Single**—Assign the firewalls that will forward logs to the local Log Collector of the primary Panorama, as illustrated in Single Local Log Collector Per Collector Group.
   - **Multiple**—Assign each firewall to both Log Collectors for redundancy. When you configure the preferences, make Log Collector 1 the first priority for half the firewalls and make Log Collector 2 the first priority for the other half, as illustrated in Multiple Local Log Collectors Per Collector Group.
6. Click **OK** to save your changes.

STEP 7 | Configure a Collector Group that contains the Log Collector of the secondary Panorama.

Required if each Collector Group has only one Log Collector.

Use the web interface of the primary Panorama to Configure a Collector Group:

1. Select **Panorama** > **Collector Groups** and **Add** the Collector Group.
2. Enter a **Name** to identify the Collector Group.
3. **Add** the local Log Collector of the secondary Panorama to the Collector Group Members list.
4. (Optional) Select **Monitoring** and configure the settings if you will use an SNMP manager to monitor Log Collector statistics and traps.
5. Select **Device Log Forwarding** and **Add** an entry to the Log Forwarding Preferences list:
   1. **Modify** the Devices list, select the firewalls that will forward logs to the local Log Collector of the secondary Panorama (see Single Local Log Collector Per Collector Group), and click **OK**.
   2. **Add** the local Log Collector of the secondary Panorama to the Collectors list and click **OK**.
6. Click **OK** to save your changes.

STEP 8 | Commit and push your changes to the Panorama configuration and the Collector Groups.

In the web interface of the primary Panorama, select **Commit** > **Commit and Push** and then **Commit and Push** your changes to Panorama and the Collector Groups you added.

 ©2025 Palo Alto Networks, Inc.

STEP 9 | Manually fail over so that the secondary Panorama becomes active.

Use the web interface of the primary Panorama to perform the following steps:

1. Select **Panorama** > **High Availability**.
2. Click **Suspend local Panorama** in the Operational Commands section.

STEP 10 | On the secondary Panorama, configure the network settings of the Log Collector that is local to the primary Panorama.

Use the web interface of the secondary Panorama to perform the following steps:

1. In the Panorama web interface, select **Panorama** > **Managed Collectors** and select the Log Collector that is local to the primary Panorama.
2. Enter the IP address or FQDN of the primary and secondary Panorama HA peers in the **Panorama Server IP** field and **Panorama Server IP 2** field respectively.

   Both of these fields are required.
3. Select **Interfaces**, click **Management**, and complete one or both of the following field sets (based on the IP protocols of your network) with the MGT interface values of the primary Panorama:
   - **IPv4**—**IP Address**, **Netmask**, and **Default Gateway**
   - **IPv6**—**IPv6 Address/Prefix Length** and **Default IPv6 Gateway**
4. Click **OK** to save your changes.
5. Select **Commit** > **Commit and Push** and then **Commit and Push** your changes to Panorama and the Collector Groups you added.

STEP 11 | Manually fail back so that the primary Panorama becomes active.

Use the web interface of the secondary Panorama to perform the following steps:

1. Select **Panorama** > **High Availability**.
2. Click **Suspend local Panorama** in the Operational Commands section.

STEP 12 | Configure log forwarding from firewalls to Panorama.

Use the web interface of the primary Panorama to:

1. Configure Log Forwarding to Panorama.
2. Verify Log Forwarding to Panorama.
3. (Optional) Configure Log Forwarding from Panorama to External Destinations.



*You can assign separate external server profiles to each Panorama HA peer. For example, you might want each peer to forward logs to a different syslog server. To make each Panorama peer forward logs to different external services, log in to the web interface of each peer, select **Panorama** > **Collector Groups**, select the Collector Group, select **Collector Log Forwarding**, assign the server profiles, and click **OK**.*

 ©2025 Palo Alto Networks, Inc.

Deploy Panorama Virtual Appliances with Local Log Collectors

You can configure firewalls to send logs to a Log Collector that runs locally on a Panorama virtual appliance in Panorama mode. In a high availability (HA) configuration, each Panorama HA peer can have a local Log Collector. You can assign the local Log Collectors on the HA peers to the same Collector Group or to separate Collector Groups, as illustrated in the following figures. Refer to the Setup Prerequisites for the Panorama Virtual Appliance to review the supported logs per second when deploying the Panorama virtual appliance with local Log Collectors in a VMware virtual infrastructure.

If you will assign more than one Log Collector to a Collector Group, see Caveats for a Collector Group with Multiple Log Collectors to understand the requirements, risks, and recommended mitigations.

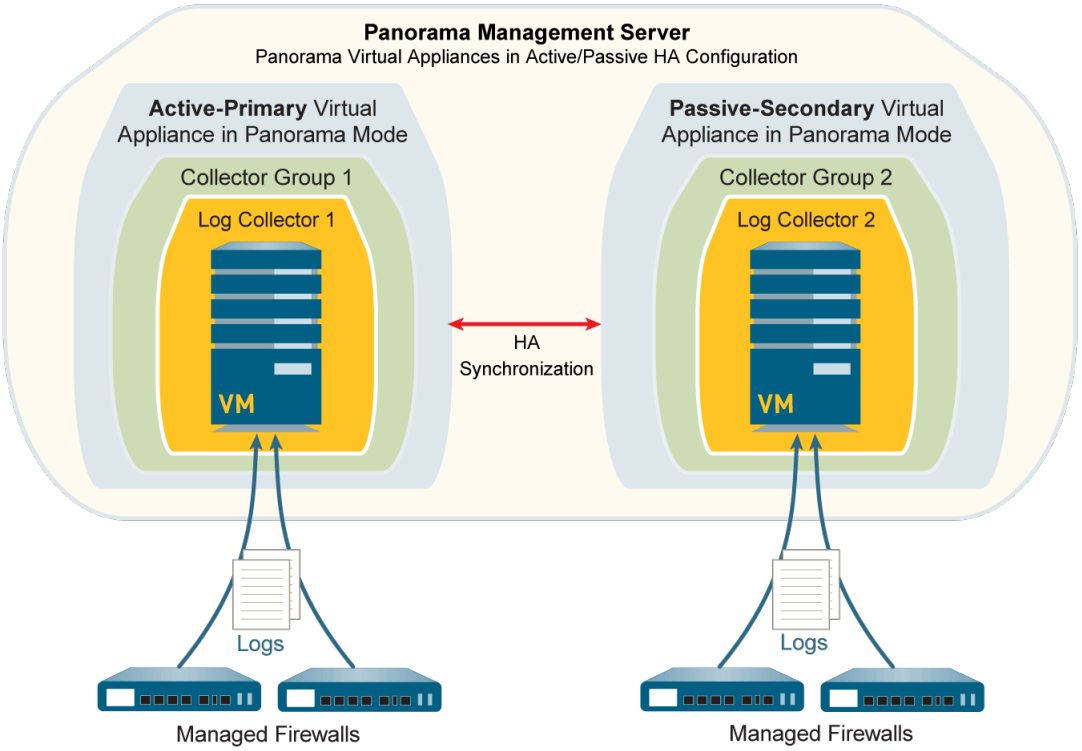


Figure 21: Single Log Collector Per Collector Group

 ©2025 Palo Alto Networks, Inc.

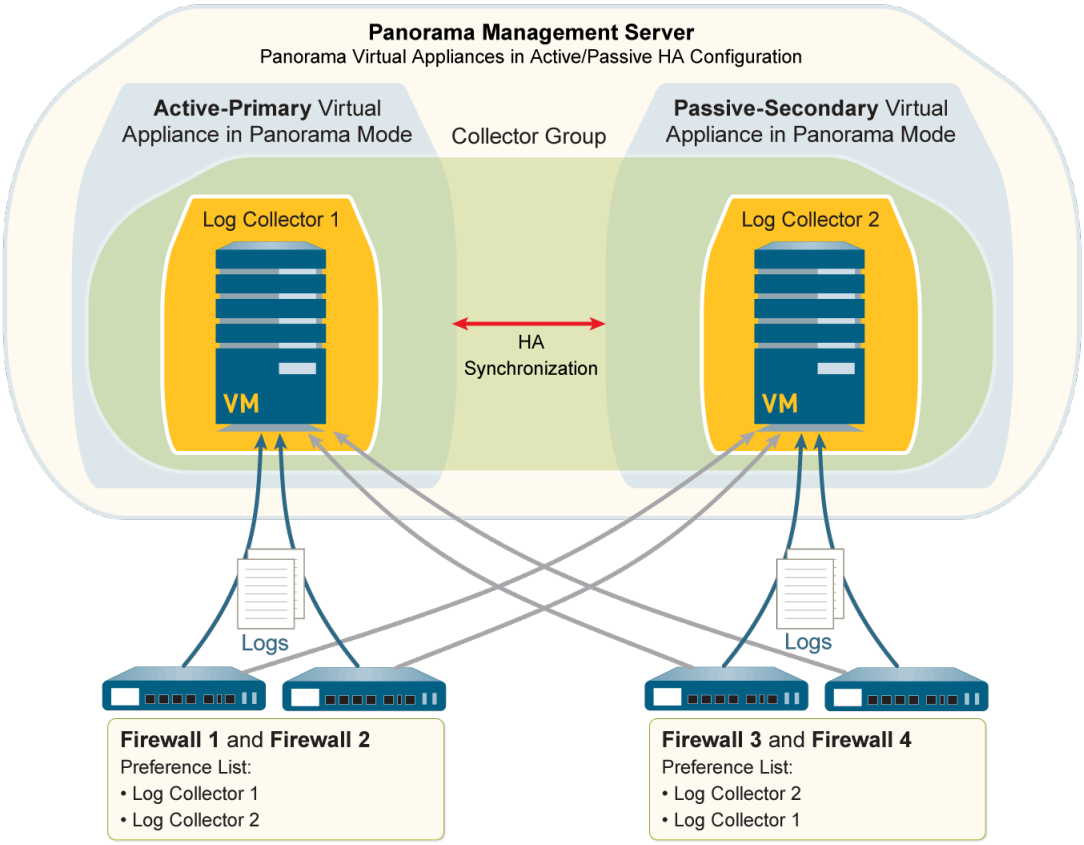


Figure 22: Multiple Log Collectors Per Collector Group

Perform the following steps to deploy Panorama with local Log Collectors. Skip any steps you have already performed (such as the initial setup).

STEP 1 | Perform the initial setup of each Panorama virtual appliance.

1. Install the Panorama Virtual Appliance. You must configure the following resources to ensure the virtual appliance starts in Panorama mode:
   - System disk with exactly 81GB of storage.
   - CPUs and memory that are sufficient for the quantity of logs that Panorama will receive and store.
   - Virtual logging disk with 2–24TB of storage.

   Panorama automatically divides the new disk into 2TB partitions, each of which will function as a separate virtual disk.
2. Perform Initial Configuration of the Panorama Virtual Appliance.
3. Register Panorama and Install Licenses.
4. Install Content and Software Updates for Panorama.

STEP 2 | Set up the Panorama virtual appliances in an HA configuration.

1. Set Up HA on Panorama.
2. Test Panorama HA Failover.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Add a Log Collector that is local to the primary Panorama.

On the primary Panorama:

1. Record the Panorama serial number.
    1. Access the Panorama web interface.
    2. Select **Dashboard** and record the **Serial #** in the General Information section.
2. Add the Log Collector as a managed collector.
    1. Select **Panorama** > **Managed Collectors** and **Add** a new Log Collector.
    2. In the **General** settings, enter the serial number (**Collector S/N**) you recorded for Panorama.
    3. Click **OK** to save your changes.
    4. Select **Commit** > **Commit to Panorama**.

        This step is required before you can add the virtual logging disks.
3. Add the virtual logging disks.
    1. Select **Panorama** > **Managed Collectors** and edit the Log Collector by clicking its name.

        The Log Collector name has the same value as the hostname of the primary Panorama.
    2. Select **Disks** and **Add** the virtual logging disks.
    3. Click **OK** to save your changes.
    4. Select **Commit** > **Commit to Panorama**.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Add a Log Collector that is local to the secondary Panorama.

Panorama treats this Log Collector as remote because it does not run locally on the primary Panorama.

1. Record the serial number of the secondary Panorama.
   1. Access the web interface of the secondary Panorama.
   2. Select **Dashboard** and record the **Serial #** in the General Information section.
2. Access the web interface of the primary Panorama.
3. Select **Panorama** > **Managed Collectors** and **Add** the Log Collector.
4. In the **General** settings, enter the serial number (**Collector S/N**) you recorded for the secondary Panorama.
5. Enter the IP address or FQDN of the primary and secondary Panorama HA peers in the **Panorama Server IP** field and **Panorama Server IP 2** field respectively.

   Both of these fields are required.
6. Click **OK** to save your changes to the Log Collector.
7. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

   This step is required before you can add the virtual logging disks.
8. Edit the Log Collector by clicking its name.

   The Log Collector name has the same value as the hostname of the secondary Panorama.
9. Select **Disks**, **Add** the virtual logging disks, and click **OK**.
10. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 5 | Add a Firewall as a Managed Device.

Use the primary Panorama to perform this task for each firewall that will forward logs to the Log Collectors.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Configure the Collector Group.

Perform this step once if you will assign both Log Collectors to the same Collector Group. Otherwise, configure a Collector Group for each Log Collector.

On the primary Panorama:

1. Select **Panorama** > **Collector Groups** and **Add** a Collector Group.
2. **Add** one or both Log Collectors as Collector Group Members.

   In any single Collector Group, all the Log Collectors must run on the same Panorama model: all M-700 appliances, all M-600 appliances, all M-500 appliances, all M-300 appliances, all M-200 appliances, or all Panorama virtual appliances.
3. (Best Practice) **Enable log redundancy across collectors** if you add multiple Log Collectors to a single Collector group. This option requires each Log Collector to have the same number of virtual logging disks.

   Enabling redundancy doubles the amount of logs and log processing traffic in a Collector Group. If necessary, Expand Log Storage Capacity on the Panorama Virtual Appliance.
4. Select **Device Log Forwarding** and configure the Log Forwarding Preferences list. This list defines which firewalls forward logs to which Log Collectors. Assign firewalls according to the number of Log Collectors in this Collector Group:
   - **Single**—Assign the firewalls that will forward logs to the Log Collector that is local to the primary Panorama, as illustrated in Single Log Collector Per Collector Group.
   - **Multiple**—Assign each firewall to both Log Collectors for redundancy. When you configure the preference list, make Log Collector 1 the first priority for half the firewalls and make Log Collector 2 the first priority for the other half, as illustrated in Multiple Log Collectors Per Collector Group.
5. Click **OK** to save your changes.
6. Select **Commit** > **Commit and Push** and then **Commit and Push** your changes to Panorama and the Collector Groups you added.

STEP 7 | Trigger failover on the primary Panorama so that the secondary Panorama becomes active.

On the primary Panorama:

1. Select **Panorama** > **High Availability**.
2. Click **Suspend local Panorama** in the Operational Commands section.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Configure the connection from the secondary Panorama to the Log Collector that is local to the primary Panorama.

On the secondary Panorama:

1. In the Panorama web interface, select **Panorama** > **Managed Collectors** and select the Log Collector that is local to the primary Panorama.
2. Enter the IP address or FQDN of the primary and secondary Panorama HA peers in the **Panorama Server IP** field and **Panorama Server IP 2** field respectively.

   Both of these fields are required.
3. Click **OK** to save your changes.
4. Select **Commit** > **Commit and Push** and then **Commit and Push** your changes to Panorama and the Collector Groups.

STEP 9 | Restore HA functionality on the primary Panorama.

1. Log in to the Panorama web interface of the primary Panorama.
2. Select **Panorama** > **High Availability**.
3. **Make local Panorama functional for high availability**.

STEP 10 | Trigger fail-back on the secondary Panorama so that the primary Panorama becomes active.

On the secondary Panorama:

1. Select **Panorama** > **High Availability**.
2. Click **Suspend local Panorama** in the Operational Commands section.
3. **Make local Panorama functional for high availability** to restore HA functionality to the secondary Panorama.
4. In the **Dashboard**, verify in the High Availability widget that the secondary Panorama is `secondary-passive`.
5. Log in to the Panorama web interface of the primary Panorama and in the **Dashboard**, verify in the High Availability widget that the primary Panorama is `primary-active`.

STEP 11 | Configure log forwarding from the firewalls to Panorama.

On the primary Panorama to:

1. Configure Log Forwarding to Panorama from firewalls.
2. Verify Log Forwarding to Panorama.

Deploy Panorama Virtual Appliances in Legacy Mode with Local Log Collection

The following figure illustrates Panorama in a centralized log collection deployment. In this example, the Panorama management server comprises two Panorama virtual appliances in Legacy mode that are deployed in an active/passive high availability (HA) configuration. This configuration suits firewall management within a VMware virtual infrastructure in which Panorama processes up to 10,000 logs/second. The firewalls send logs to the NFS datastore (ESXi server only) or virtual disk on the Panorama management server. By default, the active and passive peers both receive logs, though you can Modify Log Forwarding and Buffering Defaults

 ©2025 Palo Alto Networks, Inc.

so that only the active peer does. For the 5200 and 7000 series firewalls, only the active peer receive logs. By default, the Panorama virtual appliance in Legacy mode uses approximately 11GB on its internal disk partition for log storage, though you can Expand Log Storage Capacity on the Panorama Virtual Appliance if necessary.

If the logging rate increases beyond 10,000 logs per second, it is recommended that you Deploy Panorama with Dedicated Log Collectors.

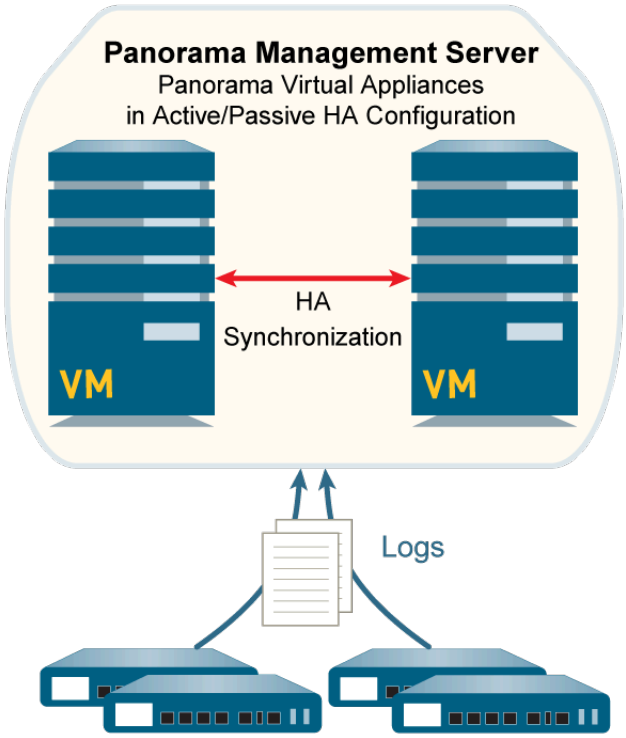


Figure 23: Panorama Virtual Appliances in Legacy Mode with Local Log Collection

Perform the following steps to deploy Panorama virtual appliances with local log collection. Skip any steps you have already performed (for example, the initial setup).

STEP 1 | Perform the initial setup of each Panorama virtual appliance.

1. Install the Panorama Virtual Appliance. To ensure the virtual appliance starts in Panorama mode, do not add a virtual logging disk during installation.

   By default, Panorama uses an 11GB partition on its system disk for log storage. If you want more storage, you can add a dedicated virtual logging disk of up to 8TB after the installation.

2. Perform Initial Configuration of the Panorama Virtual Appliance.
3. Register Panorama and Install Licenses.
4. Install Content and Software Updates for Panorama.

STEP 2 | Set up the Panorama virtual appliances in an HA configuration.

1. Set Up HA on Panorama.
2. Test Panorama HA Failover.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Perform the following steps to prepare Panorama for log collection.

1. Add a Firewall as a Managed Device for each one that will forward logs to Panorama.
2. Configure Log Forwarding to Panorama.

STEP 4 | Commit your changes.

Select **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

Manage WildFire Appliances

You can manage up to 200 standalone WildFire appliances and WildFire appliance cluster nodes centrally using a Panorama M-Series or virtual appliance. Compared to managing WildFire appliances and appliance clusters individually using the local CLI, using Panorama provides centralized management and monitoring of multiple appliances and appliance clusters. Centralized management enables you to push common configurations, configuration updates, and software upgrades to all or a subset of the managed WildFire appliances, which makes it easy to ensure that WildFire appliances and appliance clusters have consistent configurations.

When you use Panorama to manage WildFire appliance clusters, Panorama must run an equal or later version than the WildFire appliances being managed.

- Add Standalone WildFire Appliances to Manage with Panorama
- Configure Basic WildFire Appliance Settings on Panorama
- Set Up Authentication Using Custom Certificates on WildFire Appliances and Clusters
- Remove a WildFire Appliance from Panorama Management
- Manage WildFire Clusters

Add Standalone WildFire Appliances to Manage with Panorama

You can manage up to 200 WildFire® appliances with a Panorama® M-Series or virtual appliance. The WildFire 200-appliance limit is a combined total of standalone appliances and WildFire appliance cluster nodes (if you also Configure a Cluster and Add Nodes on Panorama).

Ensure that your Panorama server is running PAN-OS® 8.1.0 or a later PAN-OS version, and that any WildFire appliance you add to your Panorama management server is also running PAN-OS 8.1.0 or a later release.

A device registration authentication key is used to securely authenticate and connect the Panorama management server and the WildFire appliance on first connect. To configure the device registration authentication key, specify the key lifetime and the number of times you can use the authentication key to onboard new WildFire appliances. Additionally, you can specify one or more WildFire appliance serial numbers for which the authentication key is valid.

The authentication key expires 90 days after the key lifetime expires. After 90 days, you are prompted to re-certify the authentication key to maintain its validity. If you do not re-certify, then the authentication key becomes invalid. A system log is generated each time a WildFire appliance uses the Panorama-generated authentication key. The WildFire appliance uses the authentication key to authenticate Panorama when it delivers the device certificate that is used for all subsequent communications.

(PAN-OS 10.2 only) For WildFire appliances running a PAN-OS 10.1 release, Panorama running PAN-OS 10.2 or later release supports onboarding WildFire appliances running PAN-OS 10.1.3 or later release only. You cannot add a WildFire appliance running PAN-OS 10.1.2 or earlier PAN-OS 10.2 release to Panorama management if Panorama is running PAN-OS 10.2 or later release.

Panorama supports onboarding WildFire appliances running the following releases:

- ***Panorama running PAN-OS 10.2 or later release—** WildFire appliances running PAN-OS 10.1.3 or later release, and WildFire appliances running PAN-OS 10.0 or earlier PAN-OS release.*

There is no impact to WildFire appliances already managed by Panorama on upgrade to PAN-OS 10.2.

STEP 1 | Using the local CLI, verify that each WildFire appliance that you want to manage on a Panorama management server is running PAN-OS 8.1.0 or a later release.

```
admin@qa16> show system info | match version
sw-version: 8.0.1-c45
wf-content-version: 702-283
logdb-version: 8.0.15
```

STEP 2 | On each Panorama appliance you want to use to manage WildFire appliances, verify that the Panorama management server is running PAN-OS 8.1.0 or a later release.

Dashboard > **General Information** > **Software Version** displays the running software version.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | If you aren't sure whether a WildFire appliance belongs to a WildFire appliance cluster or is a standalone appliance on the local WildFire appliance CLI, check the `Node mode` to ensure that the status is `stand_alone` and check the `Applicationstatus` to ensure that the `global-db-service` and `global-queue-service` indicate `ReadyStandalone`.

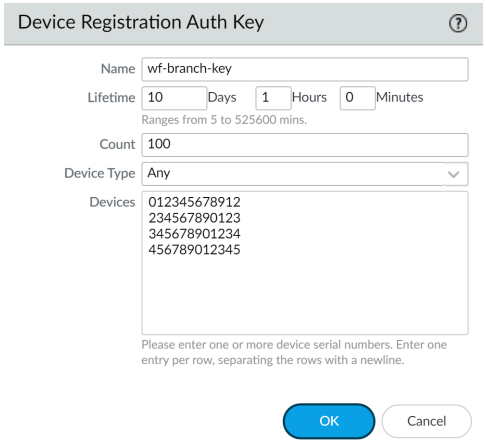
```
admin@WF-500> show cluster membership
Service Summary:  wfpc signature
Cluster name:
Address:          10.10.10.100
Host name:        WF-500
Node name:        wfpc-012345678901-internal
Serial number:    012345678901
Node mode:        stand_alone
Server role:      True
HA priority:
Last changed:     Mon, 06 Mar 2017 16:34:25 -0800
Services:         wfcore signature wfpc infra
Monitor status:
                  Serf Health Status: passing
                      Agent alive and reachable
Application status:
                  global-db-service: ReadyStandalone
                  wildfire-apps-service: Ready
                  global-queue-service: ReadyStandalone
                  wildfire-management-service: Done
                  siggen-db: ReadyMaster
Diag report:
                  10.10.10.100: reported leader '10.10.10.100', age
 0.
                  10.10.10.100: local node passed sanity check.
```

STEP 4 | If the WildFire appliances you want to manage with Panorama are new, check Get Started with WildFire to ensure that you complete basic steps such as confirming your WildFire license is active, enabling logging, connecting firewalls to WildFire appliances, and configuring basic WildFire features.

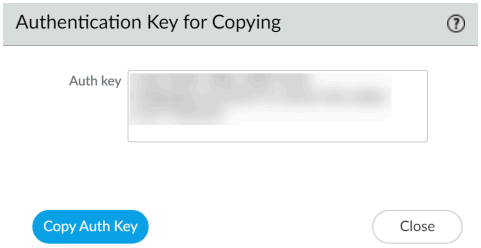
 ©2025 Palo Alto Networks, Inc.

STEP 5 | Create a device registration authentication key.

1. Select **Panorama** > **Device Registration Auth Key** and **Add** a new authentication key.
2. Configure the authentication key.
   - **Name**—Add a descriptive name for the authentication key.
   - **Lifetime**—Specify the key lifetime for how long you can use the authentication key used to onboard new WildFire appliances.
   - **Count**—Specify how many times you can use the authentication key to onboard new WildFire appliances.
   - **Device Type**—Specify that the authentication key is used to authenticate **Any** device.
   - (Optional) **Devices**—Enter one or more device serial numbers to specify for which WildFire appliances the authentication key is valid.
3. Click **OK**.



4. **Copy Auth Key** and **Close**.



STEP 6 | On the local CLI of each WildFire appliance the Panorama server will manage, configure the IP address of the Panorama server and add the device registration authentication key.

Before you register standalone WildFire appliances to a Panorama appliance, you must first configure the Panorama IP address or FQDN and add the device registration authentication key on each WildFire appliance. This enables each WildFire appliance to securely connect

 ©2025 Palo Alto Networks, Inc.

to the Panorama appliance that manages the WildFire appliance. The device registration authentication key is used only for the initial connection to the Panorama server.

1. Configure the IP address or FQDN of the management interface for the primary Panorama server.

   ```
   admin@WF-500# set deviceconfig system panorama-server <ip-address | FQDN>
   ```

2. If you use a backup Panorama appliance for high availability (recommended), configure the IP address or FQDN of the management interface for the backup Panorama server:

   ```
   admin@WF-500# set deviceconfig system panorama-server-2 <ip-address | FQDN>
   ```

3. Add the device registration authentication key.

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

STEP 7 | Register WildFire appliances on the primary Panorama appliance.

1. From the Panorama web interface, **Panorama** > **Managed WildFire Appliances** and **Add Appliance**.
2. Enter the serial number of each WildFire appliance on a separate line. If you do not have a list of serial numbers, on each WildFire appliance, run:

   ```
   admin@WF-500> show system info | match serial
   serial: 012345678901
   ```

   Several local CLI commands display the WildFire appliance serial number, including **`show cluster membership`**.
3. Click **OK**.

   If it is available, information about configuration that is already committed on the WildFire appliances displays, such as IP address and software version.

STEP 8 | (Optional) Import WildFire appliance configurations into the Panorama appliance.

1. Select the appliances that have configurations you want to import from the list of managed WildFire appliances.
2. **Import Config**.
3. Select **Yes**.

   Importing configurations updates the displayed information and makes the imported configurations part of the Panorama appliance candidate configuration.
4. **Commit to Panorama** to make the imported WildFire appliance configurations part of the Panorama running configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 9 | Configure or confirm the configuration of the WildFire appliance interfaces.

Each WildFire appliance has four interfaces: **Management** (Ethernet0), **Analysis Network Environment** (Ethernet1), **Ethernet2**, and **Ethernet3**.

1. Select **Panorama** > **Managed WildFire Appliances** and select a WildFire appliance.
2. Select **Interfaces**.
3. Select an interface to configure or edit it. You can enable the interface, set the speed and duplex, and configure the IP address and netmask, the default gateway, the MTU, the DNS server, the link state, and the **Management Services** for each interface. You can also **Add** permitted IP addresses so that an interface accepts traffic only from specified addresses.

   The **Analysis Network Environment**, **Ethernet2**, and **Ethernet3** interfaces support only **Ping** as a **Management Services** option.

   The **Management** interface supports **Ping**, **SSH**, and **SNMP** as **Management Services** options. In addition, the **Management** interface supports proxy server configuration in case a direct connection to the internet is not possible.
4. Click **OK** save your changes.

STEP 10 | Commit the configuration on the Panorama appliance and push it to the appliance or to multiple appliances.

1. **Commit and Push**.
2. If there are configurations on the Panorama appliance that you do not want to push, **Edit Selections** to choose the appliances to which you want to push configurations. The pushed configuration overwrites the running configuration on a WildFire appliance.

STEP 11 | Verify the configuration.

1. Select **Panorama** > **Managed WildFire Appliances**.
2. Check the following fields:
   - **Connected**—The state is **Connected**.
   - **Role**—The role of each WildFire appliance is **Standalone**.
   - **Config Status**—The status is `InSync`.
   - **Last Commit State**—`Commitsucceeded`.

 ©2025 Palo Alto Networks, Inc.

Configure Basic WildFire Appliance Settings on Panorama

Configuring basic settings such as content update and WildFire cloud servers, WildFire cloud services, logging, authentication, and so on, is similar to how you Configure General Cluster Settings on Panorama. Instead of selecting a cluster and configuring settings on the cluster, select a WildFire appliance and configure the individual settings for that appliance. Select and configure each WildFire appliance that you add to Panorama.

Configure the WildFire Appliance describes how to integrate a WildFire appliance into a network and perform basic setup with the CLI, but the concepts are the same as performing basic setup using Panorama.

Many settings are pre-populated with either defaults, information from previously existing settings on the WildFire appliance, or the settings you configured when adding the WildFire appliance to Panorama.

- Configure Authentication for a WildFire Appliance

Configure Authentication for a WildFire Appliance

Create and configure enhanced authentication for your WildFire appliance by configuring local administrative users with granular authentication parameters, as well as leveraging RADIUS, TACAS+, or LDAP for authorization and authentication.

When you Configure and push administrators from Panorama, you overwrite the existing administrators on the WildFire appliance with those you configure on Panorama.

- Configure An Administrative Account for a WildFire Appliance
- Configure RADIUS Authentication for a WildFire Appliance
- Configure TACACS+ Authentication for a WildFire Appliance
- Configure LDAP Authentication for a WildFire Appliance

Configure An Administrative Account for a WildFire Appliance

Create one or more administrators with granular authentication parameters for your WildFire appliance to manage from the Panorama™ management server. Additionally, you can configure local administrators from Panorama that can be configured directly on the CLI of the WildFire appliance. However, pushing new configuration changes to the WildFire appliance will overwrites local administrators with the administrators configured for the WildFire appliance.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Add Standalone WildFire Appliances to Manage with Panorama.

STEP 3 | (Optional) Configure an authentication profile to define the authentication service that validates the login credentials of the administrators who access the WildFire appliance CLI.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure one or more administrator accounts as needed.

The administrator accounts created on Panorama are later imported to the WildFire appliance and managed from Panorama.

You must configure the administrative account with Superuser admin role privileges to successfully configure authentication for the WildFire appliance.

STEP 5 | Configure the authentication for the WildFire appliance.

1. Select **Panorama** > **Managed WildFire Appliance** and select the WildFire appliance you previously added.
2. (Optional) Select the **Authentication Profile** you configured in the previous step.
3. Configure the authentication **Timeout Configuration** for the WildFire appliance.
   1. Enter the number of **Failed Attempt**s before a user is locked out of the WildFire appliance CLI.
   2. Enter the **Lockout Time**, in minutes, for which the WildFire appliance locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access the WildFire appliance.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
4. Add the WildFire appliance administrators.

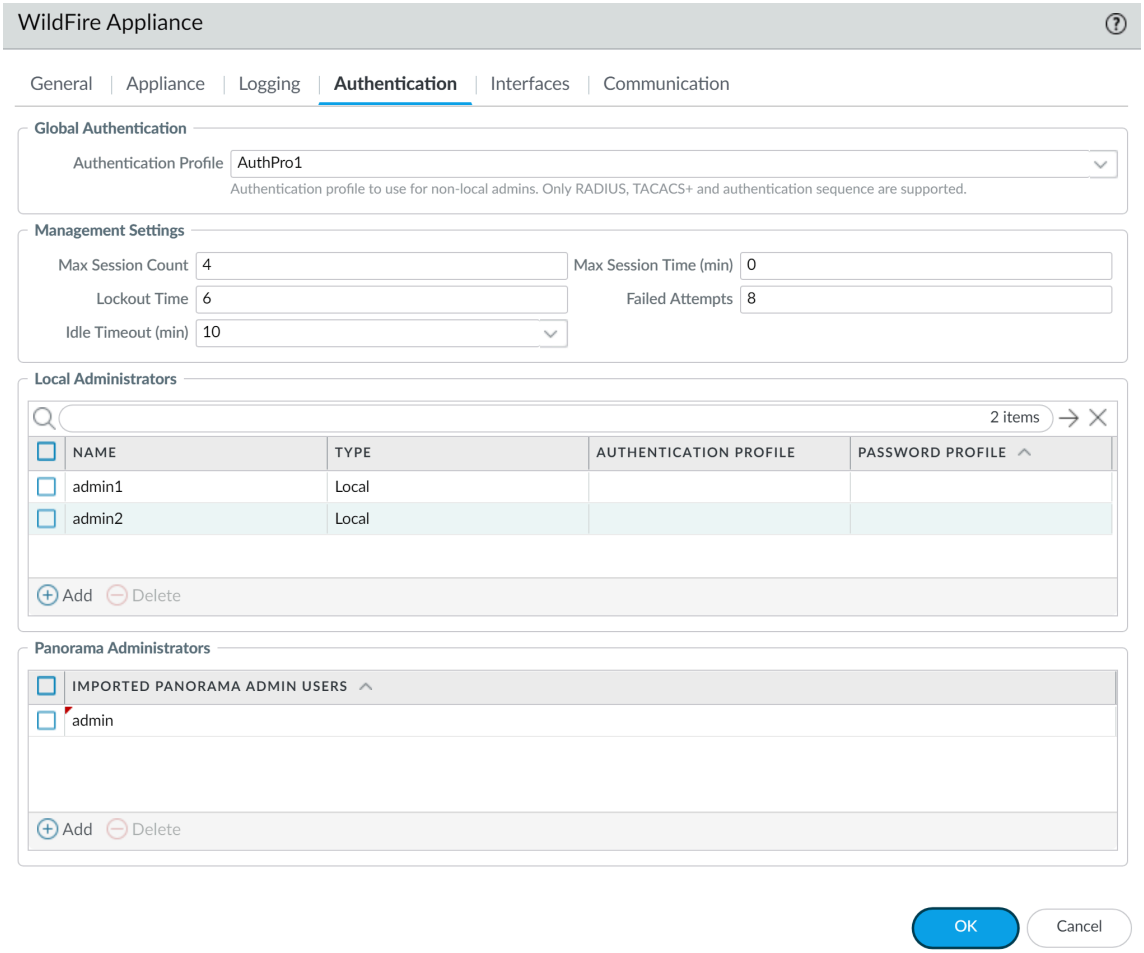
   Administrators may either be added as a local administrator or as an imported Panorama administrator—but not both. Adding the same administrator as both a local administrator and as an imported Panorama administrator is not supported and causes the Panorama

 ©2025 Palo Alto Networks, Inc.

commit to fail. For example, the commit to Panorama fails if you add **admin1** as both a local and Panorama administrator.

1. **Add** and configure new administrators unique to the WildFire appliancer. These administrators are specific to the WildFire appliance for which they are created and you manage these administrators from this table.
2. **Add** any administrators configured on Panorama. These administrators are created on Panorama and imported to the WildFire appliance.

5. Click **OK** to save the WildFire appliance authentication configuration.



STEP 6 | **Commit** and then **Commit and Push** your configuration changes.

STEP 7 | Access the WildFire appliance CLI to verify you can successfully access the WildFire appliance using the local admin user.

Configure RADIUS Authentication for a WildFire Appliance

Use a RADIUS server to authenticate administrative access to the WildFire appliance CLI. You can also define Vendor-Specific Attributes (VSAs) on the RADIUS server to manage administrator authorization. Using VSAs enables you to quickly change the roles, access domains, and user groups of administrators through your directory service, which is often easier than reconfiguring settings on the Panorama™ management server.

 ©2025 Palo Alto Networks, Inc.

You can Import the Palo Alto Networks RADIUS dictionary into RADIUS server to define the authentication attributes needed for communication between Panorama and the RADIUS server.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Add Standalone WildFire Appliances to Manage with Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Configure RADIUS authentication.

Administrator accounts configured for RADIUS authentication are required to have Superuser admin role privileges to successfully configure authentication for the Wildfire appliance.

1. Add a RADIUS server profile.

   The profile defines how the WildFire appliance connects to the RADIUS server.

   1. Select **Panorama** > **Server Profiles** > **RADIUS** and **Add** a profile.
   2. Enter a **Profile Name** to identify the server profile.
   3. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–20).
   4. Select the **Authentication Protocol** (default is **CHAP**) that the WildFire appliance uses to authenticate to the RADIUS server.

      *Select **CHAP** if the RADIUS server supports that protocol; it is more secure than **PAP**.*

   5. **Add** each RADIUS server and enter the following:
      1. **Name** to identify the server.
      2. **RADIUS Server** IP address or FQDN.
      3. **Secret**/**Confirm Secret** (a key to encrypt usernames and passwords).
      4. Server **Port** for authentication requests (default is 1812).
   6. Click **OK** to save the server profile.
2. Assign the RADIUS server profile to an authentication profile.

   The authentication profile defines authentication settings that are common to a set of administrators.

   1. Select **Panorama** > **Authentication Profile** and **Add** a profile.
   2. Enter a **Name** to identify the authentication profile.
   3. Set the **Type** to **RADIUS**.
   4. Select the **Server Profile** you configured.
   5. Select **Retrieve user group from RADIUS** to collect user group information from VSAs defined on the RADIUS server.

      Panorama matches the group information against the groups you specify in the Allow List of the authentication profile.
   6. Select **Advanced** and, in the Allow List, **Add** the administrators that are allowed to authenticate with this authentication profile.
   7. Click **OK** to save the authentication profile.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the authentication for the WildFire appliance.

1. Select **Panorama** > **Managed WildFire Appliance** and select the WildFire appliance you previously added.
2. Select the **Authentication Profile** you configured in the previous step.

   If a global authentication profile is not assigned you must assign an authentication profile to each individual local administrator to leverage remote authentication.
3. Configure the authentication **Timeout Configuration** for the WildFire appliance.
   1. Enter the number of **Failed Attempt**s before a user is locked out of the WildFire appliance CLI.
   2. Enter the **Lockout Time**, in minutes, for which the WildFire appliance locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access the WildFire appliance.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
4. Add the WildFire appliance administrators.

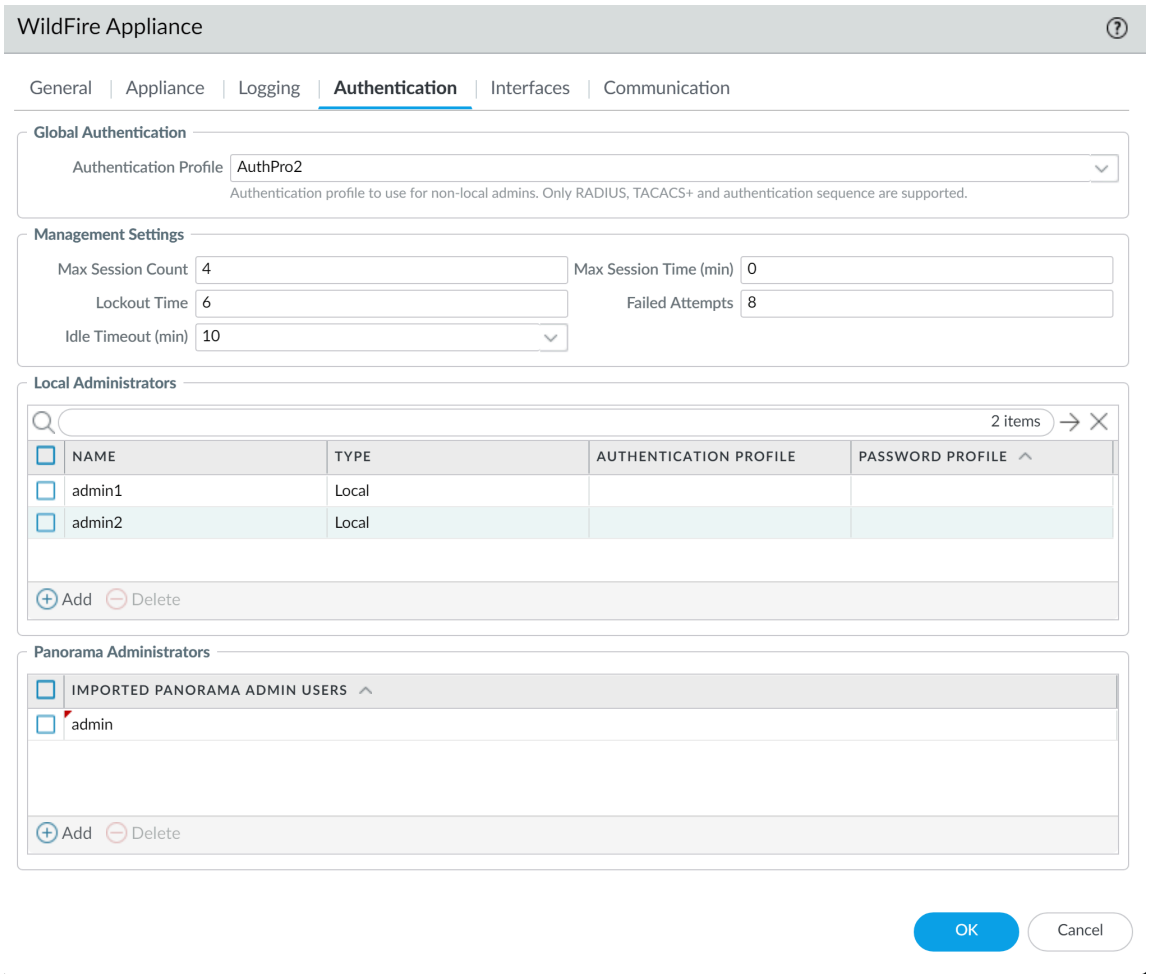
   Administrators may either be added as a local administrator or as an imported Panorama administrator—but not both. Adding the same administrator as both a local administrator and as an imported Panorama administrator is not supported and causes the Panorama

 ©2025 Palo Alto Networks, Inc.

commit to fail. For example, the commit to Panorama fails if you add **admin1** as both a local and Panorama administrator.

1. **Add** and configure new administrators unique to the WildFire appliance. These administrators are specific to the WildFire appliance for which they are created and you manage these administrators from this table.
2. **Add** any administrators configured on Panorama. These administrators are created on Panorama and imported to the WildFire appliance.

5. Click **OK** to save the WildFire appliance authentication configuration.



STEP 5 | **Commit** and then **Commit and Push** your configuration changes.

STEP 6 | Access the WildFire appliance CLI to verify you can successfully access the WildFire appliance using the local admin user.

Configure TACACS+ Authentication for a WildFire Appliance

You can use a TACACS+ server to authenticate administrative access to the WildFire appliance CLI. You can also define Vendor-Specific Attributes (VSAs) on the TACACS+ server to manage administrator authorization. Using VSAs enables you to quickly change the roles, access domains, and user groups of administrators through your directory service, which is often easier than reconfiguring settings on Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Add Standalone WildFire Appliances to Manage with Panorama.

STEP 3 | Configure TACACS+ authentication.

Administrator accounts configured for TACACS+ authentication are required to have Superuser admin role privileges to successfully configure authentication for the Wildfire appliance.

1. Add a TACACS+ server profile.

   The profile defines how the WildFire appliance connects to the TACACS+ server.

   1. Select **Panorama** > **Server Profiles** > **TACACS+** and **Add** a profile.
   2. Enter a **Profile Name** to identify the server profile.
   3. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–20).
   4. Select the **Authentication Protocol** (default is **CHAP**) that Panorama uses to authenticate to the TACACS+ server.
   5. Select **CHAP** if the TACACS+ server supports that protocol; it is more secure than **PAP**.
   6. **Add** each TACACS+ server and enter the following:
      1. **Name** to identify the server.
      2. **TACACS+ Server** IP address or FQDN.
      3. **Secret/Confirm Secret** (a key to encrypt usernames and passwords).
      4. Server **Port** for authentication requests (default is 49).
   7. Click **OK** to save the server profile.
2. Assign the TACACS+ server profile to an authentication profile.

   The authentication profile defines authentication settings that are common to a set of administrators.

   1. Select **Panorama** > **Authentication Profile** and **Add** a profile.
   2. Enter a **Name** to identify the profile.
   3. Set the **Type** to **TACACS+**.
   4. Select the **Server Profile** you configured.
   5. Select **Retrieve user group from TACACS+** to collect user group information from VSAs defined on the TACACS+ server.

      Panorama matches the group information against the groups you specify in the Allow List of the authentication profile.
   6. Select **Advanced** and, in the Allow List, **Add** the administrators that are allowed to authenticate with this authentication profile.
   7. Click **OK** to save the authentication profile.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the authentication for the WildFire appliance.

1. Select **Panorama** > **Managed WildFire Appliance** and select the WildFire appliance you previously added.
2. Select the **Authentication Profile** you configured in the previous step.

   If a global authentication profile is not assigned you must assign an authentication profile to each individual local administrator to leverage remote authentication.
3. Configure the authentication **Timeout Configuration** for the WildFire appliance.
   1. Enter the number of **Failed Attempt**s before a user is locked out of the WildFire appliance CLI.
   2. Enter the **Lockout Time**, in minutes, for which the WildFire appliance locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access the WildFire appliance.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
4. Add the WildFire appliance administrators.

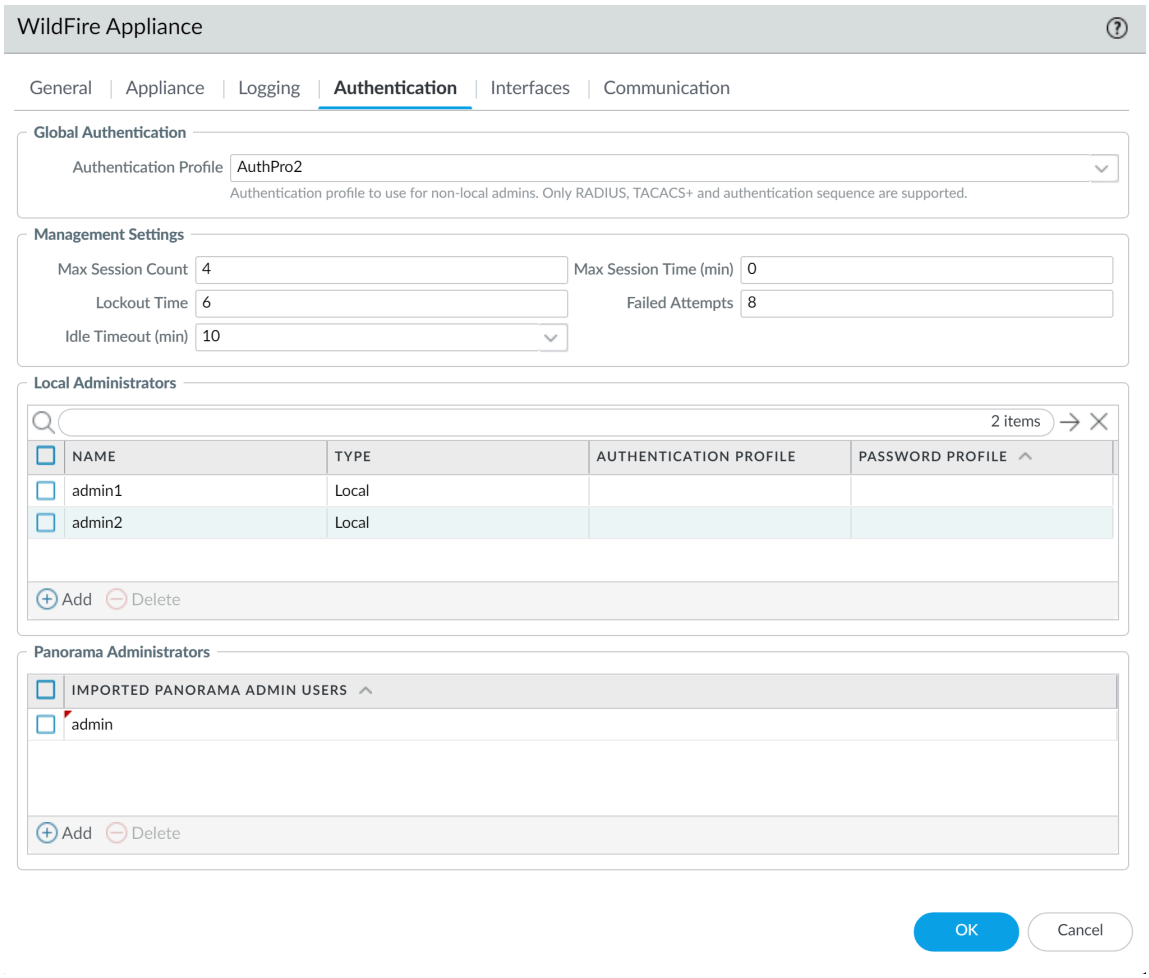
   Administrators may either be added as a local administrator or as an imported Panorama administrator—but not both. Adding the same administrator as both a local administrator and as an imported Panorama administrator is not supported and causes the Panorama

 ©2025 Palo Alto Networks, Inc.

commit to fail. For example, the commit to Panorama fails if you add **admin1** as both a local and Panorama administrator.

1. **Add** and configure new administrators unique to the WildFire appliancer. These administrators are specific to the WildFire appliance for which they are created and you manage these administrators from this table.
2. **Add** any administrators configured on Panorama. These administrators are created on Panorama and imported to the WildFire appliance.

5. Click **OK** to save the WildFire appliance authentication configuration.



STEP 5 | **Commit** and then **Commit and Push** your configuration changes.

STEP 6 | Access the WildFire appliance CLI to verify you can successfully access the WildFire appliance using the local admin user.

Configure LDAP Authentication for a WildFire Appliance

You can use LDAP to authenticate end users who access the WildFire appliance CLI.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Add Standalone WildFire Appliances to Manage with Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Add an LDAP server profile.

The profile defines how the WildFire appliance connects to the LDAP server.

Administrator accounts configured for LDAP authentication are required to have Superuser admin role privileges to successfully configure authentication for the WildFire appliance.

1. Select **Panorama** > **Server Profiles** > **LDAP** and **Add** a server profile.
2. Enter a **Profile Name** to identify the server profile.
3. **Add** the LDAP servers (up to four). For each server, enter a **Name** (to identify the server), **LDAP Server** IP address or FQDN, and server **Port** (default 389).

   If you use an FQDN address object to identify the server and you subsequently change the address, you must commit the change for the new server address to take effect.
4. Select the server **Type**.
5. Select the **Base DN**.

   To identify the Base DN of your directory, open the **Active Directory Domains and Trusts** Microsoft Management Console snap-in and use the name of the top-level domain.
6. Enter the **Bind DN** and **Password** to enable the authentication service to authenticate the firewall.

   The Bind DN account must have permission to read the LDAP directory.
7. Enter the **Bind Timeout** and **Search Timeout** in seconds (default is 30 for both).
8. Enter the **Retry Interval** in seconds (default is 60).
9. (Optional) If you want the endpoint to use SSL or TLS for a more secure connection with the directory server, enable the option to **Require SSL/TLS secured connection** (enabled by default). The protocol that the endpoint uses depends on the server port:
   - 389 (default)—TLS (Specifically, the WildFire appliance uses the StartTLS operation, which upgrades the initial plaintext connection to TLS.)
   - 636—SSL
   - Any other port—The WildFire appliance first attempts to use TLS. If the directory server doesn't support TLS, the WildFire appliance falls back to SSL.
10. (Optional) For additional security, enable to the option to **Verify Server Certificate for SSL sessions** so that the endpoint verifies the certificate that the directory server presents for SSL/TLS connections. To enable verification, you must also enable the

 ©2025 Palo Alto Networks, Inc.

option to **Require SSL/TLS secured connection**. For verification to succeed, the certificate must meet one of the following conditions:

- It is in the list of Panorama certificates: **Panorama** > **Certificate Management** > **Certificates** > **Device Certificates.** If necessary, import the certificate into Panorama.
- The certificate signer is in the list of trusted certificate authorities: **Panorama** > **Certificate Management** > **Certificates**.

11. Click **OK** to save the server profile.

STEP 4 | Configure the authentication for the WildFire appliance.

1. Select **Panorama** > **Managed WildFire Appliance** and select the WildFire appliance you previously added.
2. Configure the authentication **Timeout Configuration** for the WildFire appliance.
   1. Enter the number of **Failed Attempt**s before a user is locked out of the WildFire appliance CLI.
   2. Enter the **Lockout Time**, in minutes, for which the WildFire appliance locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access the WildFire appliance.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
3. Add the WildFire appliance administrators.

   Administrators may either be added as a local administrator or as an imported Panorama administrator—but not both. Adding the same administrator as both a local administrator and as an imported Panorama administrator is not supported and causes the Panorama

 ©2025 Palo Alto Networks, Inc.

commit to fail. For example, the commit to Panorama fails if you add **admin1** as both a local and Panorama administrator.

- Configure the local administrators.

  Configure new administrators unique to the WildFire appliances. These administrators are specific to the WildFire appliance for which they are created and you manage these administrators from this table.

  1. **Add** one or more new local administrator.
  2. Enter a **Name** for the local administrator.
  3. Assign an **Authentication Profile** you previously created.

     LDAP authentication profiles are supported only for individual local administrators.
  4. Enable (check) **Use Public Key Authentication (SSH)** to import a public key file for authentication.
  5. Select a **Password Profile** to set the expiration parameters.

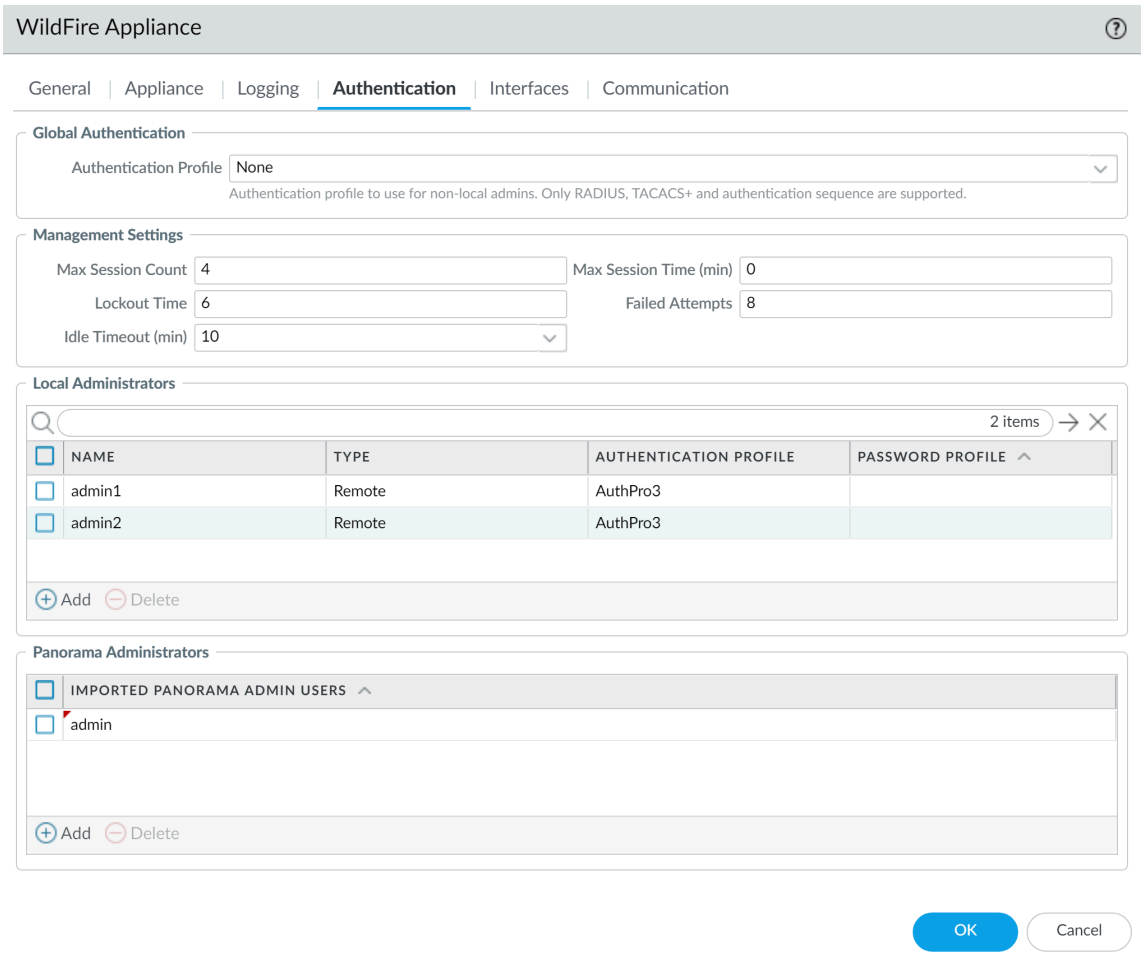
- Import existing Panorama administrators

  Import existing administrators configured on Panorama. These administrators are configured and managed on Panorama and imported to WildFire appliance.

  1. **Add** an existing Panorama administrator

4. Click **OK** to save the WildFire appliance authentication configuration.

 ©2025 Palo Alto Networks, Inc.



STEP 5 | **Commit** and then **Commit and Push** your configuration changes.

STEP 6 | Access the WildFire appliance CLI to verify you can successfully access the WildFire appliance using the local admin user.

 ©2025 Palo Alto Networks, Inc.

Set Up Authentication Using Custom Certificates on WildFire Appliances and Clusters

By default, a WildFire® appliance uses predefined certificates for mutual authentication with other Palo Alto Networks® firewalls and appliances to establish the SSL connections used for management access and inter-device communication. However, you can configure authentication using custom certificates instead. Custom certificates allow you to establish a unique chain of trust to ensure mutual authentication betweenyour WildFire appliance or WildFire cluster managed by Panorama™ and firewalls. You can generate these certificates locally on Panorama or the firewall, obtain them from a trusted third-party certificate authority (CA), or obtain certificates from enterprise private key infrastructure (PKI).

For more information about using custom certificates, see How Are SSL/TLS Connections Mutually Authenticated?

- Configure a Custom Certificate for a Panorama Managed WildFire Appliance
- Configure Authentication with a Single Custom Certificate for a WildFire Cluster
- Apply Custom Certificates on a WildFire Appliance Configured through Panorama

Configure a Custom Certificate for a Panorama Managed WildFire Appliance

If you use Panorama™ to manage your WildFire® appliance or WildFire cluster, you can configure custom certificate authentication through the Panorama web interface instead of using WildFire appliance CLI. The firewall or Panorama uses this connection to forward samples to WildFire for analysis.

This procedure describes how to install a unique certificate on a single WildFire appliance. If the WildFire appliance is part of a cluster, that device and each cluster member has a unique client certificate. To deploy a single certificate to all WildFire appliances in the cluster, see Configure Authentication with a Single Custom Certificate for a WildFire Cluster.

STEP 1 | Obtain key pairs and certificate authority (CA) certificates for the WildFire appliance and the firewall.

STEP 2 | Import the CA certificate to validate the identity of the firewall and the key pair for the WildFire appliance.

1. Select **Panorama** > **Certificate Management** > **Certificates** > **Import**.
2. Import the CA certificate and the key pair on Panorama.

STEP 3 | Configure a certificate profile that includes the root CA and intermediate CA. This certificate profile defines how the WildFire appliance and the firewalls authenticate mutually.

1. Select **Panorama** > **Certificate Management** > **Certificate Profile**.
2. Configure a certificate profile.

   If you configure an intermediate CA as part of the certificate profile, you must also include the root CA.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure an SSL/TLS profile for the WildFire appliance.

*PAN-OS 8.0 and later releases support only TLS 1.2 and higher so ou must set the max version to **TLS 1.2** or **max**.*

1. Select **Panorama** > **Certificate Management** > **SSL/TLS Service Profile**.
2. Configure an SSL/TLS service profile to define the certificate and protocol that the WildFire appliance and its the firewalls use for SSL/TLS services.

STEP 5 | Configure Secure Server Communication on WildFire.

1. Select **Panorama** > **Managed WildFire Clusters** or **Panorama** > **Managed WildFire Appliances** and select a cluster or appliance.
2. Select **Communication**.
3. Enable the **Customize Secure Server Communication** feature.
4. Select the **SSL/TLS Service Profile**. This SSL/TLS service profile applies to all SSL connection between the WildFire appliance and the firewall or Panorama.
5. Select the **Certificate Profile** you configured for communication between the WildFire appliance and the firewall or Panorama.
6. Verify that **Custom Certificates Only** is disabled (cleared). This allows the WildFire appliance to continue communicating with the firewalls with the predefined certificate while migrating to custom certificates.
7. (Optional) Configure an authorization list.
   1. **Add** an Authorization List.
   2. Select the **Subject** or **Subject Alt Name** configured in the certificate profile as the Identifier type.
   3. Enter the Common Name if the identifier is Subject or enter an IP address, hostname, or email if the identifier is Subject Alt Name.
   4. Click **OK**.
   5. Enable **Check Authorization List** to enforce the list.
8. Click **OK**.
9. **Commit** your changes.

STEP 6 | Import the CA certificate to validate the certificate for the WildFire appliance.

1. Log in to the firewall web interface.
2. Import the CA certificate.

STEP 7 | Configure a local or SCEP certificate for the firewall.

- If you are using a local certificate, import the key pair for the firewall.
- If you are using SCEP for the firewall certificate, configure a SCEP profile.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Configure the certificate profile for the firewall or Panorama. You can configure this profile on each client firewall or Panorama appliance individually or you can use a template to push the configuration from Panorama to managed firewalls.

1. Select **Device** > **Certificate Management** > **Certificate Profile** for firewalls or **Panorama** > **Certificate Management** > **Certificate Profile** for Panorama.
2. Configure a Certificate Profile.

STEP 9 | Deploy custom certificates on each firewall or Panorama appliance.

1. Log in to the firewall web interface.
2. Select **Device** > **Setup** > **Management** for a firewall or **Panorama** > **Setup** > **Management** for Panorama and **Edit** the Secure Communication Settings.
3. Select the **Certificate Type**, **Certificate**, and **Certificate Profile**.
4. In the Customize Communication settings, select **WildFire Communication**.
5. Click **OK**.
6. **Commit** your changes.

STEP 10 | After deploying custom certificates on all managed devices, enforce custom-certificate authentication.

1. Log in to Panorama.
2. Select **Panorama** > **Managed WildFire Clusters** or **Panorama** > **Managed WildFire Appliances** and select a cluster or appliance.
3. Select **Communication**.
4. Select **Custom Certificate Only**.
5. Click **OK**.
6. **Commit** your changes.

After committing this change, WildFire immediately begins the enforcement of custom certificates.

Configure Authentication with a Single Custom Certificate for a WildFire Cluster

Instead of assigning unique certificates to each WildFire® appliance in a cluster, you can assign a single, shared client certificate to the entire WildFire cluster, which, in turn, allows you to push a single certificate to all WildFire appliances in the cluster instead of configuring separate certificates for each cluster member. Because the individual WildFire appliances share a client certificate, you must configure a unique hostname (DNS name) for each WildFire appliance. Then you can add all the hostnames as certificate attributes to the shared certificate or use a one-wildcard string that matches all the custom hostnames on all the WildFire appliances in the cluster.

To configure a single custom certificate for your WildFire cluster to use when communicating with the Panorama™, complete the following procedure.

STEP 1 | Obtain a server key pair and CA certificate for Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Configure a certificate profile that includes the root certificate authority (CA) and the intermediate CA. This certificate profile defines the authentication between the WildFire cluster (client) and the Panorama appliance (server).

1. Select **Panorama** > **Certificate Management** > **Certificate Profile**.
2. Configure a certificate profile.

   If you configure an intermediate CA as part of the certificate profile, you must also include the root CA.

STEP 3 | Configure an SSL/TLS service profile.

1. Select **Panorama** > **Certificate Management** > **SSL/TLS Service Profile**.
2. Configure an SSL/TLS service profile to define the certificate and protocol that the WildFire cluster and Panorama appliance use for SSL/TLS services.

STEP 4 | Connect each node in the cluster to Panorama.

STEP 5 | Configure a unique hostname (DNS name) on each node in the cluster or use a string with a single wildcard that matches all custom DNS names set on the WildFire appliances in the cluster.

If using a single-wildcard string, see RFC-6125,Section 6.4.3 for requirements and limitations of wildcard string values. Make sure you understand these requirements and limitations when configuring your custom DNS names.

1. Log in to the WildFire CLI on a node.
2. Use the following command to assign a unique custom DNS name to the node.

```
admin@WF-500> configure
```

```
admin@WF-500# set deviceconfig setting wildfire custom-dns-name <dns-name>
```

3. **Commit** your change.
4. Repeat this process for each node in the cluster.

STEP 6 | On Panorama, generate a client certificate for all nodes in the cluster. Under Certificate Attributes, add a hostname entry for each custom DNS name you assigned to the cluster nodes or add one hostname entry with a one-wildcard string that matches all of the node hostnames, such as *.example.com; you can do this only if each custom DNS name shares a common string.

STEP 7 | On Panorama, configure the certificate profile for the cluster client certificate.

1. Select **Panorama** > **Certificate Management** > **Certificate Profile** for Panorama.
2. Configure a Certificate Profile.

 ©2025 Palo Alto Networks, Inc.

STEP 8 | Deploy custom certificates on each node. This certificate profile must contain the CA certificate that signed the Panorama server certificate.

1. Select **Panorama** > **Managed WildFire Clusters** and click on the cluster name.
2. Select **Communications**.
3. Under Secure Client Communications, select the **Certificate Type**, **Certificate**, and **Certificate Profile**.
4. Click **OK**.
5. **Commit** your changes.

STEP 9 | Configure secure server communication on Panorama.

1. Select **Panorama** > **Setup** > **Management** and **Edit** to select **Customize Secure Server Communication**.
2. Enable **Customize Secure Server Communication**.
3. Select the **SSL/TLS Service Profile**. This SSL/TLS service profile applies to all SSL connection between WildFire and Panorama.
4. Select the **Certificate Profile** for Panorama.
5. Enable **Custom Certificates Only**.
6. Click **OK**.
7. **Commit** your changes.

Apply Custom Certificates on a WildFire Appliance Configured through Panorama

By default, Panorama™ uses a predefined certificate when communicating with a WildFire® appliance to push configurations. You can alternatively configure custom certificates to establish mutual authentication for the connection Panorama uses to push configurations to a managed WildFire appliance or cluster. Complete the following procedure to configure the server certificate on Panorama and the client certificate on the WildFire appliance.

STEP 1 | Obtain key pairs and certificate authority (CA) certificates for Panorama and the WildFire appliance.

STEP 2 | Import the CA certificate to validate the identify of the WildFire appliance and the key pair for Panorama.

1. Select **Panorama** > **Certificate Management** > **Certificates** > **Import**.
2. Import the CA certificate and the key pair on Panorama.

STEP 3 | Configure a certificate profile that includes the root CA and intermediate CA. This certificate profile defines the authentication between the WildFire appliance (client) and the Panorama virtual or M-Series appliance (server).

1. Select **Panorama** > **Certificate Management** > **Certificate Profile**.
2. Configure a certificate profile.

   If you configure an intermediate CA as part of the certificate profile, you must also include the root CA.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure an SSL/TLS service profile.

1. Select **Panorama** > **Certificate Management** > **SSL/TLS Service Profile**.
2. Configure an SSL/TLS service profile to define the certificate and protocol that the WildFire and Panorama appliances use for SSL/TLS services.

STEP 5 | Configure secure server communication on the Panorama appliance.

1. Select **Panorama** > **Setup** > **Management** and **Edit** to select **Customize Secure Server Communication**.
2. Enable the **Customize Secure Server Communication** feature.
3. Select the **SSL/TLS Service Profile**.
4. Select the certificate profile from the **Certificate Profile** drop-down.
5. Verify that **Custom Certificates Only** is disabled (cleared). This allows Panorama to continue communicating with WildFire with the predefined certificate while migrating to custom certificates.
6. (Optional) Configure an authorization list.
   1. **Add** an Authorization List.
   2. Select the **Subject** or **Subject Alt Name** configured in the certificate profile as the Identifier type.
   3. Enter the **Common Name** if the identifier is `Subject` or an **IP address**, **hostname**, or **email** if the identifier is `Subject Alt Name`.
   4. Click **OK**.
   5. Enable the **Check Authorization List** option to configure Panorama to enforce the authorization list.
7. Click **OK**.
8. **Commit** your changes.

STEP 6 | Import the CA certificate to validate the certificate on Panorama.

1. Log in to the Panorama user interface.
2. Import the CA certificate.

STEP 7 | Configure a local or a SCEP certificate for the WildFire appliance.

1. If you are using a local certificate, import the key pair for the WF-500 appliance.
2. If you are using SCEP for the WildFire appliance certificate, configure a SCEP profile.

STEP 8 | Configure the certificate profile for the WildFire appliance.

1. Select **Panorama** > **Certificate Management** > **Certificate Profile**.
2. Configure a certificate profile.

 ©2025 Palo Alto Networks, Inc.

STEP 9 | Deploy custom certificates on each managed WildFire appliance.

1. Log in to Panorama.
2. Select **Panorama** > **Managed WildFire Appliances** and click on a cluster or appliance name.
3. Select **Communications**.
4. Under Secure Client Communications, select the **Certificate Type**, **Certificate**, and **Certificate Profile** from the respective drop-downs.
5. Click **OK**.
6. **Commit** your changes.

STEP 10 | After deploying custom certificates on all managed WildFire appliances, enforce custom-certificate authentication.

1. Select **Panorama** > **Setup** > **Management** and **Edit** the Secure Communications Settings.
2. **Allow Custom Certificate Only**.
3. Click **OK**.
4. **Commit** your changes.

After committing this change, the disconnect wait time begins counting down. When the wait time ends, Panorama and its managed WildFire appliances cannot connect without the configured certificates.

 ©2025 Palo Alto Networks, Inc.

Remove a WildFire Appliance from Panorama Management

You can remove WildFire standalone appliances from Panorama management. When you remove a standalone WildFire appliance from Panorama management, you no longer enjoy the benefits of centralized management and must manage the appliance using its local CLI and scripts.

STEP 1 | Select **Panorama** > **Managed WildFire Appliances**.

STEP 2 | Select the WildFire appliance or appliances you want to remove from Panorama management by selecting the checkbox next to each appliance or by clicking in an appliance's row.

STEP 3 | **Remove** the selected WildFire appliances from Panorama management.

 ©2025 Palo Alto Networks, Inc.

Manage WildFire Clusters

A WildFire appliance cluster is an interconnected group of WildFire appliances that pool resources to increase sample analysis and storage capacity, support larger groups of firewalls and simplify configuration and management of multiple WildFire appliances. For enhanced security and to maintain confidentiality of transmitted content, you can also encrypt communications between WildFire appliances in a cluster. For more information about WildFire clusters and deployment processes, refer to WildFire Appliance Clusters.

The following tasks can be performed using Panorama to manage your WildFire cluster.

- Configure a Cluster Centrally on Panorama
- View WildFire Cluster Status Using Panorama
- Configure Appliance-to-Appliance Encryption Using Predefined Certificates Centrally on Panorama
- Configure Appliance-to-Appliance Encryption Using Custom Certificates Centrally on Panorama

Configure a Cluster Centrally on Panorama

Before you configure a WildFire appliance cluster on a Panorama M-Series or virtual appliance, have two WildFire appliances available to configure as a high availability controller node pair and any additional WildFire appliances needed to serve as worker nodes to increase the analysis, storage capacity, and resiliency of the cluster.

If the WildFire appliances are new, check Get Started with WildFire to ensure that you complete basic steps such as confirming your WildFire license is active, enabling logging, connecting firewalls to WildFire appliances, and configuring basic WildFire features.

To create WildFire appliance clusters, you must upgrade all of the WildFire appliances that you want to place in a cluster to PAN-OS 8.0.1 or later. If you use Panorama to manage WildFire appliance clusters, Panorama also must run PAN-OS 8.0.1 or later. On each WildFire appliance that you want to add to a cluster, run ***`show system info | match version`*** *on the WildFire appliance CLI to ensure that the appliance is running PAN-OS 8.0.1 or later. On each Panorama appliance you use to manage clusters (or standalone appliances),* ***Dashboard*** *>* ***General Information*** *>* ***Software Version*** *displays the running software version.*

When your WildFire appliances are available, perform the appropriate tasks:

- Configure a Cluster and Add Nodes on Panorama
- Configure General Cluster Settings on Panorama
- Configure Authentication for a WildFire Cluster
- Remove a Cluster from Panorama Management

Removing a node from a cluster using Panorama is not supported. Instead, Remove a Node from a Cluster Locally using the local WildFire CLI.

 ©2025 Palo Alto Networks, Inc.

Configure a Cluster and Add Nodes on Panorama

Before configuring a WildFire appliance cluster from Panorama, you must upgrade Panorama to 8.0.1 or later and upgrade all WildFire appliances you plan to add to the cluster to 8.0.1 or later. All WildFire appliances must run the same version of PAN-OS.

You can manage up to 200 WildFire appliances with a Panorama M-Series or virtual appliance. The 200 WildFire appliance limit is the combined total of standalone appliances and WildFire appliance cluster nodes (if you also Add Standalone WildFire Appliances to Manage with Panorama). Except where noted, configuration takes place on Panorama.



Each WildFire appliance cluster node must have a static IP address in the same subnet and have low-latency connections.

STEP 1 | Using the local CLI, configure the IP address of the Panorama server that will manage the WildFire appliance cluster.

Before you register cluster or standalone WildFire appliances to a Panorama appliance, you must first configure the Panorama IP address or FQDN on each WildFire appliance using the local WildFire CLI. This is how each WildFire appliance knows which Panorama appliance manages it.

1. On each WildFire appliance, configure the IP address or FQDN of the primary Panorama appliance's management interface:

   ```
   admin@WF-500# set deviceconfig system panorama-server <ip-address | FQDN>
   ```

2. On each WildFire appliance, if you use a backup Panorama appliance for high availability (recommended), configure the IP address or FQDN of the backup Panorama appliance's management interface:

   ```
   admin@WF-500# set deviceconfig system panorama-server-2 <ip-address | FQDN>
   ```

3. Commit the configuration on each WildFire appliance:

   ```
   admin@WF-500# commit
   ```

STEP 2 | On the primary Panorama appliance, Register the WildFire appliances.

The newly registered appliances are in standalone mode unless they already belong to a cluster due to local cluster configuration.

1. Select **Panorama** > **Managed WildFire Appliances** and **Add Appliance**.
2. Enter the serial number of each WildFire appliance on a separate line. If you do not have a list of WildFire appliance serial numbers, using the local CLI, run **show system info** on each WildFire appliance to obtain the serial number.
3. Click **OK**.

   If it is available, information about configuration that is already committed on the WildFire appliances displays, such as IP address and software version. WildFire

 ©2025 Palo Alto Networks, Inc.

appliances that already belong to a cluster (for example, because of local cluster configuration) display their cluster information and connection status.

STEP 3 | (Optional) Import WildFire appliance configurations into the Panorama appliance.

Importing configurations saves time because you can reuse or edit the configurations on Panorama and then push them to one or more WildFire appliance clusters or standalone WildFire appliances. If there are no configurations you want to import, skip this step. When you push a configuration from Panorama, the pushed configuration overwrites the local configuration.

1. Select **Panorama** > **Managed WildFire Appliances**, and select the appliances that have configurations you want to import from the list of managed WildFire appliances.
2. **Import Config**.
3. Select **Yes**.

   Importing configurations updates the displayed information and makes the imported configurations part of the Panorama appliance candidate configuration.
4. **Commit to Panorama** to make the imported WildFire appliance configurations part of the Panorama running configuration.

STEP 4 | Create a new WildFire appliance cluster.

1. Select **Managed WildFire Clusters**.

   Appliance > **No Cluster Assigned** displays standalone WildFire appliances (nodes) and indicates how many available nodes are not assigned to a cluster.
2. **Create Cluster**.
3. Enter an alphanumeric cluster **Name** of up to 63 characters in length. The **Name** can contain lower-case characters and numbers, and hyphens and periods if they are not the first or last character. No spaces or other characters are allowed.
4. Click **OK**.

   The new cluster name displays but has no assigned WildFire nodes.

STEP 5 | Add WildFire appliances to the new cluster.

The first WildFire appliance added to the cluster automatically becomes the controller node, and the second WildFire appliance added to the cluster automatically becomes the controller backup node. All subsequent WildFire appliances added to the cluster become worker nodes. Worker nodes use the controller node settings so that the cluster has a consistent configuration.

1. Select the new cluster.
2. Select **Clustering**.
3. **Browse** the list of WildFire appliances that do not belong to clusters.
4. Add (⊕) each WildFire appliance you want to include in the cluster. You can add up to twenty nodes to a cluster. Each WildFire appliance that you add to the cluster is displayed along with its automatically assigned role.
5. Click **OK**.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Configure the **Management**, **Analysis Environment Network**, HA, and cluster management interfaces.

Configure the **Management**, **Analysis Environment Network**, and cluster management interfaces on each cluster member (controller and worker nodes) if they are not already configured. The cluster management interface is a dedicated interface for management and communication within the cluster and is not the same as the Management interface.

Configure the HA interfaces individually on both the controller node and the controller backup node. The HA interfaces connect the primary and backup controller nodes and enable them to remain in sync and ready to respond to a failover.

Cluster nodes need IP addresses for each of the four WildFire appliance interfaces. You cannot configure HA services on worker nodes.

1. Select the new cluster.
2. Select **Clustering**.
3. If the management interface is not configured on a cluster node, select **Interface Name** > **Management** and enter the IP address, netmask, services, and other information for the interface.
4. If the interface for the Analysis Environment Network is not configured on a cluster node, select **Interface Name** > **Analysis Environment Network** and enter the IP address, netmask, services, and other information for the interface.
5. On both the controller node and controller backup node, select the interface to use for the HA control link. You must configure the same interface on both controller nodes for the HA service. For example, on the controller node and then on the controller backup node, select **Ethernet3**.
6. For each controller node, select **Clustering Services** > **HA**. (The **HA** option is not available for worker nodes.) If you also want the ability to ping the interface, select **Management Services** > **Ping**.
7. Click **OK**.
8. (Recommended) Select the interface to use as the backup HA control link between the controller node and the controller backup node. You must use the same interface on both nodes for the HA backup service. For example, on both nodes, select **Management**.

   Select **Clustering Services** > **HA Backup** for both nodes. You can also select **Ping**, **SSH**, and **SNMP** if you want those **Management Services** on the interface.

   *The **Analysis Environment Network** interface cannot be an HA or HA Backup interface or a cluster management interface.*
9. Select the dedicated interface to use for management and communication within the cluster. You must use the same interface on both nodes, for example, **Ethernet2**.
10. Select **Clustering Services** > **Cluster Management** for both nodes. If you also want the ability to ping on the interface, select **Management Services** > **Ping**.

    Worker nodes in the cluster automatically inherit the controller node's settings for the dedicated management and communication interface.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Commit the configuration on the Panorama appliance and push it to the cluster.

1. **Commit and Push**.
2. If there are configurations on the Panorama appliance that you do not want to push, **Edit Selections** to choose the appliances to which you push configurations. The pushed configuration overwrites the running configuration on the cluster nodes so that all cluster nodes run the same configuration.

STEP 8 | Verify the configuration.

1. Select **Panorama** > **Managed WildFire Clusters**.
2. Check the following fields:
   - **Appliance**—Instead of displaying as standalone appliances, the WildFire nodes added to the cluster display under the cluster name.
   - **Cluster Name**—The cluster name displays for each node.
   - **Role**—The appropriate role (**Controller**, **Controller Backup**, or **Worker**) displays for each node.
   - **Config Status**—Status is `InSync`.
   - **Last Commit State**—`Commitsucceeded`.

STEP 9 | Using the local CLI on the primary controller node (not the Panorama web interface), check to ensure that the configurations are synchronized.

If they are not synchronized, manually synchronize the high availability configurations on the controller nodes and commit the configuration.

Even though you can perform most other configuration on Panorama, synchronizing the controller node high availability configurations must be done on the primary controller node's CLI.

1. On the primary controller node, check to ensure that the configurations are synchronized:

   ```
   admin@WF-500(active-controller)> show high-availability all
   ```

   At the end of the output, look for the `ConfigurationSynchronization` output:

   ```
   Configuration Synchronization:
       Enabled: yes
   ```

 ©2025 Palo Alto Networks, Inc.

```
    Running Configuration: synchronized
```

If the running configuration is synchronized, you do not need to manually synchronize the configuration. However, if the configuration is not synchronized, you need to synchronize the configuration manually.

2. If the configuration is not synchronized, on the primary controller node, synchronize the high availability configuration to the remote peer controller node:

```
admin@WF-500(active-controller)> request high-availability
 sync-to-remote running-config
```

If there is a mismatch between the primary controller node's configuration and the configuration on the controller backup node, the configuration on the primary controller node overrides the configuration on the controller backup node.

3. Commit the configuration:

```
admin@WF-500# commit
```

Configure General Cluster Settings on Panorama

Some general settings are optional and some general settings are pre-populated with default values. It's best to at least check these settings to ensure that the cluster configuration matches your needs. General settings include:

- Connecting to the WildFire public cloud and submitting samples to the public cloud.
- Configuring data retention policies.
- Configuring logging.
- Setting the analysis environment (the VM image that best matches your environment) and customizing the analysis environment to best service the types of samples the firewalls submit to WildFire.
- Set IP addresses for the DNS server, NTP server, and more.

STEP 1 | Configure settings for the WildFire appliance cluster nodes.

Many settings are pre-populated with either defaults, information from previously existing settings on the controller node, or the settings you just configured.

1. Select the cluster.
2. Select **Appliance**.
3. Enter new information, keep the pre-populated information from the cluster controller node, or edit the pre-populated information, including:
   - **Domain** name.
   - IP address of the **Primary DNS Server** and the **Secondary DNS Server**.
   - **NTP Server Address** and **Authentication Type** of the **Primary NTP Server** and the **Secondary NTP Server**. The **Authentication Type** options are **None**, **Symmetric Key**, and **AutoKey**.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Configure general cluster settings.

Many settings are pre-populated with either defaults, information from previously existing settings on the controller node, or the settings you just configured.

1. Select the new cluster > **General**.
2. (Optional) **Enable DNS** for the controller node to advertise the service status using DNS protocol. The cluster controller provides DNS services on the management (MGT) interface port.
3. **Register Firewall To** the use the service advertised by the cluster controller(s). Palo Alto Networks recommends adding both controllers as authority servers as this provides the benefit of high-availability. Use the form:

   ```
   wfpc.service.<cluster-name>.<domain>
   ```

   For example, a cluster named *mycluster* in the *paloaltonetworks.com* domain would have the domain name:

   ```
   wfpc.service.mycluster.paloaltonetworks.com
   ```
4. Enter the **Content Update Server** for the cluster. Use the default `updates.paloaltonetworks.com` FQDN to connect to the closest server. **Check Server Identity** to confirm the update server identity by matching the common name (CN) in the certificate with the IP address or FQDN of the server (this is checked by default).
5. (Optional) Enter the public **WildFire Cloud Server** location or use the default `wildfire.paloaltonetworks.com` so that the cluster (or standalone appliance managed by Panorama) can send information to the closest WildFire cloud server. If you leave this field blank and do not connect to a WildFire cloud server, the cluster can't receive signature updates directly from the WildFire public cloud, and can't send samples for analysis or contribute data to the public cloud.
6. If you connect the cluster to the public WildFire cloud, select the cloud services you want to enable:
   - **Send Analysis Data**—Send an XML report about local malware analysis. If you send the actual samples, the cluster doesn't send reports.
   - **Send Malicious Samples**—Send malware samples.
   - **Send Diagnostics**—Send diagnostic data.
   - **Verdict Lookup**—Automatically query the WildFire public cloud for verdicts before performing local analysis to reduce the load on the local WildFire appliance cluster.
7. Select the **Sample Analysis Image** to use, based on the types of samples the cluster will analyze.
8. Configure the amount of time for the cluster to retain **Benign/Grayware** sample data (1-90 day range, 14 day default) and **Malicious** sample data (minimum 1 day, no

 ©2025 Palo Alto Networks, Inc.

maximum (indefinite), default is indefinite). Malicious sample data includes phishing verdicts.

9. (Optional) Select **Preferred Analysis Environment** to allocate more resources to **Executables** or **Documents**, depending on your environment. The **Default** allocation is balanced between **Executables** and **Documents**. The available resource amount depends on the number of WildFire nodes in the cluster.

STEP 3 | Check to ensure that the primary and backup Panorama servers are configured.

If you did not configure a backup Panorama server and want to do so, you can add the backup Panorama server.

1. Select the cluster.
2. Select **Appliance**.
3. Check (or enter) the IP address or FQDN of the primary **Panorama Server** and of the backup **Panorama Server 2** if you are using a high availability configuration for centralized cluster management.

STEP 4 | (Optional) Configure system and configuration log settings for the cluster, including log forwarding.

1. Select the cluster.
2. Select **Logging**.
3. Select **System** or **Configuration** to configure a system or configuration log, respectively. The process for configuring them is similar.
4. **Add** (⊕) and **Name** the log forwarding instance, select the **Filter**, and configure the **Forward Method** (**SNMP**, **Email**, **Syslog**, or **HTTP**).

STEP 5 | Configure administrator authentication.

1. Select the cluster.
2. Select **Authentication**.
3. Select the **Authentication Profile**, either **None** or **radius**. RADIUS is the only supported external authentication method.
4. Set the **Local Authentication** mode for admin users as either **Password** or **Password Hash**, and enter the **Password**.

STEP 6 | Commit the configuration on the Panorama appliance and push it to the cluster.

1. **Commit and Push**.
2. If there are configurations on the Panorama appliance that you do not want to push, **Edit Selections** to choose the appliances to which you push configurations. The pushed configuration overwrites the running configuration on the cluster nodes so that all cluster nodes run the same configuration.

Configure Authentication for a WildFire Cluster

Create and configure enhanced authentication for all WildFire appliances in a WildFire cluster by configuring local administrative users with granular authentication parameters, as well as leveraging RADIUS, TACAS+, or LDAP for authorization and authentication.

 ©2025 Palo Alto Networks, Inc.

When you Configure and push administrators from Panorama, you overwrite the existing administrators for all WildFire appliances in the WildFire cluster with those you configure on Panorama.

- Configure an Administrative Account for a WildFire Cluster
- Configure RADIUS Authentication for a WildFire Cluster
- Configure TACACS+ Authentication for a WildFire Cluster
- Configure LDAP Authentication for a WildFire Cluster

Configure an Administrative Account for a WildFire Cluster

Create one or more administrators with granular authentication parameters for all WildFire appliances in a WildFire cluster to manage from the Panorama™ management server. Additionally, you can configure local administrators from Panorama that can be configured directly on the CLI of the WildFire appliance. However, pushing new configuration changes to the WildFire appliance will overwrites local administrators with the administrators configured for the WildFire appliance.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Configure a Cluster Centrally on Panorama.

STEP 3 | (Optional) Configure an authentication profile to define the authentication service that validates the login credentials of the administrators who access the WildFire appliance CLI.

STEP 4 | Configure one or more administrator accounts as needed.

The administrator accounts created on Panorama are later imported to the WildFire appliances in the WildFire Cluster and managed from Panorama.

You must configure the administrative account with Superuser admin role privileges to successfully configure authentication for Wildfire appliances in the WildFire cluster.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Configure the authentication for the WildFire appliances in the WildFire cluster.

1. Select **Panorama** > **Managed WildFire Clusters** and select the WildFire cluster you previously configured.
2. (Optional) Select the **Authentication Profile** you configured in the previous step.
3. Configure the authentication **Timeout Configuration** for the WildFire appliances.
   1. Enter the number of **Failed Attempt**s before a user is locked out of the WildFire appliance CLI.
   2. Enter the **Lockout Time**, in minutes, for which a WildFire appliance locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access a WildFire appliance.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
4. Add the WildFire appliance administrators.

   Administrators may either be added as a local administrator or as an imported Panorama administrator—but not both. Adding the same administrator as both a local administrator and as an imported Panorama administrator is not supported and causes the Panorama commit to fail. For example, the commit to Panorama fails if you add **admin1** as both a local and Panorama administrator.

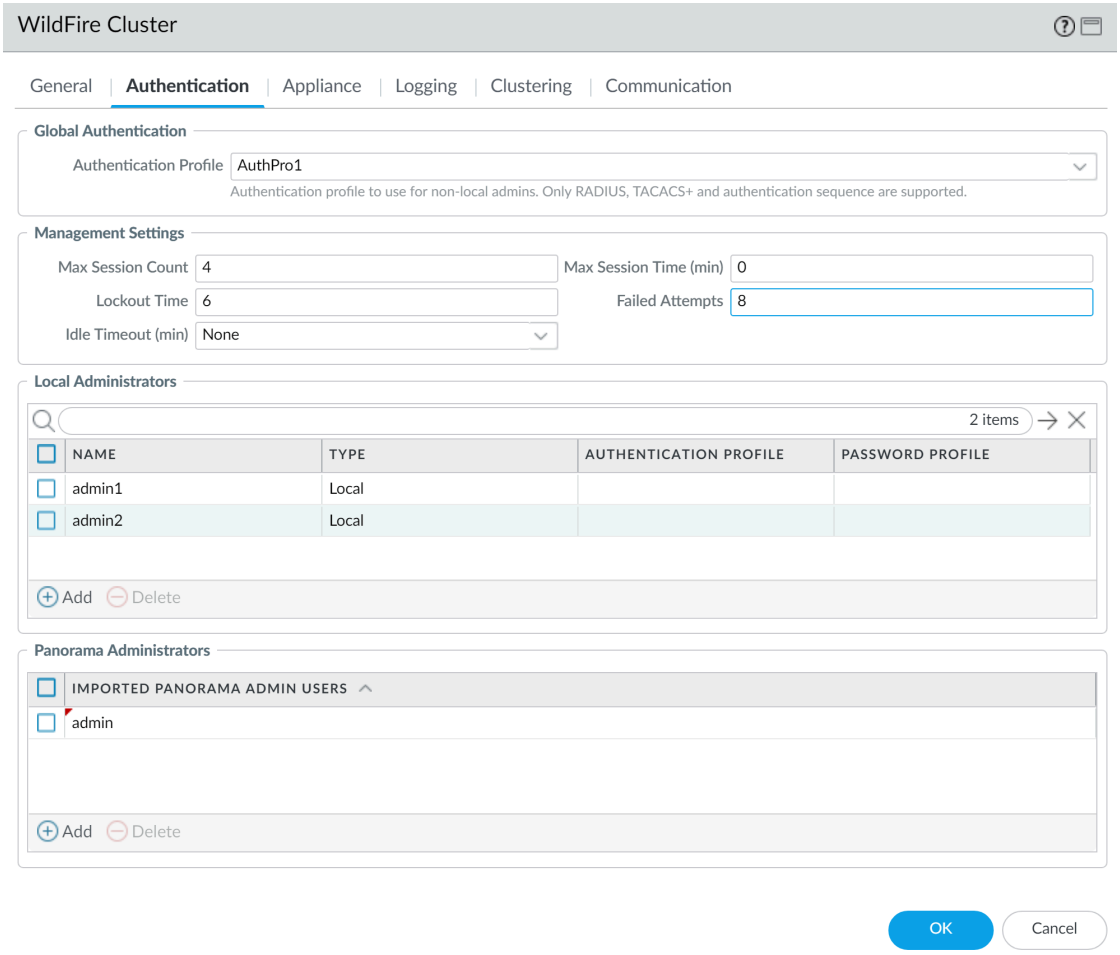
   1. **Add** and configure new administrators unique to the WildFire appliances in the WildFire cluster. These administrators are specific to the WildFire appliances in the

 ©2025 Palo Alto Networks, Inc.

WildFire cluster for which they are created and you manage these administrators from this table.

2. **Add** any administrators configured on Panorama. These administrators are created on Panorama and imported to the WildFire appliances in the WildFire cluster.

5. Click **OK** to save the WildFire cluster authentication configuration.



STEP 6 | **Commit** and then **Commit and Push** your configuration changes.

STEP 7 | Access the WildFire appliance CLI to verify you can successfully access a WildFire appliance using the local admin user.

Configure RADIUS Authentication for a WildFire Cluster

Use a RADIUS server to authenticate administrative access to the CLI of the WildFire appliances in a WildFire cluster. You can also define Vendor-Specific Attributes (VSAs) on the RADIUS server to manage administrator authorization. Using VSAs enables you to quickly change the roles, access domains, and user groups of administrators through your directory service, which is often easier than reconfiguring settings on the Panorama™ management server.

You can Import the Palo Alto Networks RADIUS dictionary into RADIUS server to define the authentication attributes needed for communication between Panorama and the RADIUS server.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Configure a Cluster Centrally on Panorama.

STEP 3 | Configure RADIUS authentication.

Administrator accounts configured for RADIUS authentication are required to have Superuser admin role privileges to successfully configure authentication for Wildfire appliances in the WildFire cluster.

1. Add a RADIUS server profile.

   The profile defines how the WildFire appliances in the WildFire clust connect to the RADIUS server.

   1. Select **Panorama** > **Server Profiles** > **RADIUS** and **Add** a profile.
   2. Enter a **Profile Name** to identify the server profile.
   3. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–20).
   4. Select the **Authentication Protocol** (default is **CHAP**) that a WildFire appliance uses to authenticate to the RADIUS server.

      *Select **CHAP** if the RADIUS server supports that protocol; it is more secure than **PAP**.*

   5. **Add** each RADIUS server and enter the following:
      1. **Name** to identify the server.
      2. **RADIUS Server** IP address or FQDN.
      3. **Secret**/**Confirm Secret** (a key to encrypt usernames and passwords).
      4. Server **Port** for authentication requests (default is 1812).
   6. Click **OK** to save the server profile.
2. Assign the RADIUS server profile to an authentication profile.

   The authentication profile defines authentication settings that are common to a set of administrators.

   1. Select **Panorama** > **Authentication Profile** and **Add** a profile.
   2. Enter a **Name** to identify the authentication profile.
   3. Set the **Type** to **RADIUS**.
   4. Select the **Server Profile** you configured.
   5. Select **Retrieve user group from RADIUS** to collect user group information from VSAs defined on the RADIUS server.

      Panorama matches the group information against the groups you specify in the Allow List of the authentication profile.
   6. Select **Advanced** and, in the Allow List, **Add** the administrators that are allowed to authenticate with this authentication profile.
   7. Click **OK** to save the authentication profile.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the authentication for the WildFire cluster.

1. Select **Panorama** > **Managed WildFire Clusters** and select the WildFire cluster you previously added.
2. Select the **Authentication Profile** you configured in the previous step.

   If a global authentication profile is not assigned you must assign an authentication profile to each individual local administrator to leverage remote authentication.
3. Configure the authentication **Timeout Configuration** for a WildFire appliance.
   1. Enter the number of **Failed Attempt**s before a user is locked out of a WildFire appliance CLI.
   2. Enter the **Lockout Time**, in minutes, for which a WildFire appliance locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access a WildFire appliance.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
4. Add the WildFire appliance administrators.

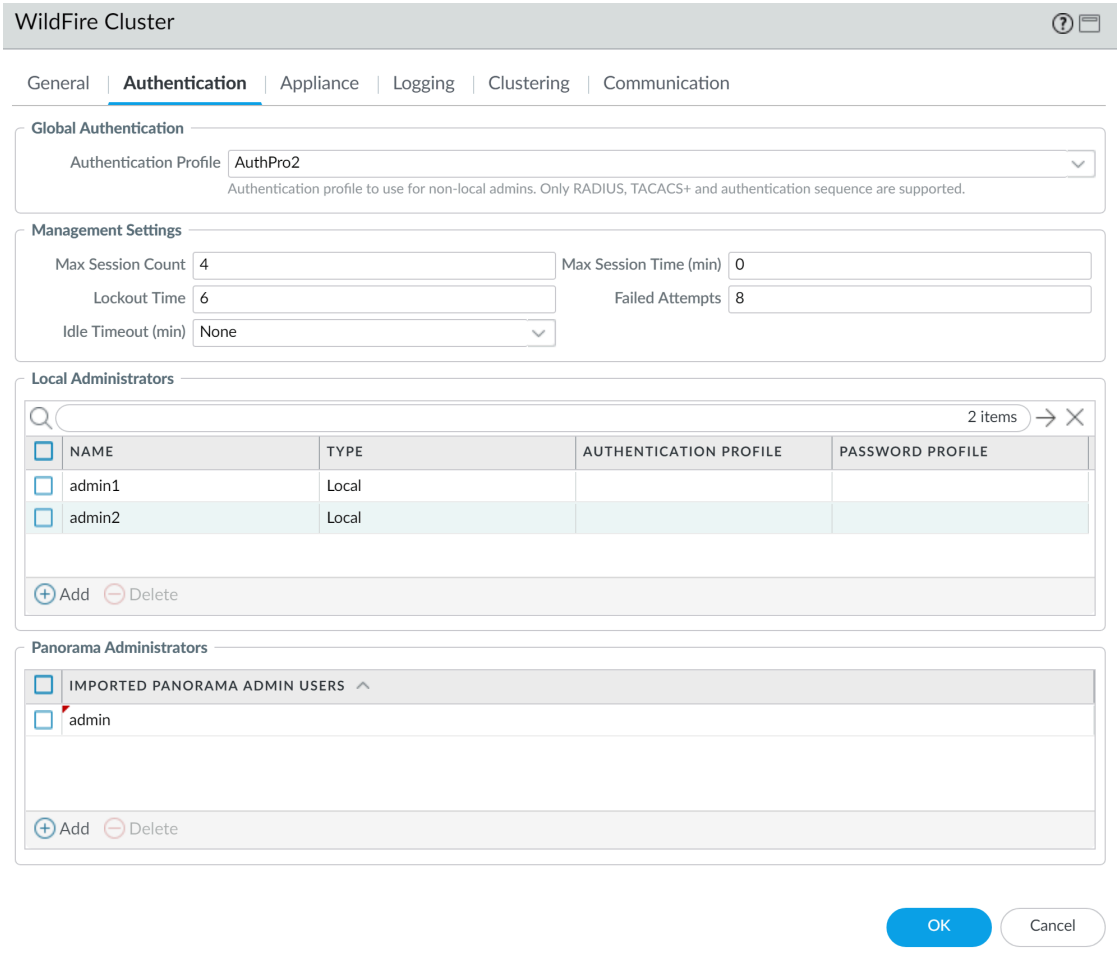
   Administrators may either be added as a local administrator or as an imported Panorama administrator—but not both. Adding the same administrator as both a local administrator and as an imported Panorama administrator is not supported and causes the Panorama commit to fail. For example, the commit to Panorama fails if you add **`admin1`** as both a local and Panorama administrator.
   1. **Add** and configure new administrators unique to the WildFire appliances in the WildFire cluster. These administrators are specific to the WildFire appliances in the

 ©2025 Palo Alto Networks, Inc.

WildFire cluster for which they are created and you manage these administrators from this table.

2. **Add** any administrators configured on Panorama. These administrators are created on Panorama and imported to the WildFire appliances in the WildFire cluster.

5. Click **OK** to save the WildFire cluster authentication configuration.



STEP 5 | **Commit** and then **Commit and Push** your configuration changes.

STEP 6 | Access the WildFire appliance CLI to verify you can successfully access a WildFire appliance using the local admin user.

Configure TACACS+ Authentication for a WildFire Cluster

You can use a TACACS+ server to authenticate administrative access to the CLI of the WildFire appliances in a WildFire cluster. You can also define Vendor-Specific Attributes (VSAs) on the TACACS+ server to manage administrator authorization. Using VSAs enables you to quickly change the roles, access domains, and user groups of administrators through your directory service, which is often easier than reconfiguring settings on Panorama.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Configure a Cluster Centrally on Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Configure TACACS+ authentication.

Administrator accounts configured for TACACS+ authentication are required to have Superuser *admin role privileges to successfully configure authentication for Wildfire appliances in the WildFire cluster.*

1. Add a TACACS+ server profile.

   The profile defines how a WildFire appliance connects to the TACACS+ server.

   1. Select **Panorama** > **Server Profiles** > **TACACS+** and **Add** a profile.
   2. Enter a **Profile Name** to identify the server profile.
   3. Enter a **Timeout** interval in seconds after which an authentication request times out (default is 3; range is 1–20).
   4. Select the **Authentication Protocol** (default is **CHAP**) that Panorama uses to authenticate to the TACACS+ server.
   5. Select **CHAP** if the TACACS+ server supports that protocol; it is more secure than **PAP**.
   6. **Add** each TACACS+ server and enter the following:
      1. **Name** to identify the server.
      2. **TACACS+ Server** IP address or FQDN.
      3. **Secret/Confirm Secret** (a key to encrypt usernames and passwords).
      4. Server **Port** for authentication requests (default is 49).
   7. Click **OK** to save the server profile.
2. Assign the TACACS+ server profile to an authentication profile.

   The authentication profile defines authentication settings that are common to a set of administrators.

   1. Select **Panorama** > **Authentication Profile** and **Add** a profile.
   2. Enter a **Name** to identify the profile.
   3. Set the **Type** to **TACACS+**.
   4. Select the **Server Profile** you configured.
   5. Select **Retrieve user group from TACACS+** to collect user group information from VSAs defined on the TACACS+ server.

      Panorama matches the group information against the groups you specify in the Allow List of the authentication profile.
   6. Select **Advanced** and, in the Allow List, **Add** the administrators that are allowed to authenticate with this authentication profile.
   7. Click **OK** to save the authentication profile.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the authentication for the WildFire cluster.

1. Select **Panorama** > **Managed WildFire Clusters** and select the WildFire cluster you previously added.
2. Select the **Authentication Profile** you configured in the previous step.

   If a global authentication profile is not assigned you must assign an authentication profile to each individual local administrator to leverage remote authentication.
3. Configure the authentication **Timeout Configuration** for a WildFire appliance.
   1. Enter the number of **Failed Attempt**s before a user is locked out of a WildFire appliance CLI.
   2. Enter the **Lockout Time**, in minutes, for which a WildFire appliance locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access a WildFire appliance.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
4. Add the WildFire appliance administrators.

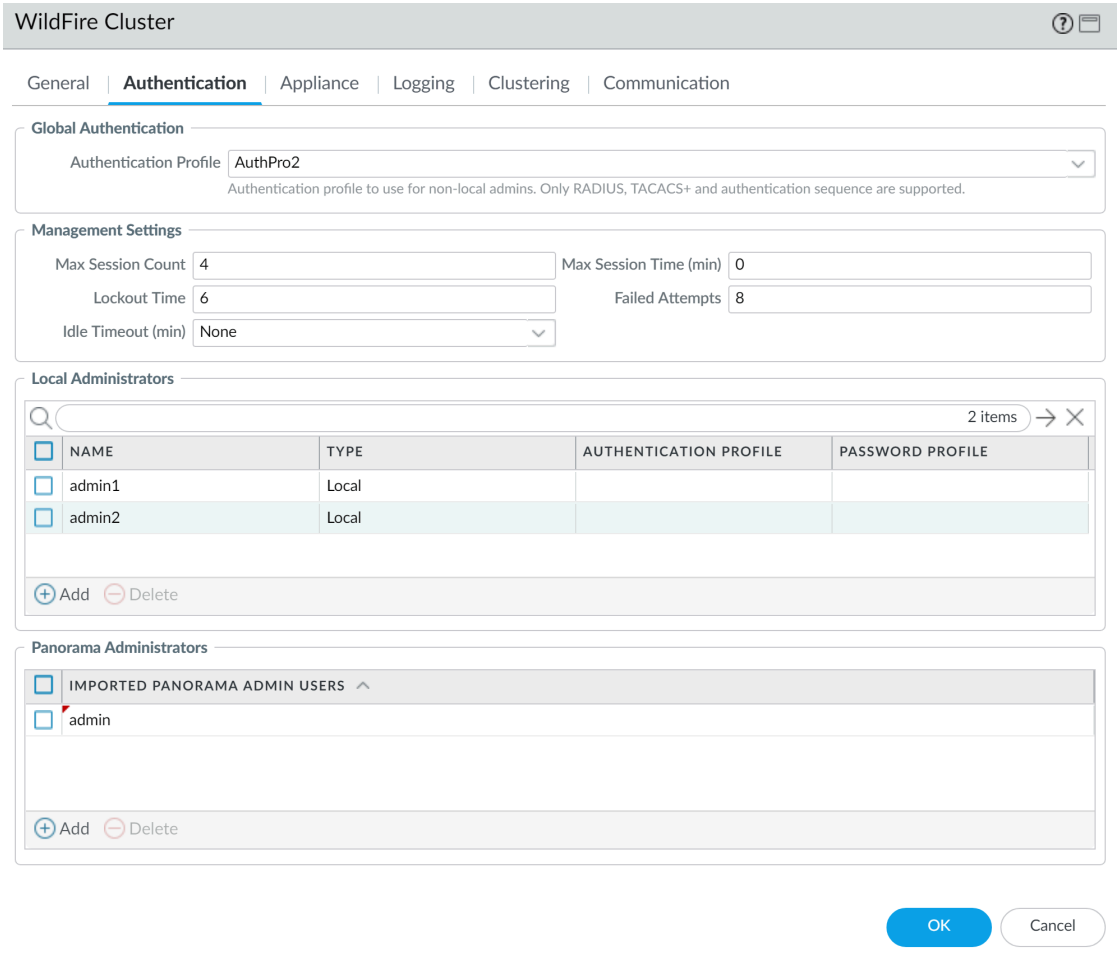
   Administrators may either be added as a local administrator or as an imported Panorama administrator—but not both. Adding the same administrator as both a local administrator and as an imported Panorama administrator is not supported and causes the Panorama commit to fail. For example, the commit to Panorama fails if you add **`admin1`** as both a local and Panorama administrator.
   1. **Add** and configure new administrators unique to the WildFire appliances in the WildFire cluster. These administrators are specific to the WildFire appliances in the

 ©2025 Palo Alto Networks, Inc.

WildFire cluster for which they are created and you manage these administrators from this table.

2. **Add** any administrators configured on Panorama. These administrators are created on Panorama and imported to the WildFire appliances in the WildFire cluster.

5. Click **OK** to save the WildFire cluster authentication configuration.



STEP 5 | **Commit** and then **Commit and Push** your configuration changes.

STEP 6 | Access the WildFire appliance CLI to verify you can successfully access a WildFire appliance using the local admin user.

Configure LDAP Authentication for a WildFire Cluster

You can use LDAP to authenticate end users to access the CLI of the WildFire appliances in a WildFire cluster.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Configure a Cluster Centrally on Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Add an LDAP server profile.

The profile defines how a WildFire appliance connects to the LDAP server.

Administrator accounts configured for LDAP authentication are required to have Superuser admin role privileges to successfully configure authentication for WildFire appliances in the WildFire cluster.

1. Select **Panorama** > **Server Profiles** > **LDAP** and **Add** a server profile.
2. Enter a **Profile Name** to identify the server profile.
3. **Add** the LDAP servers (up to four). For each server, enter a **Name** (to identify the server), **LDAP Server** IP address or FQDN, and server **Port** (default 389).

   If you use an FQDN address object to identify the server and you subsequently change the address, you must commit the change for the new server address to take effect.
4. Select the server **Type**.
5. Select the **Base DN**.

   To identify the Base DN of your directory, open the **Active Directory Domains and Trusts** Microsoft Management Console snap-in and use the name of the top-level domain.
6. Enter the **Bind DN** and **Password** to enable the authentication service to authenticate the firewall.

   The Bind DN account must have permission to read the LDAP directory.
7. Enter the **Bind Timeout** and **Search Timeout** in seconds (default is 30 for both).
8. Enter the **Retry Interval** in seconds (default is 60).
9. (Optional) If you want the endpoint to use SSL or TLS for a more secure connection with the directory server, enable the option to **Require SSL/TLS secured connection** (enabled by default). The protocol that the endpoint uses depends on the server port:
   - 389 (default)—TLS (Specifically, the WildFire appliance uses the StartTLS operation, which upgrades the initial plaintext connection to TLS.)
   - 636—SSL
   - Any other port—The WildFire appliance first attempts to use TLS. If the directory server doesn't support TLS, the WildFire appliance falls back to SSL.
10. (Optional) For additional security, enable to the option to **Verify Server Certificate for SSL sessions** so that the endpoint verifies the certificate that the directory server presents for SSL/TLS connections. To enable verification, you must also enable the

 ©2025 Palo Alto Networks, Inc.

option to **Require SSL/TLS secured connection**. For verification to succeed, the certificate must meet one of the following conditions:

- It is in the list of Panorama certificates: **Panorama** > **Certificate Management** > **Certificates** > **Device Certificates.** If necessary, import the certificate into Panorama.
- The certificate signer is in the list of trusted certificate authorities: **Panorama** > **Certificate Management** > **Certificates**.

11. Click **OK** to save the server profile.

STEP 4 | Configure the authentication for the WildFire cluster.

1. Select **Panorama** > **Managed WildFire Clusters** and select the WildFire cluster you previously added.
2. Configure the authentication **Timeout Configuration** for a WildFire appliance.
   1. Enter the number of **Failed Attempt**s before a user is locked out of a WildFire appliance CLI.
   2. Enter the **Lockout Time**, in minutes, for which a WildFire appliance locks out a user account after that user reaches the configured number of **Failed Attempts**.
   3. Enter the **Idle Timeout**, in minutes, before the user account is automatically logged out due to inactivity.
   4. Enter the **Max Session Count** to set how many user accounts can simultaneously access a WildFire appliance.
   5. Enter the **Max Session Time** the administrator can be logged in before being automatically logged out.
3. Add the WildFire appliance administrators.

   Administrators may either be added as a local administrator or as an imported Panorama administrator—but not both. Adding the same administrator as both a local administrator and as an imported Panorama administrator is not supported and causes the Panorama commit to fail. For example, the commit to Panorama fails if you add **`admin1`** as both a local and Panorama administrator.

   - Configure the local administrators.

     Configure new administrators unique to the WildFire appliances in the WildFire cluster. These administrators are specific to the WildFire appliances in the WildFire

 ©2025 Palo Alto Networks, Inc.

cluster for which they are created and you manage these administrators from this table.

1. **Add** one or more new local administrator.
2. Enter a **Name** for the local administrator.
3. Assign an **Authentication Profile** you previously created.

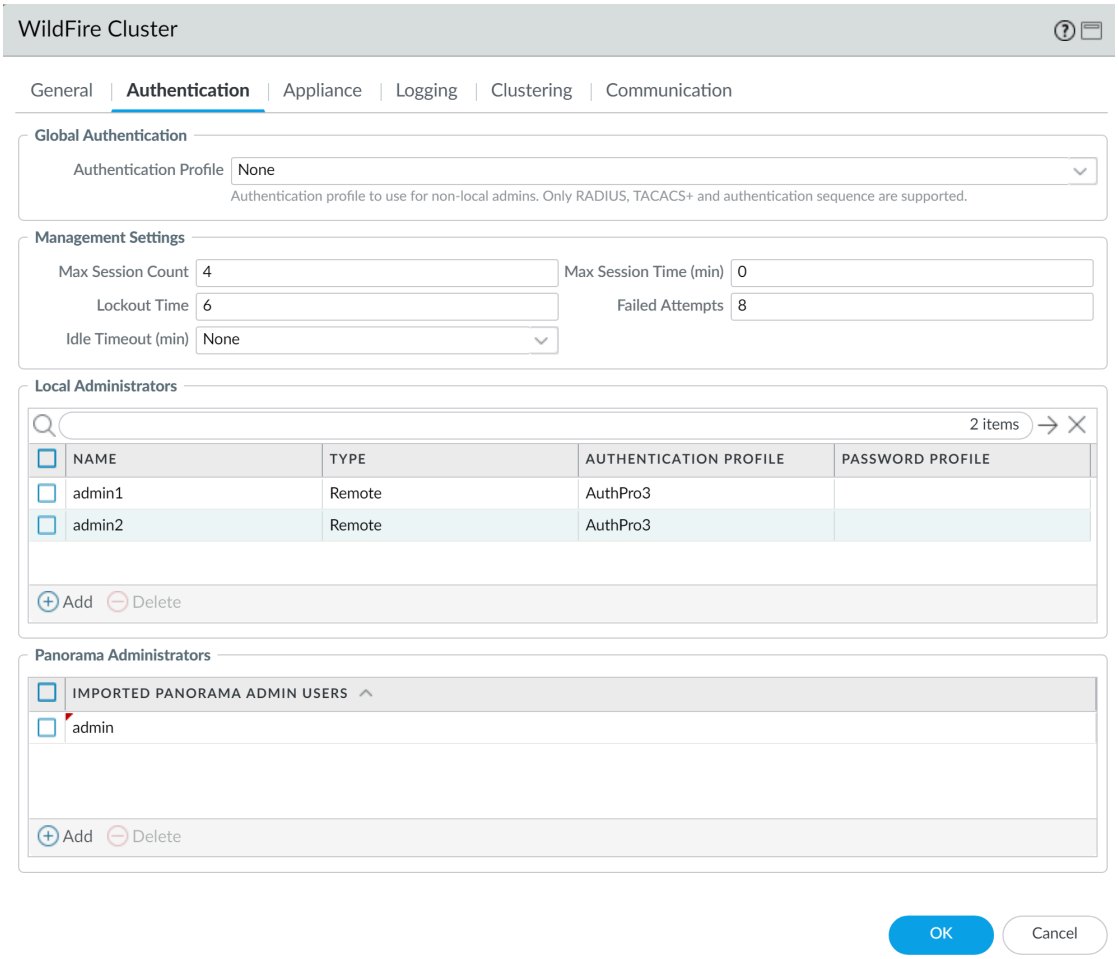
   LDAP authentication profiles are supported only for individual local administrators.
4. Enable (check) **Use Public Key Authentication (SSH)** to import a public key file for authentication.
5. Select a **Password Profile** to set the expiration parameters.

- Import existing Panorama administrators

  Import existing administrators configured on Panorama. These administrators are configured and managed on Panorama and imported to all WildFire appliances in the WildFire cluster.

1. **Add** an existing Panorama administrator

4. Click **OK** to save the WildFire cluster authentication configuration.



 ©2025 Palo Alto Networks, Inc.

STEP 5 | **Commit** and then **Commit and Push** your configuration changes.

STEP 6 | Access the WildFire appliance CLI to verify you can successfully access the WildFire appliance using the local admin user.

Remove a Cluster from Panorama Management

To remove a cluster from Panorama management, **Panorama** > **Managed WildFire Clusters** and select the row of the cluster you want to remove (do not click the cluster name) and **Remove From Panorama**.

If you remove a WildFire appliance cluster from Panorama management, the Panorama web interface places the WildFire appliances in that cluster into read-only mode. Although the WildFire appliances in the removed cluster display in the Panorama web interface, when in read-only mode, you can't push configurations to the WildFire appliances or manage them with Panorama. After being removed from Panorama management, the WildFire appliance cluster members use the local cluster configuration and you can manage the cluster using the local CLI.

To manage the WildFire appliances in the cluster with Panorama after you remove the cluster from Panorama management, import the cluster back into Panorama (**Panorama** > **Managed WildFire Clusters** > **Import Cluster Config**).

STEP 1 | Select the cluster's controller node. The cluster name populates **Cluster** automatically.

STEP 2 | Click **OK**. The cluster backup controller node and worker nodes populate automatically.

STEP 3 | Click **OK** to import the cluster.

STEP 4 | **Commit** the changes.

Configure Appliance-to-Appliance Encryption Using Predefined Certificates Centrally on Panorama

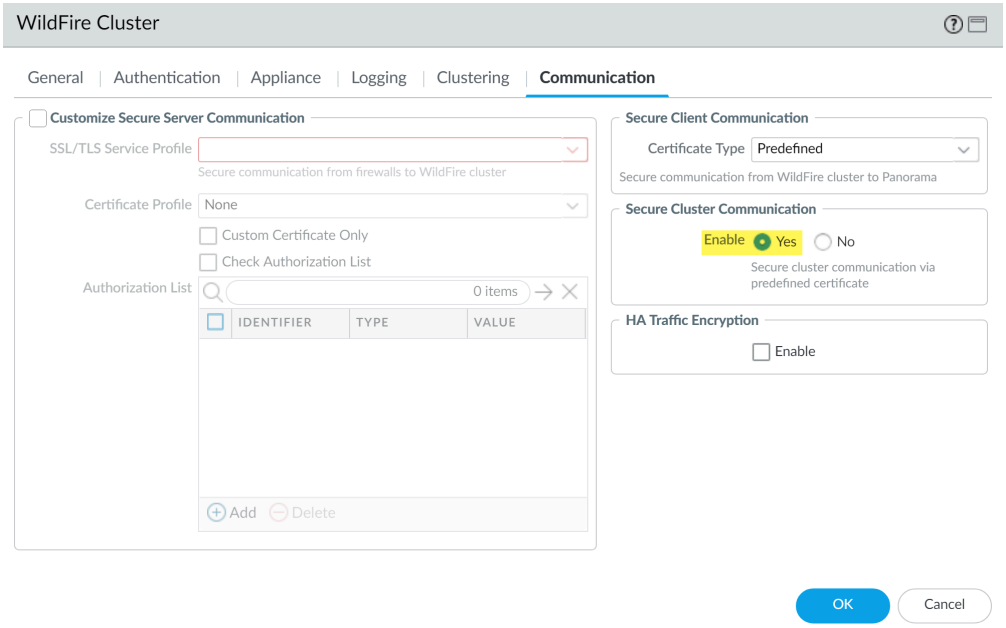
STEP 1 | Upgrade each managed WildFire appliance to PAN-OS 8.1.x. All managed appliances must be running PAN-OS 8.1 or later to enable appliance-to-appliance encryption.

STEP 2 | Verify that your WildFire appliance cluster has been properly configured and is operating in a healthy state.

STEP 3 | On Panorama, select **Panorama > Managed WildFire Clusters > *WF_cluster_name*> Communication.**

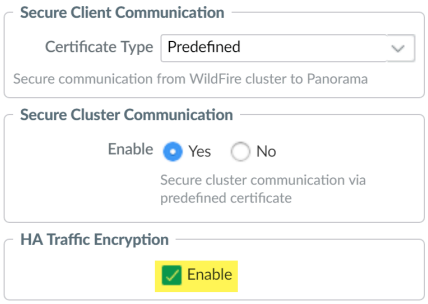
 ©2025 Palo Alto Networks, Inc.

STEP 4 | **Enable** Secure Cluster Communication.



STEP 5 | (Recommended) **Enable** HA Traffic Encryption. This optional setting encrypts the HA traffic between the HA pair and is a Palo Alto Networks recommended best practice.

HA Traffic Encryption cannot be disabled when operating in FIPS/CC mode.



STEP 6 | Click **OK** to save the **WildFire Cluster** settings.

STEP 7 | **Commit** your changes.

Configure Appliance-to-Appliance Encryption Using Custom Certificates Centrally on Panorama

STEP 1 | Upgrade each managed WildFire appliance to PAN-OS 8.1.x. All managed appliances must be running PAN-OS 8.1 or later to enable appliance-to-appliance encryption.

STEP 2 | Verify that your WildFire appliance cluster has been properly configured and is operating in a healthy state.

STEP 3 | Review your existing WildFire secure communications configuration. Keep in mind, if you previously configured the WildFire appliance and the firewall for secure

 ©2025 Palo Alto Networks, Inc.

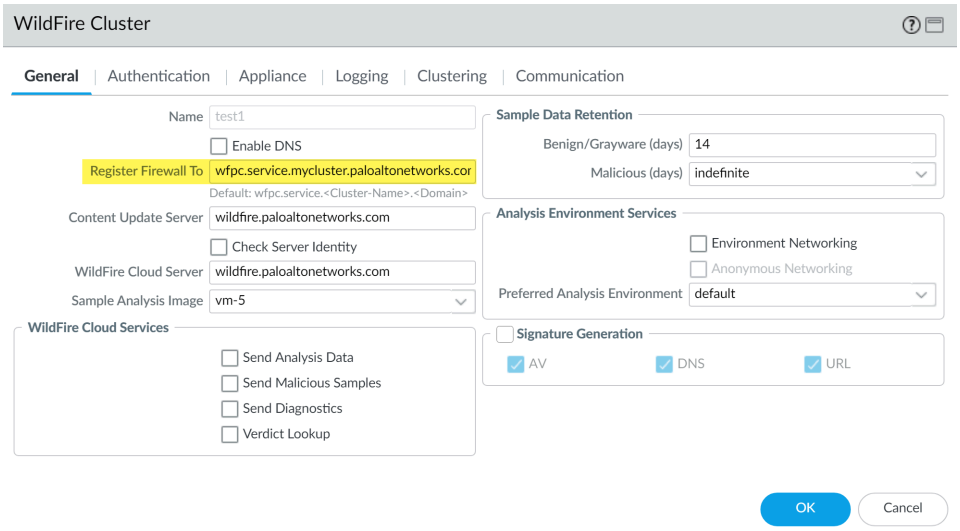
communications using a custom certificate, you can also use that custom certificate for secure communications between WildFire appliances.

1. Select **Panorama >Managed WildFire Clusters> *WF_cluster_name*> Communication.**
2. If **Customize Secure Server Communication** has been enabled and you would like to use that certificate, identify the details of the custom certificate being used. Otherwise proceed to Step 5 to begin the process of installing a new custom certificate.
3. Determine the custom certificate FQDN (DNS name) that will be used to define the firewall registration address in Step 4.

Make sure to note the custom certificate name and the associated FQDN. These are referenced several times during the configuration process.

STEP 4 | Configure the firewall registration address on Panorama.

1. On Panorama, select **Panorama >Managed WildFire Clusters> *WF_cluster_name*> General.**
2. In the **Register Firewall To** field, specify the DNS name used for authentication found in the custom certificate (typically the SubjectName or the SubjectAltName). For example, the default domain name is **`wfpc.service.mycluster.paloaltonetworks.com`**



STEP 5 | Configure WildFire **Secure Server Communication** settings on Panorama. If you already configured secure communications between the firewall and the WildFire cluster and are using the existing custom certificate, proceed to Step 4 below.

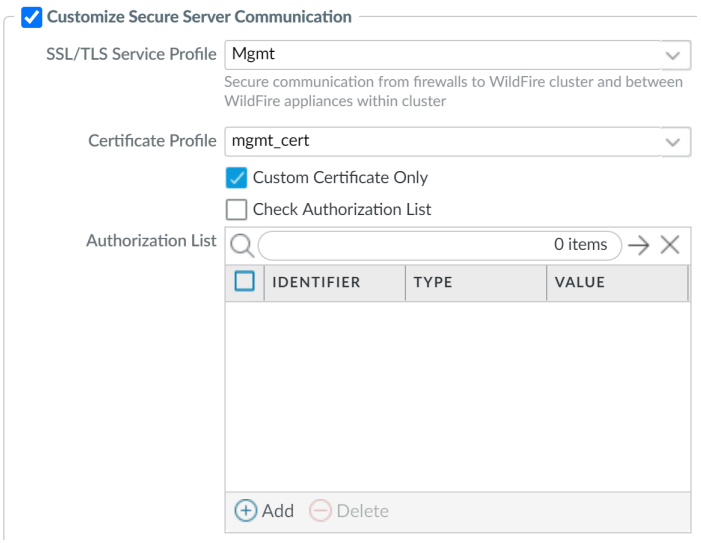
1. On Panorama, select **Panorama> Managed WildFire Clusters> *WF_cluster_name*> Communication.**
2. Click **Customize Secure Server Communication**.
3. Configure and deploy custom certificates used by the WildFire appliances and the associated firewall. The SSL/TLS service profile defines the custom certificate used by WildFire appliances to communicate with WildFire appliance peers and to the firewall.

 ©2025 Palo Alto Networks, Inc.

You must also configure the custom certificate settings on the firewall associated with the WildFire appliance cluster. This is configured later in Step 9.

1. Open the SSL/TLS Service Profile drop-down and click SSL/TLS Service Profile. Configure an SSL/TLS service profile with the custom certificate that you want to use. After you configure the SSL/TLS service profile, click OK and select the newly created SSL/TLS Service profile.
2. Open the Certificate Profile drop-down and click Certificate Profile. Configure a Certificate Profile that identifies the custom certificate used to establish secure connections between the firewall and WildFire appliances, as well as between peer WildFire appliances. After you configure the Certificate Profile, click OK and select the newly created profile.

4. Select the **Custom Certificate Only** check box. This allows you to use the custom certificates that you configured instead of the default preconfigured certificates.
5. (Optional) Configure an authorization list. The authorization list checks the custom certificate Subject or Subject Alt Name; if the **Subject** or **Subject Alt Name** presented with the custom certificate does not match an identifier on the authorization list, authentication is denied.
   1. **Add** an Authorization List.
   2. Select the **Subject** or **Subject Alt Name** configured in the custom certificate profile as the Identifier type.
   3. Enter the Common Name if the identifier is Subject or and IP address, hostname or email if the identifier is Subject Alt Name.
   4. Click **OK**.
   5. Select **Check Authorization List** to enforce the authorization list.
6. Click **OK**.



STEP 6 | **Enable** Secure Cluster Communication.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | (Recommended) **Enable** HA Traffic Encryption. This optional setting encrypts the HA traffic between the HA pair and is a Palo Alto Networks recommended best practice.

HA Traffic Encryption cannot be disabled when operating in FIPS/CC mode.

STEP 8 | Click **OK** to save the **WildFire Cluster** settings.

STEP 9 | Configure the firewall **Secure Communication Settings** on Panorama to associate the WildFire appliance cluster with the firewall custom certificate. This provides a secure communications channel between the firewall and WildFire appliance cluster. If you already configured secure communications between the firewall and the WildFire appliance cluster and are using the existing custom certificate, proceed to the next step.

1. Select **Device** > **Setup** > **Management** > **Secure Communication Settings** and click the **Edit** icon in **Secure Communication Settings** to configure the firewall custom certificate settings.
2. Select the **Certificate Type**, **Certificate**, and **Certificate Profile** from the respective drop-downs and configure them to use the custom certificate.
3. Under Customize Communication, select **WildFire Communication**.
4. Click **OK**.

STEP 10 | **Commit** your changes.

View WildFire Cluster Status Using Panorama

To confirm that a configured WildFire appliance cluster is operating correctly, you can view the current status using the Panorama appliance.

Palo Alto Networks recommends using the WildFire appliance CLI to verify the status of your WildFire cluster. Additional status details that are not visible from Panorama are displayed in the command output.

STEP 1 | On the primary Panorama appliance, select **Panorama** > **Managed WildFire Clusters**.

STEP 2 | In the **Cluster Status** column, verify that:

1. The wfpc and signature services are running.
2. No other operations are present. Abnormal operations and their status conditions include:
   - Decommission [requested / ongoing / denied / success / fail]
   - Suspend [requested / ongoing / denied / success / fail]
   - Reboot [requested / ongoing / denied / success / fail]
   - Cluster [offline / splitbrain / unready]
   - Service [suspended / none]
   - HA [peer-offline / cfg-not-sync / cfg-sync-off]

 ©2025 Palo Alto Networks, Inc.

STEP 3 | In the **Config Status** column, verify that:

1. The appliance configuration is **In Sync** with the configuration stored on the Panorama appliance.
2. No other status is present. Abnormal status conditions include:
   - **Out of Sync** [The appliance configuration is not in sync with its saved configuration on Panorama. You can mouse over the magnifying glass to display the cause of the sync failure].

STEP 4 | In the **Connected** column, verify that the configured WildFire appliances show a status of **Connected**.

 ©2025 Palo Alto Networks, Inc.

Manage Licenses and Updates

You can use the Panorama™ management server to centrally manage licenses, software updates, and content updates on firewalls and Dedicated Log Collectors. When you deploy licenses or updates, Panorama checks in with the Palo Alto Networks® licensing server or update server, verifies the request validity, and then allows retrieval and installation of the license or update. This capability facilitates deployment by eliminating the need to repeat the same tasks on each firewall or Dedicated Log Collector. It is particularly useful for managing firewalls that don't have direct internet access or for managing Dedicated Log Collectors, which don't have a web interface.

Before deploying updates, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility for important details about update version compatibility.

You must activate a support subscription directly on each firewall; you cannot use Panorama to deploy support subscriptions.

To activate licenses or install updates on the Panorama management server, see Register Panorama and Install Licenses and Install Content and Software Updates for Panorama.

- Manage Licenses on Firewalls Using Panorama

Manage Licenses on Firewalls Using Panorama

The following steps describe how to retrieve new licenses using an authentication (*auth*) code and push the license keys to managed firewalls. It also describes how to manually update (refresh) the license status of firewalls that both have direct internet access and those that do not have direct internet access. Panorama™ automatically performs a daily check-in with the licensing server, retrieves license updates and renewals, and pushes them to the firewalls. The check-in is hard-coded to occur between 1 a.m. and 2 a.m.; you cannot change this schedule.

You cannot use Panorama to activate the support license for firewalls. You must access the firewalls individually to activate their support licenses.

To activate licenses for Panorama, see Register Panorama and Install Licenses.

- Activate newly purchased licenses.
  1. Select **Panorama** > **Device Deployment** > **Licenses** and **Activate**.
  2. Enter the **Auth Code** that Palo Alto Networks® provided for each firewall that has a new license.
  3. **Activate** the license.
  4. (WildFire® subscriptions only) Perform a commit on each firewall that has a new WildFire subscription to complete the activation:
     - **Commit** any pending changes. You must access each firewall web interface to do this.
     - If no configuration changes are pending, make a minor change and **Commit**. For example, update a rule description and commit the change. If the firewalls belong to the same device group, you can push the rule change from Panorama to initiate a commit on all those firewalls instead of accessing each firewall separately.

     Check that the WildFire Analysis profile rules *include the advanced file types that the WildFire subscription supports.*

- Update the license status of firewalls.
  1. Select **Panorama** > **Device Deployment** > **Licenses**.

     Each entry on the page indicates whether the license is active or inactive and displays the expiration date for active licenses.
  2. If you previously activated auth codes for the support subscription directly on the firewalls, click **Refresh** and select the firewalls from the list. Panorama retrieves the

 ©2025 Palo Alto Networks, Inc.

license, deploys it to the firewalls, and updates the licensing status on the Panorama web interface.

3. (Enterprise Data Loss Prevention (DLP) license only) Push the updated license to the managed firewalls leveraging Enterprise DLP.
   1. Select **Commit** and **Commit to Panorama.**
   2. Select **Commit** > **Push to Devices** and **Edit Selections**.
   3. Select **Templates** and select the template stack associated with the managed firewalls leveraging Enterprise DLP.

      Click **OK** to continue.
   4. **Push** the template configuration to successfully update the Enterprise DLP license.

 ©2025 Palo Alto Networks, Inc.

 ©2025 Palo Alto Networks, Inc.

Monitor Network Activity

The Panorama™ management server provides a comprehensive, graphical view of network traffic. Using the visibility tools on Panorama—the Application Command Center (ACC), logs, and report generation capabilities—you can centrally analyze, investigate and report on all network activity, identify areas with potential security impact, and translate them into secure application enablement policies.

This section covers the following topics:

- Use Panorama for Visibility
- Ingest Traps ESM Logs on Panorama
- Use Case: Monitor Applications Using Panorama
- Use Case: Respond to an Incident Using Panorama

Use Panorama for Visibility

In addition to its central deployment and firewall configuration features, Panorama also allows you to monitor and report on all traffic that traverses your network. While the reporting capabilities on Panorama and the firewall are very similar, the advantage that Panorama provides is that it is a single pane view of aggregated information across all your managed firewalls. This aggregated view provides actionable information on trends in user activity, traffic patterns, and potential threats across your entire network.

Using the Application Command Center (ACC), the App-Scope, the log viewer, and the standard, customizable reporting options on Panorama, you can quickly learn more about the traffic traversing the network. The ability to view this information allows you to evaluate where your current policies are adequate and where they are insufficient. You can then use this data to augment your network security strategy. For example, you can enhance the security rules to increase compliance and accountability for all users across the network, or manage network capacity and minimize risks to assets while meeting the rich application needs for the users in your network.

The following topics provide a high-level view of the reporting capabilities on Panorama, including a couple of use cases to illustrate how you can use these capabilities within your own network infrastructure. For a complete list of the available reports and charts and the description of each, refer to the online help.

- Monitor the Network with the ACC and AppScope
- Analyze Log Data
- Generate, Schedule, and Email Reports
- Configure Key Limits for Scheduled Reports

Monitor the Network with the ACC and AppScope

Both the ACC and the AppScope allow you to monitor and report on the data recorded from traffic that traverses your network.

The ACC on Panorama displays a summary of network traffic. Panorama can dynamically query data from all the managed firewalls on the network and display it in the ACC. This display allows you to monitor the traffic by applications, users, and content activity—URL categories, threats, security policies that effectively block data or files—across the entire network of Palo Alto Networks next-generation firewalls.

The AppScope helps identify unexpected or unusual behavior on the network at a glance. It includes an array of charts and reports—Summary Report, Change Monitor, Threat Monitor, Threat Map, Network Monitor, Traffic Map—that allow you to analyze traffic flows by threat or application, or by the source or destination for the flows. You can also sort by session or byte count.

Device Group and Template admins can only network and ACC data for device groups within their access domains.

Use the ACC and the AppScope to answer questions such as:

 ©2025 Palo Alto Networks, Inc.

ACC	Monitor > AppScope
• What are the top applications used on the network and how many are high-risk applications? Who are the top users of high-risk applications on the network? • What are the top URL categories being viewed in the last hour?	• What are the application usage trends —what are the top five applications that have gained use and the top five that have decreased in use? • How has user activity changed over the current week as compared to last week or last month?
• What are the top bandwidth-using applications? Who are the users/hosts that consume the highest bandwidth? • What content or files are being blocked and are there specific users who trigger this File Blocking/Data Filtering rule? • What is the amount of traffic exchanged between two specific IP addresses or generated by a specific user? Where is the destination server or client located geographically?	• Which users and applications take up most of the network bandwidth? And how has this consumption changed over the last 30 days? • What are the threats on the network, and how are these incoming and outgoing traffic threats distributed geographically?

You can then use the information to maintain or enforce changes to the traffic patterns on your network. See Use Case: Monitor Applications Using Panorama for a glimpse into how the visibility tools on Panorama can influence how you shape the acceptable use policies for your network.

Here are a few tips to help you navigate the ACC:

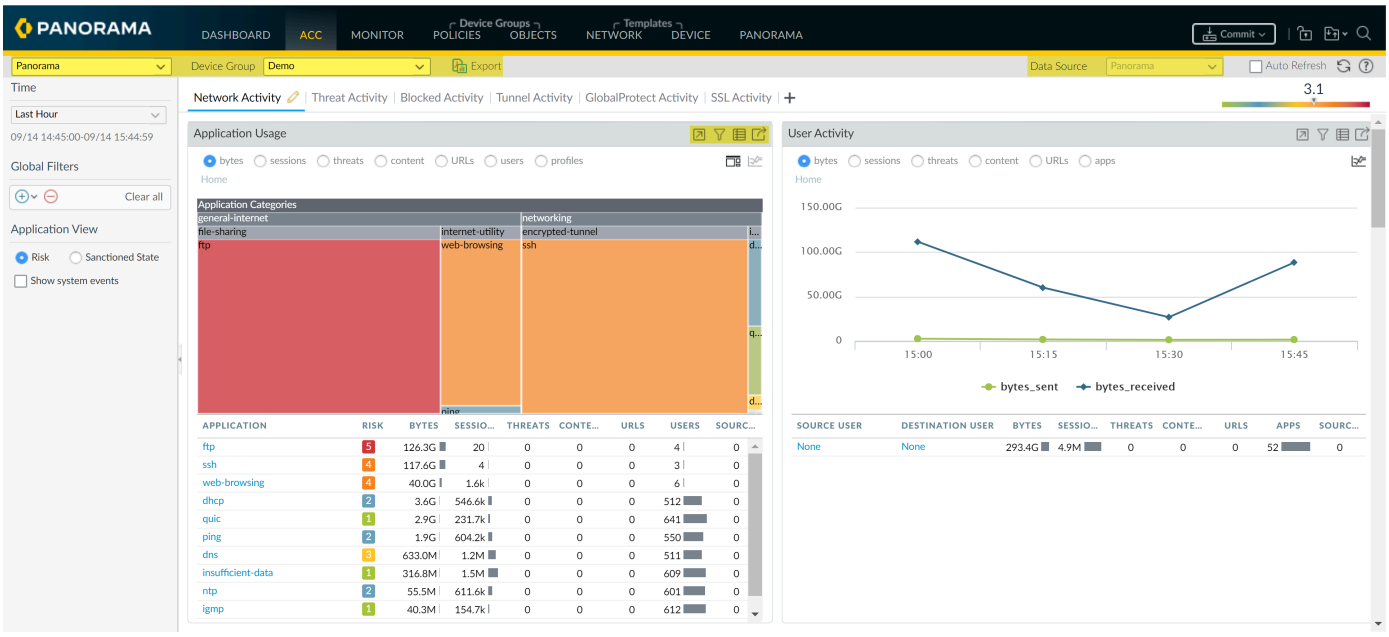


Figure 24: ACC Navigation Tips

 ©2025 Palo Alto Networks, Inc.

- **Switch from a Panorama view to a Device view**—Use the **Context** drop-down to access the web interface of any managed firewall. For details, see Context Switch—Firewall or Panorama.
- **Change Device Group and Data Source**—The default **Data Source** used to display the statistics on the charts in the ACC is **Panorama** local data, and the default **Device Group** setting is **All**. Using the local data on Panorama provides a quick load time for the charts. You can, however, change the data source to **Remote Device Data** if all the managed firewalls are on PAN-OS 7.0 or a later release. If the managed firewalls have a mix of PAN-OS 7.0 and earlier releases, you can only view Panorama data. When configured to use Remote Device Data, Panorama will poll all the managed firewalls and present an aggregated view of the data. The onscreen display indicates the total number of firewalls being polled and the number of firewalls that have responded to the query for information.
- **Select the Tabs and Widgets to View**—The ACC includes three tabs and an array of widgets that allow you to find the information that you care about. With the exception of the application usage widget and host information widget, all the other widgets display data only if the corresponding feature has been licensed on the firewall, and you have enabled logging.
- **Tweak Time Frame and Refine Data**—The reporting time period in the ACC ranges from the last 15 minutes to the last hour, day, week, month, or any custom-defined time. By default, each widget displays the top 10 items and aggregates all the remaining items as **others**. You can sort the data in each widget using various attributes—for example, sessions, bytes, threats, content, and URLs. You can also set local filters to filter the display within the table and graph in a widget, and then promote the widget filter as a global filter to pivot the view across all the widgets in the ACC.

Analyze Log Data

The **Monitor** tab on Panorama provides access to log data; these logs are an archived list of sessions that have been processed by the managed firewalls and forwarded to Panorama.

Log data can be broadly grouped into two types: those that detail information on traffic flows on your network such as applications, threats, host information profiles, URL categories, content/file types and those that record system events, configuration changes, and User-ID™ mapping information.

Based on the log forwarding configuration on the managed firewalls, the **Monitor** > **Logs** tab can include logs for traffic flows, threats, URL filtering, data filtering, host information profile (HIP) matches, and WildFire™ submissions. You can review the logs to verify a wealth of information on a given session or transaction. Some examples of this information are the user who initiated the session, the action (allow or deny) that the firewall performed on the session, and the source and destination ports, zones, and addresses. The System and Config logs can indicate a configuration change or an alarm that the firewall triggered when a configured threshold was exceeded.

*If Panorama will manage firewalls running software versions earlier than PAN-OS 7.0, specify a WildFire server from which Panorama can gather analysis information for WildFire samples that those firewalls submit. Panorama uses the information to complete WildFire Submissions logs that are missing field values introduced in PAN-OS 7.0. Firewalls running earlier releases won't populate those fields. To specify the server, select **Panorama** > **Setup** > **WildFire**, edit the General Settings, and enter the **WildFire Private Cloud** name. The default is **wildfire-public-cloud**, which is the WildFire cloud hosted in the United States.*

 ©2025 Palo Alto Networks, Inc.

Generate, Schedule, and Email Reports

You can configure reports to run immediately or schedule them to run at specific intervals. You can save and export the reports or email them to specific recipients. Emailing is particularly useful if you want to share reports with administrators who do not have access to Panorama. Panorama supports the same report types as the Palo Alto Networks firewall.

Beginning with Panorama 10.0.2 and Cloud Services plugin version 1.8.0, you can generate scheduled reports on Strata Logging Service data. This setting is disabled by default and must be manually enabled on Panorama. This setting status (enabled or disabled) persists across PAN-OS upgrades, downgrades, and if you uninstall the Cloud Services plugin. To generate reports from Strata Logging Service data, you must first enable the feature from the Panorama CLI.

It is recommended that you install matching software releases on Panorama and the firewalls for which you will generate reports. For example, if the Panorama management server runs Panorama 10.0, install PAN-OS 10.2 on its managed firewalls before generating the reports. This practice avoids issues that might occur if you create reports that include fields supported in the Panorama release but not supported in an earlier PAN-OS release on the firewalls.

STEP 1 | Enable scheduled reports on Panorama.

1. Log in to the Panorama CLI.
2. Enable the scheduled reports setting.

```
admin> request plugins cloud_services logging-service sched-report-enable
```

3. Commit the configuration change.

```
admin>  configure
```

```
admin#  commit force
```

4. Verify the scheduled report setting is enabled.

```
admin> show system state | match sched-report
```

The output displays `cfg.report.lcass-sched-reports-enabled: True` if enabled.

The output displays `cfg.report.lcass-sched-reports-enabled: False` or will not return an output if disabled.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Configure Panorama predefined reports.

1. Select **Panorama** > **Setup** > **Management** and edit **Logging and Reporting**.
2. Select **Log Export and Reporting** and enable (check) **Use Data for Pre-Defined Reports** to offload hourly report aggregation to Log Collectors.

   (Strata Logging Service only) This step is required to generate scheduled reports for logs stored on Strata Logging Service.

   Enabling this setting is recommended for VM-50, VM-50 Lite and PA-200 firewalls. Enabling this setting is optional for all other managed firewall models.
3. Select **Pre-Defined Reports** and enable (check) predefined reports to push from Panorama.
4. Select **Commit** > **Commit to Panorama** and **Commit** your configuration changes.
5. (VM-50, VM-50 Lite, and PA-200 firewalls only) Access the firewall CLI to enable predefined reports.

   This command must be run on each VM-50, VM-50 Lite, and PA-200 firewall.

   ```
   admin> debug run-panorama-predefined-report yes
   ```

STEP 3 | Configure Panorama to receive and store user and user group information that it receives from firewalls.

Required to generate reports based on usernames and groups instead of just IP addresses.

1. If you want Panorama to include user group information in reports, upgrade the managed firewalls to PAN-OS 8.1 or a later release. Panorama cannot synchronize group information from firewalls running earlier releases.
2. Select **Panorama** > **Setup** > **Management**, edit the Panorama Settings, and **Enable reporting and filtering on groups**.
3. Add a Device Group if you haven't already. For each device group:
   - Select a **Master Device**, which is the firewall that provides user and user group information to Panorama.
   - Enable Panorama to **Store users and groups from Master Device**.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Generate reports.

Scheduled and Run Now summary reports for the same database and timeframe have discrepancies in the data displayed in each report. This is due to how Log Collectors and firewalls aggregate logs during hourly aggregation.

The steps to generate a report depend on the type.

- Custom report:
  1. Select **Monitor** > **Manage Custom Reports** and **Add** the report.
  2. Enter a **Name** to identify the report.
  3. Select a **Database** for the report.

     You can base the report on **Summary Databases** or **Detailed Logs** databases.

     To base the report on logs stored on the Panorama management server and Log Collectors, select **Panorama Data** (recommended for faster performance).

     To base the reports on logs stored on the managed firewalls, select **Remote Device Data**. This option is for cases where the firewalls might have logs that were not yet forwarded to Panorama. However, because Panorama must query the firewalls directly, this option is slower.
  4. Select **Scheduled**.
  5. Define your log filtering criteria by selecting the **Time Frame**, **Sort By** order, **Group By** preference, and the columns (log attributes) that the report will display.

     *Selecting the **Sort By** order is required in order to generate an accurate report. If you do not select a **Sort By** order, the generated custom report is populated with the most recent log matches for the selected database.*
  6. From the **Available Columns** list, select the columns that you want to display in the report, and add them to the **Selected Columns** list.

     The columns that you add, appear in the **Selected Column** list. You must add at least one column to save or run the report.

     *To generate scheduled reports, ensure that you add the field you selected in the **Sort By** list to the **Selected Columns** list.*
  7. (Optional) Use the **Query Builder** to further refine the log filtering criteria based on log attributes.
  8. To test the report settings, select **Run Now**. If necessary, modify the settings to change the information that the report displays.
  9. Click **OK** to save the custom report.
- **PDF Summary Report**:
  1. Select **Monitor** > **PDF Reports** > **Manage PDF Summary** and add the report.
  2. Enter a **Name** to identify the report.
  3. Use the drop-down for each report group and select one or more of the elements to design the PDF Summary Report. You can include up to 18 elements.
  4. Click **OK** to save the settings.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Configure a **Report Group**.

It can include predefined reports, PDF Summary reports, and custom reports. Panorama compiles all the included reports into a single PDF.

1. Select **Monitor** > **PDF Reports** > **Report Groups** and **Add** a report group.
2. Enter a **Name** to identify the report group.
3. (Optional) Select **Title Page** and add a **Title** for the PDF output.
4. Select reports in the Predefined Report, Custom Report, and PDF Summary Report lists.
5. **Add** the selected reports to the report group.
6. Click **OK** to save the settings.

STEP 6 | Configure an Email server profile.

The profile defines how the firewall connects to the server and sends email.

1. Select **Panorama** > **Server Profiles** > **Email** and **Add** a server profile.
2. Enter a **Name** to identify the profile.
3. **Add** up to four SMTP servers and **Add** the following information for each one:
   - **Name**—A name to identify the SMTP server (1 to 31 characters). This field is just a label and doesn't have to be the hostname of an existing server.
   - **Email Display Name**—The name to display in the From field of the email.
   - **From**—The email address where notification emails will be sent from.
   - **To**—The email address to which notification emails will be sent.
   - **Additional Recipient**—To send notifications to a second account, enter the additional address here.
   - **Email Gateway**—The IP address or hostname of the SMTP gateway to use to send the emails.
4. Click **OK** to save the profile.

STEP 7 | Schedule the report for email delivery.

1. Select **Monitor** > **PDF Reports** > **Email Scheduler** and **Add** an email scheduler profile.
2. Enter a **Name** to identify the profile.
3. Select the **Report Group**, the Email server profile you just created (**Email Profile**), and the **Recurrence** for the report (default is **Disable**).
4. **Send test email** to verify that the email settings are accurate.
5. Click **OK** to save your changes.
6. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

Configure Key Limits for Scheduled Reports

The Panorama™ management server and the PA-7000 Series firewall reports utilize keys (unique values on which you can aggregate) from one or more Log Collector to build and generate reports. To improve the accuracy of scheduled reports, you can now configurable the maximum and minimum key limits. By increasing the number of keys supported, scheduled reports can now include more data that can be aggregated, sorted, and grouped.

 ©2025 Palo Alto Networks, Inc.

The default minimum key limit is based on the **Sort By** and **Group By** values configured for the scheduled report using the following calculation:

<Sort By value> x 100 x <Group By value>

For example, if **Sort By** is configured as **Top 25** and **Group By** is configured as **5 Groups**, the default minimum key limit is 12,500 keys. The **Group By** value is not factored into the calculation when set to **None**. The default minimum key limit is limited to and cannot exceed the maximum key limit.

You can only configure the key limits for the M-Series appliances and Panorama virtual appliances. The PA-7000 series key limits are not configurable.

The supported maximum and minimum keys are increased for the following Panorama models:

Panorama Model	Minimum Key Limit	Maximum Key Limit
PA-7000 Series	1,000 - Default, not configurable	25,000 - Default, not configurable
M-200	15,000	50,000
M-500	15,000	50,000
M-600	15,000	50,000
Panorama Virtual Appliance in Legacy mode	5,000	25,000
Panorama Virtual Appliance (all supported models)	15,000	50,000

STEP 1 | Log in to the Panorama CLI.

STEP 2 | Configure the maximum key limit using the following command:

You can set the maximum key limit between 0 and 50, where 50 equals 50,000 keys. In this example, we are setting the maximum key limit for the Panorama virtual appliance to 30,000 keys.

```
admin@Panorama> request max-report-keys set limit <Key Limit>
```

```
admin@Panorama> request max-report-keys set limit 30

cfg.report.max-keys-limit: 30
```

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Configure the minimum key limit using the following command:

You can set the minimum key limit between 0 and 15, where 15 equals 15,000 keys. In this example, we are setting the minimum key limit for the Panorama virtual appliance to 15,000 keys.

```
admin@Panorama> request min-report-keys set limit <Key Limit>
```

```
admin@Panorama> request min-report-keys set limit 15

cfg.report.min-keys-limit: 15
```

STEP 4 | (Optional) Set the minimum key limit to the default setting.

```
admin@Panorama> request min-report-keys set limit 0
```

STEP 5 | Commit the new maximum and minimum key limits to Panorama using the following command:

```
admin@Panorama> commit-all
```

 ©2025 Palo Alto Networks, Inc.

Ingest Traps ESM Logs on Panorama

Visibility is a critical first step in preventing and reducing the impact of an attack. To help you meet this challenge, Panorama provides an integrated view of firewall logs (events on the network) and Traps™ ESM Server logs (security events on the endpoints) so that you can trace any suspicious or malicious activity.

For awareness and context on the events observed on the network and on your endpoints, forward security events that the Traps agents report to the ESM Server on to Panorama. Panorama can serve as a Syslog receiver that ingests these logs from the Traps ESM components using Syslog over TCP, UDP, or SSL. Then, Panorama can correlate discrete security events that occur on the endpoints with what's happening on the network and generate match evidence. This evidence gives you more context on the chronology and flow of events to investigate issues and fix security gaps in your network.

STEP 1 | Define the log ingestion profile on Panorama and attach it to a Collector Group.

Panorama virtual appliance in legacy mode cannot ingest Traps logs.

1. Select **Panorama** > **Log Ingestion Profile**, and click **Add**.
2. Enter a **Name** for the profile.
3. Click **Add** and enter the details for the ESM Server. You can add up to four ESM Servers to a profile.
   1. Enter a **Source Name**.
   2. Specify the **Port** on which Panorama will be listening for syslog messages. The range is 23000 to 23999.
   3. Select the **Transport** layer protocol—TCP, UDP, or SSL.
   4. Select Traps_ESM for **External Log type** and your Traps ESM **Version**. For example, for Traps ESM 4.0 or 4.1, select **3.4.1+**.

      As Traps log formats are updated, the updated log definitions will be available through content updates on Panorama.
4. Select **Panorama** > **Collector Groups** > **Log Ingestion** and **Add** the log ingestion profile so that the Collector Group can receive logs from the ESM Server(s) listed in the profile.

   If you are enabling SSL for secure syslog communication between Panorama and the ESM Server(s), you must attach a certificate to the Managed Collectors that belong to the Collector Group (**Panorama** > **Managed Collectors** > **General**, and select the certificate to use for **Inbound Certificate for Secure Syslog**).
5. **Commit** changes to Panorama and the Collector Group.

STEP 2 | Configure Panorama as a Syslog receiver on the ESM Server.

Traps ESM 4.0 and later supports log forwarding to both an external syslog receiver and Panorama. Because earlier Traps ESM releases do not support log forwarding to multiple

 ©2025 Palo Alto Networks, Inc.

syslog receivers, you must configure Panorama as a syslog receiver in the **Syslog** settings (for Traps ESM 3.4, see Enable Log Forwarding to an External Logging Platform).

For Traps ESM 4.0 and later releases:

1. From the ESM Console, select **Settings** > **ESM** > **Panorama**, and **Enable log forwarding to Panorama**.
2. Enter the Panorama hostname or IP address as the **Panorama Server** and the **Panorama Server Port** on which Panorama is listening. Repeat this step for an optional **Panorama Failover Server**.
3. Select the Transport layer **Communication Protocol**: TCP, TCP with SSL, or UDP. If you select TCP with SSL, the ESM Server requires a server certificate to enable client authentication.

   From Panorama, you must export the root CA certificate for the Inbound Certificate for Secure Syslog, and import the certificate in to the trusted root certificate store of the host on which you have installed the ESM Server.

STEP 3 | View ESM logs and correlated events.

1. Select **Monitor** > **External Logs** > **Traps ESM** to view the logs ingested in to Panorama.
2. Select **Monitor** > **Automated Correlation Engine** > **Correlated Events**, and filter on the **Wildfire and Traps ESM Correlated C2** correlation object name to find correlated events. Panorama generates correlated events when a host on your network exhibits command and control activity that matches the behavior observed for a malicious file in the WildFire virtual environment. This correlated event alerts you to suspicious activity that a Traps agent and the firewall have observed from one or more infected hosts on your network.

 ©2025 Palo Alto Networks, Inc.

Use Case: Monitor Applications Using Panorama

This example takes you through the process of assessing the efficiency of your current policies and determining where you need to adjust them to fortify the acceptable use policies for your network.

When you log in to Panorama, the **Top Applications** widget on the **Dashboard** gives a preview of the most used applications over the last hour. To display the widget, select **Widgets** > **Application** > **Top Applications** in the toolbar. You can either glance over the list of top applications and mouse over each application block for which you want to review the details, or you can select the **ACC** tab to view the same information as an ordered list. The following image is a view of the **Top Applications** widget on the **Dashboard**.

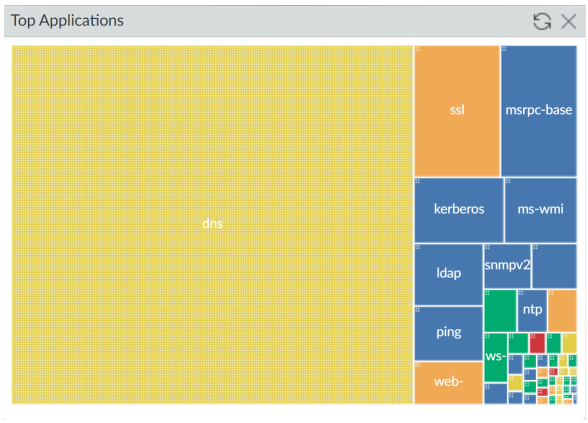


Figure 25: Top Applications Widget

The data source for this display is the application statistics database; it does not use the Traffic logs and is generated whether or not you have enabled logging for security rules. This view into the traffic on your network depicts everything that is allowed on your network and is flowing through unblocked by any policy rules that you have defined.

In the **ACC** tab, you can select and toggle the **Data Source** to be local on **Panorama** or you can query the managed firewalls (**Remote Device Data**) for the data; Panorama automatically aggregates and displays the information. For a speedier flow, consider using Panorama as the data source (with log forwarding to Panorama enabled) because the time to load data from the managed firewalls varies by the time period for which you choose to view data and the volume of traffic that is generated on your network. If your managed firewalls have a combination of PAN-OS 7.0 and earlier versions, **Remote Device Data** is not available.

The **Dashboard** example in Figure 1 shows DNS as a popular application. If you click the DNS application block, Panorama opens the **ACC** > **Network Activity** tab with DNS applied as a global filter and shows information on the application, users who accessed the application, and the details on the risk level and characteristics of the application.

 ©2025 Palo Alto Networks, Inc.

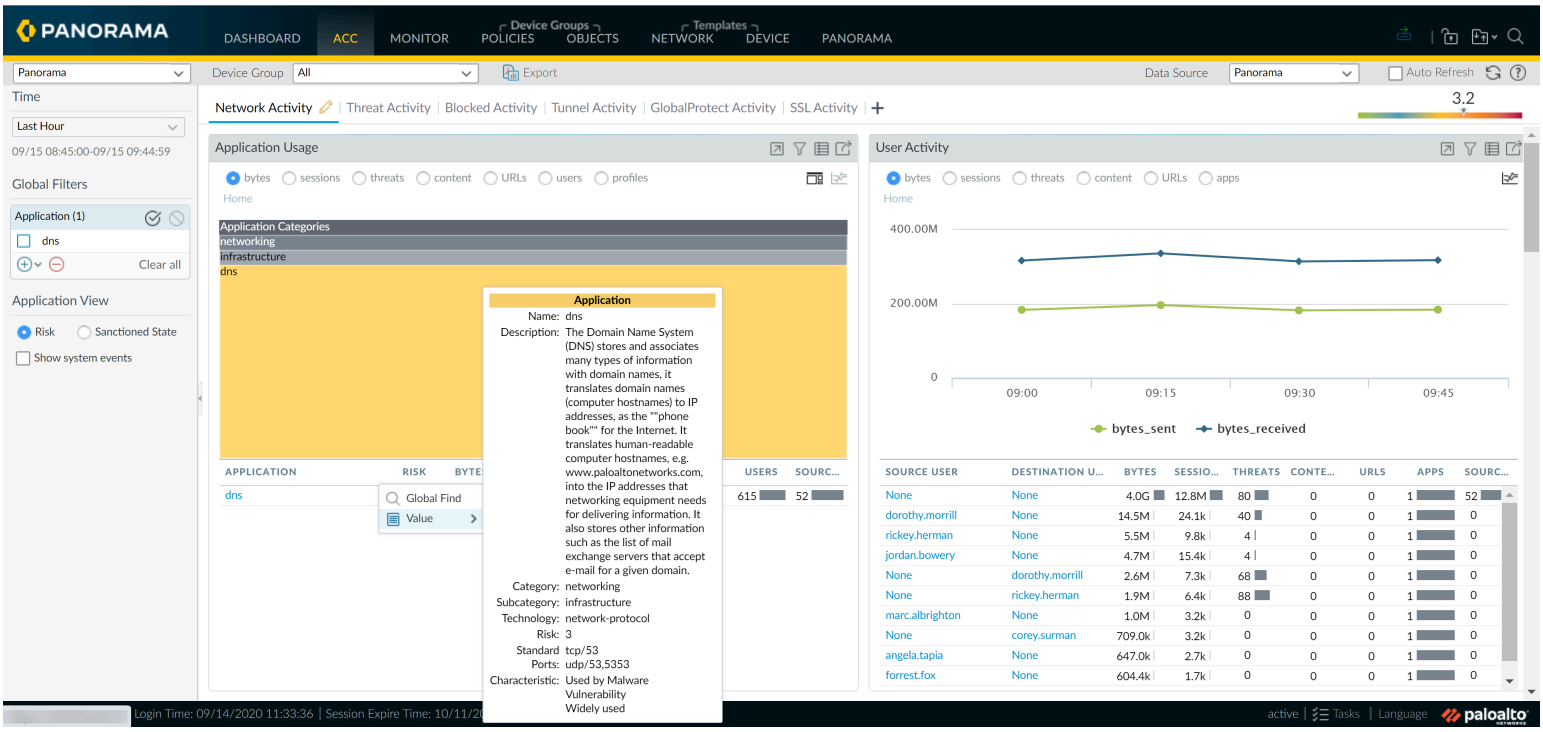


Figure 26: Network Activity Tab

In the **User Activity** widget, you can see how many users are using DNS and the volume of traffic being generated. If you have enabled User-ID, you can view the names of the users who are generating this traffic, and drill in to review all the sessions, content or threats associated with each user.

In the **Threat Activity** tab, view the **Compromised Hosts** widget to see what correlation objects were matched on, and view the match evidence associated with the user and application. You can also view the threat name, category and ID in the **Threat Activity** widget.

With DNS set as a global filter, use the **Destination IP Activity** and the **Destination Regions** widgets to verify where the traffic was destined. You can also view the ingress and egress zones and the security rule that is letting this connection through.

For more detailed information, jump into the Traffic logs for a filtered view and review each log entry for ports used, packets sent, bytes sent and received. Adjust the columns to view more information or less information based on your needs.

The **Monitor** > **App-Scope** > **Traffic Map** tab displays a geographical map of the traffic flow and provides a view of incoming versus outgoing traffic. You can also use the **Monitor** > **App-Scope** > **Change Monitor** tab to view changes in traffic patterns. For example, compare the top applications used over this hour to the last week or month to determine if there is a pattern or trend.

With all the information you have now uncovered, you can evaluate what changes to make to your policy configurations. Here are some suggestions to consider:

- Be restrictive and create a pre-rule on Panorama to block or allow all DNS traffic. Then use Panorama device groups to create and push this policy rule to one or more firewalls.

 ©2025 Palo Alto Networks, Inc.

- Enforce bandwidth use limits and create a QoS profile and policy rule that de-prioritizes non-business traffic. Use Panorama device groups and templates to configure QoS and then push rules to one or more firewalls.
- Schedule a custom report group that pulls together the activity for the specific user and that of top applications used on your network to observe that pattern for another week or two before taking action.

Besides checking for a specific application, you can also check for any unknown applications in the list of top applications. These are applications that did not match a defined App-ID™ signature and display as unknown-udp and unknown-tcp. To delve into these unknown applications, click on the name to drill down to the details for the unclassified traffic.

Use the same process to investigate the top source IP addresses of the hosts that initiated the unknown traffic along with the IP address of the destination host to which the session was established. For unknown traffic, the traffic logs, by default, perform a packet capture (pcap) when an unknown application is detected. The green arrow in the left column represents the packet capture snippet of the application data. Clicking on the green arrow displays the pcap in the browser.

Having the IP addresses of the servers (destination IP), the destination port, and the packet captures, you will be better positioned to identify the application and make a decision on how you would like to take action on your network. For example, you can create a custom application that identifies this traffic instead of labeling it as unknown TCP or UDP traffic. Refer to the article Identifying Unknown Applications for more information on identifying unknown application and Custom Application Signatures for information on developing custom signatures to discern the application.

 ©2025 Palo Alto Networks, Inc.

Use Case: Respond to an Incident Using Panorama

Network threats can originate from different vectors, including malware and spyware infections due to drive-by downloads, phishing attacks, unpatched servers, and random or targeted denial of service (DoS) attacks, to name a few methods of attack. The ability to react to a network attack or infection requires processes and systems that alert the administrator to an attack and provide the necessary forensics evidence to track the source and methods used to launch the attack.

The advantage that Panorama provides is a centralized and consolidated view of the patterns and logs collected from the managed firewalls across your network. You can use the information from the automated correlation engine alone or in conjunction with the reports and logs generated from a Security Information Event Manager (SIEM), to investigate how an attack was triggered and how to prevent future attacks and loss of damage to your network.

The questions that this use case probes are:

- How are you notified of an incident?
- How do you corroborate that the incident is not a false positive?
- What is your immediate course of action?
- How do you use the available information to reconstruct the sequence of events that preceded or followed the triggering event?
- What are the changes you need to consider for securing your network?

This use case traces a specific incident and shows how the visibility tools on Panorama can help you respond to the report.

- Incident Notification
- Review the Widgets in the ACC
- Review Threat Logs
- Review WildFire Logs
- Review Data Filtering Logs
- Update Security Rules

Incident Notification

There are several ways that you could be alerted to an incident depending on how you've configured the Palo Alto Networks firewalls and which third-party tools are available for further analysis. You might receive an email notification that was triggered by a log entry recorded to Panorama or to your syslog server, or you might be informed through a specialized report generated on your SIEM solution, or a third-party paid service or agency might notify you. For this example, let's say that you receive an email notification from Panorama. The email informs you of an event that was triggered by an alert for a Zero Access gent.Gen Command And Control Traffic that matched against a spyware signature. Also listed in the email are the IP address of the source and destination for the session, a threat ID and the timestamp of when the event was logged.

 ©2025 Palo Alto Networks, Inc.

Review the Widgets in the ACC

In the **ACC** > **Threat Activity** tab, check the **Compromised Hosts** widget and **Threat Activity** widget for any critical or high severity threats. In the **Compromised Hosts** widget, look into the Matching Objects and click a Match Count value to view the match evidence for the associated incident.

Review Threat Logs

To begin investigating the alert, use the threat ID to search the Threat logs on Panorama (**Monitor** > **Logs** > **Threat**). From the Threat logs, you can find the IP address of the victim, export the packet capture (PCAP) by clicking the download icon in the log entry, and use a network analyzer tool such as WireShark to review the packet details. In the HTTP case, look for a malformed or bogus HTTP REFERER in the protocol, suspicious host, URL strings, the user agent, the IP address and port in order to validate the incident. Data from these pcaps is also useful in searching for similar data patterns and creating custom signatures or modifying security policy to better address the threat in the future.

As a result of this manual review, if you feel confident about the signature, consider transitioning the signature from an alert action to a block action for a more aggressive approach. In some cases, you may choose to add the attacker IP to an IP block list to prevent further traffic from that IP address from reaching the internal network.

*If you see a DNS-based spyware signature, the IP address of your local DNS server might display as the **Victim IP** address. Often this is because the firewall is located north of the local DNS server, and so DNS queries show the local DNS server as the source IP rather than showing the IP address of the client that originated the request.*

If you see this issue, enable the DNS sinkholing action in the Anti-Spyware profile in security rules to identify the infected hosts on your network. DNS sinkholing allows you to control outbound connections to malicious domains and redirect DNS queries to an internal IP address that is unused; the sinkhole that does not put out a response. When a compromised host initiates a connection to a malicious domain, instead of going out to the internet, the firewall redirects the request to the IP address you defined and it is sinkholed. Now, reviewing the traffic logs for all hosts that connected to the sinkhole allows you locate all compromised hosts and take remedial action to prevent the spread.

To continue with the investigation on the incident, use the information on the attacker and the victim IP address to find out more information, such as:

- Where is the attacker located geographically? Is the IP address an individual IP address or a NATed IP address?
- Was the event caused by a user being tricked into going to a website, a download, or was it sent through an email attachment?
- Is the malware being propagated? Are there other compromised hosts/endpoints on the network?
- Is it a zero-day vulnerability?

The log details for each log entry display the related logs for the event. This information points you to the Traffic, Threat, URL Filtering or other logs that you can review and correlate the events that led to the incident. For example, filter the Traffic log (**Monitor** > **Logs** > **Traffic**) using the IP

 ©2025 Palo Alto Networks, Inc.

address as both the source and the destination IP to get a complete picture of all the external and internal hosts/clients with which this victim IP address has established a connection.

Review WildFire Logs

In addition to the Threat logs, use the victim IP address to filter through the WildFire Submissions logs. The WildFire Submissions logs contain information on files uploaded to the WildFire service for analysis. Because spyware typically embeds itself covertly, reviewing the WildFire Submissions logs tells you whether the victim recently downloaded a suspicious file. The WildFire forensics report displays information on the URL from which the file or .exe was obtained, and the behavior of the content. It informs you if the file is malicious, if it modified registry keys, read/wrote into files, created new files, opened network communication channels, caused application crashes, spawned processes, downloaded files, or exhibited other malicious behavior. Use this information to determine whether to block the application that caused the infection (web-browsing, SMTP, FTP), make more stringent URL Filtering rules, or restrict some applications/actions (for example, file downloads to specific user groups).

Access to the WildFire logs from Panorama requires the following: a WildFire subscription, a File Blocking profile that is attached to a Security rule, and Threat log forwarding to Panorama.

If Panorama will manage firewalls running software versions earlier than PAN-OS 7.0, specify a WildFire server from which Panorama can gather analysis information for WildFire samples that those firewalls submit. Panorama uses the information to complete WildFire Submissions logs that are missing field values introduced in PAN-OS 7.0. Firewalls running earlier releases won't populate those fields. To specify the server, select ***Panorama*** *>* ***Setup*** *>* ***WildFire****, edit the General Settings, and enter the* ***WildFire Private Cloud*** *name. The default is* ***wildfire-public-cloud****, which is the WildFire cloud hosted in the United States.*

If WildFire determines that a file is malicious, a new antivirus signature is created within 24-48 hours and made available to you. If you have a WildFire subscription, the signature is made available within 30-60 minutes as part of the next WildFire signature update. As soon as the Palo Alto Networks next-generation firewall has received a signature for it, if your configuration is configured to block malware, the file will be blocked and the information on the blocked file will be visible in your threat logs. This process is tightly integrated to protect you from this threat and stems the spread of malware on your network.

Review Data Filtering Logs

The Data Filtering log (**Monitor** > **Logs** > **Data Filtering**) is another valuable source for investigating malicious network activity. While you can periodically review the logs for all the files that you are being alerted on, you can also use the logs to trace file and data transfers to or from the victim IP address or user, and verify the direction and flow of traffic: server to client or client to server. To recreate the events that preceded and followed an event, filter the logs for the victim IP address as a destination, and review the logs for network activity.

Because Panorama aggregates information from all managed firewalls, it presents a good overview of all activity in your network. Some of the other visual tools that you can use to survey traffic on your network are the **Threat Map**, **Traffic Map**, and the **Threat Monitor**. The threat map and traffic map (**Monitor** > **AppScope** > **Threat Map** or **Traffic Map**) allow you to visualize

 ©2025 Palo Alto Networks, Inc.

the geographic regions for incoming and outgoing traffic. It is particularly useful for viewing unusual activity that could indicate a possible attack from outside, such as a DDoS attack. If, for example, you do not have many business transactions with Eastern Europe, and the map reveals an abnormal level of traffic to that region, click into the corresponding area of the map to launch and view the ACC information on the top applications, traffic details on the session count, bytes sent and received, top sources and destinations, users or IP addresses, and the severity of the threats detected, if any. The threat monitor (**Monitor** > **AppScope** > **Threat Monitor**) displays the top ten threats on your network, or the list of top attackers or top victims on the network.

Update Security Rules

With all the information you have now uncovered, you can sketch together how the threat impacts your network—the scale of the attack, the source, the compromised hosts, the risk factor—and evaluate what changes, if any, to follow through. Here are some suggestions to consider:

- Forestall DDoS attacks by enhancing your DoS Protection profile to configure random early drop or to drop SYN cookies for TCP floods. Consider placing limits on ICMP and UDP traffic. Evaluate the options available to you based on the trends and patterns you noticed in your logs, and implement the changes using Panorama templates.

  Create a dynamic block list (**Objects** > **Dynamic Block Lists**), to block specific IP addresses that you have uncovered from several intelligence sources: analysis of your own threat logs, DDoS attacks from specific IP addresses, or a third-party IP block list.

  The list must be a text file that is located on a web server. Using device groups on Panorama, push the object to the managed firewalls so that the firewalls can access the web server and import the list at a defined frequency. After creating a dynamic block list object, define a Security rule that uses the address object in the source and destination fields to block traffic from or to the IP address, range, or subnet defined. This approach allows you to block intruders until you resolve the issue and make larger policy changes to secure your network.
- Determine whether to create shared policy rules or device group rules to block specific applications that caused the infection (web-browsing, SMTP, FTP), make more stringent URL Filtering rules, or restrict some applications/actions (for example, file downloads to specific user groups).
- On Panorama, you can also switch to the firewall context and configure the firewall for Botnet reports that identify potential botnet-infected hosts on the network.

 ©2025 Palo Alto Networks, Inc.

 ©2025 Palo Alto Networks, Inc.

Panorama High Availability

To provide redundancy in case of a system or network failure, you can deploy two Panorama™ management servers in a high availability (HA) configuration. Panorama supports an HA configuration in which one peer is the active-primary and the other is the passive-secondary. If a failure occurs on the primary peer, it automatically fails over and the secondary peer becomes active.

- Panorama HA Prerequisites
- Priority and Failover on Panorama in HA
- Failover Triggers
- Logging Considerations in Panorama HA
- Synchronization Between Panorama HA Peers
- Manage a Panorama HA Pair

Panorama HA Prerequisites

To configure Panorama in HA, you require a pair of identical Panorama servers with the following requirements on each:

- **The same form factor**—The peers must be the same model: both M-700 appliances, both M-600 appliances, both M-500 appliances, both M-300 appliances, both M-200 appliances, or both deployed on the same supported hypervisor for Panorama virtual appliances. For example, to successfully configure HA for a Panorama virtual appliance deployed on AWS in Panorama mode, the HA peer must also be deployed on AWS and be in Panorama mode.
- **The same mode**—The peers must be in the same Panorama mode: both running in Panorama mode, Management Only mode, or Legacy mode (ESXi and vCloud Air only).

  Panorama appliances in Log Collector mode do not support HA.
- **The same Panorama OS version**—Must run the same Panorama version to synchronize configuration information and maintain parity for a seamless failover.
- **The same set of licenses**—Must have the same firewall management capacity license.
- (Panorama virtual appliance only) **FIPCS-CC Mode**—FIPS-CC mode must be enabled or disabled on both Panorama HA peers.
- (Panorama virtual appliance only) **Virtual Appliance Resources**—Must have the same number of vCPU cores and memory allocated to successfully synchronize configuration information.
- (Panorama virtual appliance only) **Unique serial number**—Must have unique serial numbers; if the serial number is the same for both Panorama instances, they will be in suspended mode until you resolve the issue.

While it is recommended to match the number of logging disk and the logging disk capacities between the Panorama HA peers, having a different number logging disks or different logging disk capacities between the Panorama HA peers does not impact configuration synchronization or HA failover

.

 ©2025 Palo Alto Networks, Inc.

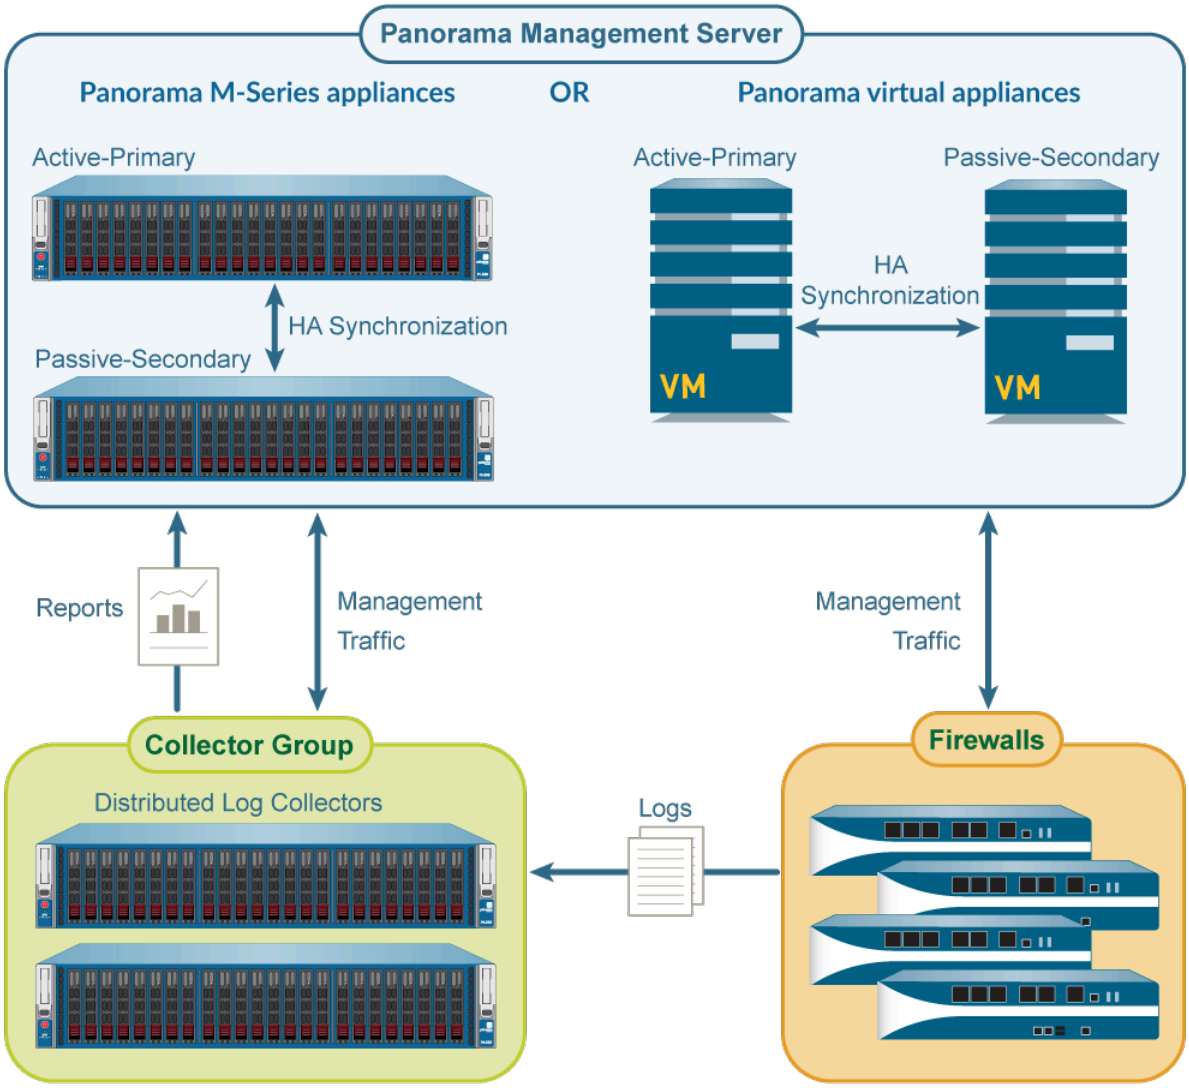


Figure 27: Panorama HA Organization

The Panorama servers in the HA configuration are peers and you can use either (active or passive) to centrally manage the firewalls, Log Collectors, and WildFire appliances and appliance clusters, with a few exceptions (see Synchronization Between Panorama HA Peers). The HA peers use the management (MGT) interface to synchronize the configuration elements pushed to the managed firewalls, Log Collectors, and WildFire appliances and appliance clusters to maintain state information. Typically, Panorama HA peers are geographically located in different sites, so you need to make sure that the MGT interface IP address assigned to each peer is routable through your network. HA connectivity uses TCP port 28 with encryption enabled. If encryption is not enabled, ports 28769 and 28260 are used for HA connectivity and to synchronize configuration between the HA peers. We recommend less than 500ms latency between the peers. To determine the latency, use Ping during a period of normal traffic.

Palo Alto Networks recommends you add at least three Log Collectors to your Collector Groups to avoid the Collector Group becoming inoperable if one Log Collector becomes inaccessible. See Changes to Default Behavior for Collector Groups for more information.

 ©2025 Palo Alto Networks, Inc.

Priority and Failover on Panorama in HA

Each Panorama peer in the HA pair is assigned a *priority* value. The priority value of the primary or secondary peer determines which will be eligible for being the main point of administration and log management. The peer set as primary assumes the active state, and the secondary becomes passive.

Typically, the active peer is responsible for handling all configuration changes and pushing them to the managed firewalls. While it is possible for the passive peer to push configuration changes, this is not recommended as it may lead to issues within the HA setup.

The passive peer is synchronized and ready to transition to the active state if a path, link, system, or network failure occur on the active Panorama.

When a failover occurs, only the state (active or passive) of the Panorama peer changes; the priority (primary and secondary) does not. For example, when the primary peer fails, its status changes from active-primary to passive-primary.

A peer in the active-secondary state can perform all functions with two exceptions:

- It cannot manage firewall or Log Collector deployment functions such as license updates or software upgrades.
- It cannot log to an NFS until you manually change its priority to primary. Only the Panorama virtual appliance in Legacy mode supports NFS.

The following table lists the capabilities of Panorama based on its state and priority settings:

Capability	active-primary	passive-primary passive-secondary	active-secondary
Switch device context	■	■	■
Perform distributed reporting	■	■	■
Manage shared policy	■	■	■
Log to local disk	■	■ (Optional on the Panorama virtual appliance only)	■ (Optional on the Panorama virtual appliance only)
Log to an NFS partition (Panorama virtual appliance only)	■	■	■
Deploy software and licenses	■	■	■
Export Panorama configuration	■	■	■

Figure 28: Panorama HA Capabilities

©2025 Palo Alto Networks, Inc.

For more information, see Panorama HA Prerequisites or Set Up HA on Panorama.

 ©2025 Palo Alto Networks, Inc.

Failover Triggers

When a failure occurs on the active Panorama and the passive Panorama takes over the task of managing the firewalls, the event is called a failover. A failover is triggered when a monitored metric on the active Panorama fails. This failure transitions the state on the primary Panorama from active-primary to passive-primary, and the secondary Panorama becomes active-secondary.

The conditions that trigger a failover are:

- The Panorama peers cannot communicate with each other and the active peer does not respond to health and status polls; the metric used is HA Heartbeat Polling and Hello Messages.

  When the Panorama peers cannot communicate with each other, the active one monitors whether the peers are still connected before a failover is triggered. This check helps in avoiding a failover and causing a split-brain scenario, where both Panorama peers are in an active state.
- One or more of the destinations (IP addresses) specified on the active peer cannot be reached; the metric used is HA Path Monitoring.

In addition to the failover triggers listed above, a failover also occurs when the administrator places the Panorama peer in a suspended state or when preemption occurs. Preemption is a preference for the primary Panorama to resume the active role after recovering from a failure (or user-initiated suspension). By default, preemption is enabled and when the primary Panorama recovers from a failure and becomes available, the secondary Panorama relinquishes control and returns to the passive state. When preemption occurs, the event is logged in the System log.

If you are logging to an NFS datastore, do not disable preemption because it allows the primary peer (that is mounted to the NFS) to resume the active role and write to the NFS datastore. For all other deployments, preemption is only required if you want to make sure that a specific Panorama is the preferred active peer.

HA Heartbeat Polling and Hello Messages

The HA peers use hello messages and heartbeats to verify that the peer is responsive and operational. Hello messages are sent from one peer to the other at the configured Hello Interval to verify the state of the other. The heartbeat is an ICMP ping to the HA peer, and the peer responds to the ping to establish that the peers are connected and responsive. By default, the interval is 1,000 milliseconds for the heartbeat and 8,000ms for hello messages.

HA Path Monitoring

Path monitoring checks the network connectivity and link state for an IP address or group of IP addresses (path group). The active peer uses ICMP pings to verify that one or more destination IP addresses can be reached. For example, you can monitor the availability of interconnected networking devices like a router or a switch, connectivity to a server, or some other vital device that is in the flow of traffic. Make sure that the node/device configured for monitoring is not likely to be unresponsive, especially when it comes under load, as this could cause a path monitoring failure and trigger a failover.

The default ping interval is 5,000ms. An IP address is considered unreachable when three consecutive pings (the default value) fail, and a peer failure is triggered when any or all of the IP

 ©2025 Palo Alto Networks, Inc.

addresses monitored become unreachable. By default, if any one of the IP addresses becomes unreachable, the HA state transitions to non-functional.

 ©2025 Palo Alto Networks, Inc.

Logging Considerations in Panorama HA

Setting up Panorama in an HA configuration provides redundancy for log collection. Because the managed firewalls are connected to both Panorama peers over SSL, when a state change occurs, each Panorama sends a message to the managed firewalls. The firewalls are notified of the Panorama HA state and can forward logs accordingly.

By default, when the managed firewalls cannot connect to Panorama, they buffer the logs; when the connection is restored, they resume sending logs from where it was last left off.

The logging options on the hardware-based Panorama and on the Panorama virtual appliance differ:

- Logging Failover on a Panorama Virtual Appliance in Legacy Mode
- Logging Failover on an M-Series Appliance or Panorama Virtual Appliance in Panorama Mode

Logging Failover on a Panorama Virtual Appliance in Legacy Mode

The Panorama virtual appliance in Legacy mode provides the following log failover options:

Log Storage Type	Description
Virtual disk	By default, the managed firewalls send logs as independent streams to each Panorama HA peer. By default, if a peer becomes unavailable, the managed firewalls buffer the logs and when the peer reconnects it resumes sending logs from where it had left off (subject to disk storage capacity and duration of the disconnection). The maximum log storage capacity depends on the virtual platform (VMware ESXi or vCloud Air); see Panorama Models for details. *You can choose whether to forward logs only to the active peer (see Modify Log Forwarding and Buffering Defaults). However, Panorama does not support log aggregation across the HA pair. Therefore, if you log to a virtual disk, for monitoring and reporting you must query the Panorama peer that collects the logs from the managed firewalls.*
Network File System (NFS)	You can mount NFS storage only to a Panorama virtual appliance that runs on a VMware ESXi server. Only the active-primary Panorama mounts to the NFS-based log partition and can receive logs. On failover, the primary device goes into a passive-primary state. In this scenario, until preemption occurs, the active-secondary Panorama manages the firewalls, but it does not receive the logs and it cannot write to the NFS. To allow the active-secondary peer to log to the NFS, you must manually switch it to primary so that it can mount to the NFS partition. For instructions, see Switch Priority after Panorama Failover to Resume NFS Logging.

 ©2025 Palo Alto Networks, Inc.

Logging Failover on an M-Series Appliance or Panorama Virtual Appliance in Panorama Mode

If you forward firewall logs to the local Log Collectors on an HA pair of M-700 appliances, M-600 appliances, M-500 appliances, M-300 appliances, M-200 appliances, or Panorama virtual appliances in Panorama mode, you specify which firewalls send logs to which Log Collectors when you Configure a Collector Group. You can configure a separate Collector Group for the Log Collector of each Panorama peer or configure a single Collector Group to contain the Log Collectors of both peers. In a Collector Group that contains both local Log Collectors, the log forwarding preference list determines which Log Collector receives logs from firewalls. For all managed firewalls, you have the option to send logs to all the Log Collectors in the Collector Group, in which case Panorama uses round-robin load balancing to select which Log Collector receives the logs at any given moment.

You can enable log redundancy so that each log will have a copy 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 information. Log redundancy is available only if each Log Collector in the Collector Group has the same number of disks.

To utilize log redundancy and ensure logging failover, you must add at least three Log Collectors to a Collector Group to meet the Log Collector n/2+1 quorum requirement introduced in PAN-OS 10.0.

All the Log Collectors for any particular Collector Group must be the same model: all M-200 appliances, all M-300 appliances, all M-500 appliances, all M-600 appliances, all M-700 appliances, or all Panorama virtual appliances in Panorama mode.

Because enabling redundancy creates more logs, this configuration requires more storage capacity. Enabling redundancy doubles the log processing traffic in a Collector Group, which reduces its maximum logging rate by half, as each Log Collector must distribute a copy of each log it receives. (When a Collector Group runs out of space, it deletes older logs.)

 ©2025 Palo Alto Networks, Inc.

Synchronization Between Panorama HA Peers

The Panorama HA peers synchronize the running configuration each time you commit changes on the active Panorama peer. The candidate configuration is synchronized between the peers each time you save the configuration on the active peer or just before a failover occurs.

Settings that are common across the pair, such as shared objects and policy rules, device group objects and rules, template configuration, certificates and SSL/TLS service profiles, and administrative access configuration, are synchronized between the Panorama HA peers.

When you Enable Automated Commit Recovery, HA synchronization occurs only after the firewall successfully tests the connection between itself and Panorama after a push from Panorama.

The settings that are not synchronized are those that are unique to each peer, such as the following:

- Panorama HA configuration—Priority setting, peer IP address, path monitoring groups and IP addresses
- Panorama configuration—Management interface IP address, FQDN settings, login banner, NTP server, time zone, geographic location, DNS server, permitted IP addresses for accessing Panorama, SNMP system settings, and dynamic content update schedules
- Scheduled configuration exports
- NFS partition configuration and all disk quota allocation for logging. This applies only to a Panorama virtual appliance in Legacy mode that runs on a VMware ESXi server
- Disk quota allocation for the different types of logs and databases on the Panorama local storage (SSD)

  If you use a master key to encrypt the private keys and certificates on Panorama, you must use the same master key on both HA peers. If the master keys differ, Panorama cannot synchronize the HA peers.
- Password for the Panorama `admin` administrator

For more information, see Panorama HA Prerequisites or Set Up HA on Panorama.

 ©2025 Palo Alto Networks, Inc.

Manage a Panorama HA Pair

- Set Up HA on Panorama
- Set Up Authentication Using Custom Certificates Between HA Peers
- Test Panorama HA Failover
- Switch Priority after Panorama Failover to Resume NFS Logging
- Restore the Primary Panorama to the Active State

To install software or content updates, see Install Updates for Panorama in an HA Configuration.

Set Up HA on Panorama

Review the Panorama HA Prerequisites before performing the following steps.

If you configure Secure Communication Settings between Panorama HA peers*, the Panorama HA peers use the custom certificate specified for authentication one another. Otherwise, the Panorama HA peers use the predefined certificate for authentication.*

Regardless of how you configure the Panorama HA peers to authenticate communication, neither will impact the ability for the Panorama HA peers to communicate with one another.

STEP 1 | Set up connectivity between the MGT ports on the HA peers.

The Panorama peers communicate with each other using the MGT port. Make sure that the IP addresses you assign to the MGT port on the Panorama servers in the HA pair are routable and that the peers can communicate with each other across your network. To set up the MGT port, see Perform Initial Configuration of the Panorama Virtual Appliance or Perform Initial Configuration of the M-Series Appliance.

Pick a Panorama peer in the pair and complete the remaining tasks.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Enable HA and (optionally) enable encryption for the HA connection.

1. Select **Panorama** > **High Availability** and edit the **Setup** section.
2. Select **Enable HA**.
3. In the **Peer HA IP Address** field, enter the IP address assigned to the peer Panorama.
4. In the **Peer HA Serial** field, enter the serial number of the peer Panorama.

   Entering the Panorama HA peer serial number reduces your attack surface against brute force attacks on the Panorama IP.
5. In the **Monitor Hold Time** field, enter the length of time (milliseconds) that the system will wait before acting on a control link failure (range is 1000-60000, default is 3000).
6. If you do not want encryption, clear the **Encryption Enabled** check box and click **OK**: no more steps are required. If you do want encryption, select the **Encryption Enabled** check box, click **OK**, and perform the following tasks:

   1. Select **Panorama** > **Certificate Management** > **Certificates**.
   2. Select **Export HA key**. Save the HA key to a network location that the peer Panorama can access.
   3. On the peer Panorama, navigate to **Panorama** > **Certificate Management** > **Certificates**, select **Import HA key**, browse to the location where you saved the key, and import it.

STEP 3 | Set the HA priority.

1. In **Panorama** > **High Availability**, edit the **Election Settings** section.
2. Define the **Device Priority** as **Primary** or **Secondary**. Make sure to set one peer as primary and the other as secondary.

   If both peers have the same priority setting, the peer with the higher serial number will be placed in a suspended state.
3. Define the **Preemptive** behavior. By default preemption is enabled. The preemption selection—enabled or disabled—must be the same on both peers.

   If you are using an NFS for logging and you have disabled preemption, to resume logging to the NFS see Switch Priority after Panorama Failover to Resume NFS Logging.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | To configure path monitoring, define one or more path groups.

The path group lists the destination IP addresses (nodes) that Panorama must ping to verify network connectivity.

Perform the following steps for each path group that includes the nodes that you want to monitor.

1. Select **Panorama** > **High Availability** and, in the Path Group section, click **Add**.
2. Enter a **Name** for the path group.
3. Select a **Failure Condition** for this group:
   - **any** triggers a path monitoring failure if any one of the IP addresses becomes unreachable.
   - **all** triggers a path monitoring failure only when none of the IP addresses are reachable.
4. **Add** each destination IP address you want to monitor.
5. Click **OK**. The Path Group section displays the new group.

STEP 5 | (Optional) Select the failure condition for path monitoring on Panorama.

1. Select **Panorama** > **High Availability** and edit the Path Monitoring section.
2. Select a **Failure Condition**:
   - **all** triggers a failover only when all monitored path groups fail.
   - **any** triggers a failover when any monitored path group fails.
3. Click **OK**.

STEP 6 | Commit your configuration changes.

Select **Commit** > **Commit to Panorama** and **Commit** your changes.

STEP 7 | Configure the other Panorama peer.

Repeat Step 2 through Step 6 on the other peer in the HA pair.

STEP 8 | Synchronize the Panorama peers.

1. Access the **Dashboard** on the active Panorama and select **Widgets** > **System** > **High Availability** to display the HA widget.
2. **Sync to peer**, click **Yes**, and wait for the **Running Config** to display **Synchronized**.
3. Access the **Dashboard** on the passive Panorama and select **Widgets** > **System** > **High Availability** to display the HA widget.
4. Verify that the **Running Config** displays **Synchronized**.

STEP 9 | (Optional) Set Up Authentication Using Custom Certificates Between HA Peers.

You must configure the Secure Communication Settings for both Panorama HA peers. Configuring Secure Communication Settings for Panorama in HA configuration does not impact HA connectivity between the HA peers. However, functionality that goes over the Secure Communication link may fail if the Secure Communication Settings are configured

 ©2025 Palo Alto Networks, Inc.

incorrectly, or if the HA peer or managed firewalls do not have the correct certificate, or have an expired certificate.

All traffic on the link established by configuring the Secure Communication Settings is always encrypted.

*If you configure Secure Communication Settings for Panorama in a HA configuration, it is required to **Customize Secure Server Communication** as well. Otherwise, managed firewalls and WildFire appliances are unable to connect to Panorama and PAN-OS functionality is impacted.*

Set Up Authentication Using Custom Certificates Between HA Peers

You can Set Up Authentication Using Custom Certificates for securing the HA connection between Panorama HA peers.

STEP 1 | Generate a certificate authority (CA) certificate on Panorama.

1. Select **Panorama** > **Certificate Management** > **Certificates**.
2. Create a self-signed root CA certificate or import a certificate from your enterprise CA.

STEP 2 | Configure a certificate profile that includes the root CA and intermediate CA.

1. Select **Panorama** > **Certificate Management** > **Certificate Profile**.
2. Configure a certificate profile.

STEP 3 | Configure an SSL/TLS service profile.

1. Select **Panorama** > **Certificate Management** > **SSL/TLS Service Profile**.
2. Configure an SSL/TLS profile to define the certificate and protocol that Panorama and its manage devices use for SSL/TLS services.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure Secure Communication Settings on Panorama on the primary HA peer.

*If you configure Secure Communication Settings on Panorama for Panorama in a HA configuration, it is required to **Customize Secure Server Communication** as well. Otherwise, managed firewalls, Dedicated Log Collectors, and WildFire appliances are unable to connect to Panorama and PAN-OS functionality is impacted.*

1. Select **Panorama** > **Setup** > **Management** and **Edit** the Secure Communication Settings.
2. For the Certificate Type, select **Local**.
3. Select the **Certificate** and **Certificate Profile** you configured in the previous steps.
4. Check (enable) **HA Communication**, **WildFire Communication**, and **Data Redistribution**.
5. Check (enable) **Customize Secure Server Communication**.
6. Select the SSL/TLS service profile from the **SSL/TLS Service Profile** drop-down. This SSL/TLS service profile applies to all SSL connections between Panorama, firewalls, Log Collectors, and Panorama's HA peers.
7. Select the certificate profile from the **Certificate Profile** drop-down.
8. Configure an authorization list.

   When you configure Secure Communication Setting for Panorama in a HA configuration, you are required to add the Panorama HA peer to the authorization list.

   1. Click **Add** under Authorization List.
   2. Select the **Subject** or **Subject Alt Name** as the Identifier type.
   3. Enter the Common Name
9. (Optional) Verify that **Allow Custom Certificate Only** check box is not selected. This allows you to continue managing all devices while migrating to custom certificates.

   *When **Allow Custom Certificate Only** check box is selected, Panorama does not authenticate and cannot manage devices using predefined certificates.*
10. In **Disconnect Wait Time (min)**, enter the number of minutes Panorama should before breaking and reestablishing the connection with its managed devices. This field is blank by default and the range is 0 to 44,640 minutes.

    The disconnect wait time does not begin counting down until you commit the new configuration.
1. Click **OK**.
2. **Commit** and **Commit to Panorama**.
3. Repeat this step on the secondary Panorama HA peer.

   When you configure Secure Communication Settings on the secondary Panorama HA peer, add the primary HA peer to the authorization list as described above.

STEP 5 | Upgrade the client-side Panorama to PAN-OS 10.1.

Upgrade Panorama.

 ©2025 Palo Alto Networks, Inc.

Test Panorama HA Failover

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

STEP 1 | Log in to the active Panorama peer.

You can verify the state of the Panorama server in the bottom right corner of the web interface.

STEP 2 | Suspend the active Panorama peer.

Select **Panorama** > **High Availability**, and then click the **Suspend local Panorama** link in the Operational Commands section.

STEP 3 | Verify that the passive Panorama peer has taken over as active.

On the Panorama **Dashboard**, **High Availability** widget, verify that the state of the **Local** passive server is **active** and the state of the **Peer** is **suspended**.

STEP 4 | Restore the suspended peer to a functional state. Wait for a couple minutes, and then verify that preemption has occurred, if preemptive is enabled.

On the Panorama you previously suspended:

1. Select **Panorama** > **High Availability** and, in the Operational Commands section, click **Make local Panorama functional**.
2. In the **High Availability** widget on the **Dashboard**, confirm that this (Local) Panorama has taken over as the active peer and that the other peer is now in a passive state.

Switch Priority after Panorama Failover to Resume NFS Logging

The Panorama virtual appliance in Legacy mode running on an ESXi server can use an NFS datastore for logging. In an HA configuration, only the primary Panorama peer is mounted to the NFS-based log partition and can write to the NFS. When a failover occurs and the passive Panorama becomes active, its state becomes active-secondary. Although a secondary Panorama peer can actively manage the firewalls, it cannot receive logs or write to the NFS because it does not own the NFS partition. When the firewalls cannot forward logs to the primary Panorama peer, each firewall writes the logs to its local disk. The firewalls maintain a pointer for the last set of log entries that they forwarded to Panorama so that when the passive-primary Panorama becomes available again, they can resume forwarding logs to it.

Use the instructions in this section to manually switch priority on the active-secondary Panorama peer so that it can begin logging to the NFS partition. The typical scenarios in which you might need to trigger this change are as follows:

- Preemption is disabled. By default, preemption is enabled on Panorama and the primary peer resumes as active when it becomes available again. When preemption is disabled, you need to switch the priority on the secondary peer to primary so that it can mount the NFS partition, receive logs from the managed firewalls, and write to the NFS partition.
- The active Panorama fails and cannot recover from the failure in the short term. If you do not switch the priority, when the maximum log storage capacity on the firewall is reached, the oldest logs will be overwritten to enable it to continue logging to its local disk. This situation can lead to loss of logs.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Log in to the currently passive-primary Panorama, select **Panorama** > **Setup** > **Operations** and, in the Device Operations section, click **Shutdown Panorama**.

STEP 2 | Log in to the active-secondary Panorama, select **Panorama** > **High Availability**, edit the Election Settings, and set the **Priority** to **Primary**.

STEP 3 | Click **OK** to save your changes.

STEP 4 | Select **Commit** > **Commit to Panorama** and **Commit** your changes.

Do not reboot when prompted.

STEP 5 | Log in to the Panorama CLI and enter the following command to change the ownership of the NFS partition to this peer: **`request high-availability convert-to-primary`**

STEP 6 | Select **Panorama** > **Setup** > **Operations** and, in the Device Operations section, click **Reboot Panorama**.

STEP 7 | Power on the Panorama peer that you powered off in step 1. This peer will now be in a passive-secondary state.

Restore the Primary Panorama to the Active State

By default, the preemptive capability on Panorama allows the primary Panorama to resume functioning as the active peer as soon as it becomes available. However, if preemption is disabled, the only way to force the primary Panorama to become active after recovering from a failure, a non-functional, or a suspended state, is by suspending the secondary Panorama peer.

Before the active-secondary Panorama goes into a suspended state, it transfers the candidate configuration to the passive Panorama so that all your uncommitted configuration changes are saved and can be accessed on the other peer.

STEP 1 | Suspend Panorama.

1. Log in to the Panorama peer that you want to place in a suspended state.
2. Select **Panorama** > **High Availability**, and click the **Suspend local Panorama** link in the Operational Commands section.

STEP 2 | Verify that the status indicates that the Panorama was suspended at user request.

On the **Dashboard**, **High Availability** widget, verify that the **Local** state is **suspended**.

A failover is triggered when you suspend a peer, and the other Panorama takes over as the active peer.

STEP 3 | Restore the suspended Panorama to a functional state.

1. In the **Panorama** > **High Availability** tab, Operational Commands section, click the **Make local Panorama functional** link.
2. On the **Dashboard**, **High Availability** widget, confirm that the Panorama has transitioned to either the active or passive state.

 ©2025 Palo Alto Networks, Inc.

 ©2025 Palo Alto Networks, Inc.

Administer Panorama

This section describes how to administer and maintain the Panorama™ management server. It includes the following topics:

- Preview, Validate, or Commit Configuration Changes
- Commit Selective Configuration Changes for Managed Devices
- Push Selective Configuration Changes to Managed Devices
- Enable Automated Commit Recovery
- Manage Panorama and Firewall Configuration Backups
- Compare Changes in Panorama Configurations
- Manage Locks for Restricting Configuration Changes
- Add Custom Logos to Panorama
- Use the Panorama Task Manager
- Manage Storage Quotas and Expiration Periods for Logs and Reports
- Monitor Panorama
- Reboot or Shut Down Panorama
- Configure Panorama Password Profiles and Complexity

For instructions on completing initial setup, including defining network access settings, licensing, upgrading the Panorama software version, and setting up administrative access to Panorama, see Set Up Panorama.

Preview, Validate, or Commit Configuration Changes

You can perform Panorama Commit, Validation, and Preview Operations on pending changes to the Panorama configuration and then push those changes to the devices that Panorama manages, including firewalls, Log Collectors, and WildFire appliances and appliance clusters. You can filter the pending changes by administrator or *location* and then commit, push, validate, or preview only those changes. The locations can be specific device groups, templates, Collector Groups, Log Collectors, shared settings, or the Panorama management server.

Because Panorama pushes its running configuration, you cannot push changes to devices until you first commit them to Panorama. If the changes are not ready to activate on devices, you can select **Commit** > **Commit to Panorama** to commit the changes to the Panorama configuration without pushing them to devices. Later, when the changes are ready to activate on devices, you can select **Commit** > **Push to Devices**. If the changes are ready to activate on both Panorama and the devices, select **Commit** > **Commit and Push** as described in the following procedure.



(Device Groups only) When you make a configuration change to a parent device group, pushing these changes to managed firewalls associated with its child device groups is supported only from the Panorama web interface. In this instance, device group configuration changes pushed from the Panorama CLI are pushed only to managed firewalls directly associated with the impacted device group and not to any managed firewalls associated with its child device groups. Child device groups are selected by default when you push a device group configuration change for a parent device group from the Panorama web interface and are not included by default when you push from the Panorama CLI.

For example, you create `ParentDG1` with two firewall associated and its `ChildDG2` with two different managed firewalls associated. You make configuration changes to `ParentDG1`.

*In this scenario, **Push to Devices** and **Commit and Push** from the Panorama web interface successfully push the `ParentDG1` changes to all four firewalls. However, the **`commit-all`** operation from the Panorama CLI pushes only to managed firewalls associated with `ParentDG1`.*

STEP 1 | Configure the scope of configuration changes that you will commit, validate, or preview.

1. Click **Commit** at the top of the web interface.
2. Select one of the following options:
   - **Commit All Changes** (default)—Applies the commit to all changes for which you have administrative privileges. You cannot manually filter the commit scope when you

 ©2025 Palo Alto Networks, Inc.

select this option. Instead, the administrator role assigned to the account you used to log in determines the commit scope.

- **Commit Changes Made By**—Enables you to filter the commit scope by administrator or location. The administrative role assigned to the account you used to log in determines which changes you can filter.

*To commit the changes of other administrators, the account you used to log in must be assigned the Superuser role or an Admin Role profile with the **Commit For Other Admins** privilege enabled.*

3. (Optional) To filter the commit scope by administrator, select **Commit Changes Made By**, click the adjacent link, select the administrators, and click **OK**.
4. (Optional) To filter by location, select **Commit Changes Made By** and clear any changes that you want to exclude from the Commit Scope.

If dependencies between the configuration changes you included and excluded cause a validation error, perform the commit with all the changes included. For example, when you commit changes to a device group, you must include the changes of all administrators who added, deleted, or repositioned rules for the same rulebase in that device group.

STEP 2 | Preview the changes that the commit will activate.

When you preview changes after you delete and then re-add the same device to a policy rule, Panorama displays that same device as both deleted in the running configuration and as added in the candidate configuration. Additionally, the order of devices in the device target list in the running configuration may then be different from the candidate configuration and display as a change when you preview changes even when there aren't any configuration changes.

This can be useful if, for example, you don't remember all your changes and you're not sure you want to activate all of them.

Panorama compares the configurations you selected in the Commit Scope to the running configuration. The preview window displays the configurations side-by-side and uses color coding to indicate which changes are additions (green), modifications (yellow), or deletions (red).

Preview Changes and select the **Lines of Context**, which is the number of lines from the compared configuration files to display before and after the highlighted differences. These

 ©2025 Palo Alto Networks, Inc.

lines help you correlate the preview output to settings in the web interface. Close the preview window when you finish reviewing the changes.

*After you upgrade Panorama to PAN-OS 10.1 or later release, **Preview Changes** shows that HIP Profiles called `source-hip-any` and `destination-hip-any` were added to each Security policy rule for any managed firewall running PAN-OS 9.1 or earlier release instead of `hip-profiles-any`. This is due to a change to the XML file Panorama uses to compare the running and candidate configurations in PAN-OS 10.0 and later releases. You can ignore this error as the push will succeed.*

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

STEP 3 | Preview the individual settings for which you are committing changes.

This can be useful if you want to know details about the changes, such as the types of settings and who changed them.

1. Click **Change Summary**.
2. (Optional) **Group By** a column name (such as the **Type** of setting).
3. **Close** the Change Summary dialog when you finish reviewing the changes.

STEP 4 | Validate the changes before committing to ensure the commit will succeed.

1. **Validate Changes**.

   The results display all the errors and warnings that an actual commit would display.
2. Resolve any errors that the validation results identify.

STEP 5 | (Optional) Modify the Push Scope.

By default, the Push Scope includes all locations with changes that require a Panorama commit.

*If you select **Commit** > **Push to Devices**, the push scope includes all locations associated with devices that are out of sync with the Panorama running configuration.*

1. **No Default Selections** to manually select specific devices. The default devices Panorama pushes to are based on the impacted device group and template configuration changes.
2. **Edit Selections** and select:
   - **Device Groups**—Select device groups or individual firewalls or virtual systems.
   - **Templates**—Select templates, template stacks, or individual firewalls.
   - **Collector Groups**—Select Collector Groups.
   - **Merge with Device Candidate Config**—This setting is enabled by default and merges any pending local firewall configurations with the configuration push from Panorama. The local firewall configuration is merged and committed regardless of the admin

 ©2025 Palo Alto Networks, Inc.

pushing the changes from Panorama or the admin who made the local firewall configuration changes.

Disable this setting if you manage and commit local firewall configuration changes independently of the Panorama managed configuration.

3. Click **OK** to save your changes to the Push Scope.

STEP 6 | Validate the changes you will push to device groups or templates.

1. **Validate Device Group Push** or **Validate Template Push**.

   The results display all the errors and warnings that an actual push operation would display.

2. Resolve any errors that the validation results identify.

STEP 7 | Commit your changes to Panorama and push the changes to devices.

Commit and Push the configuration changes.



Use the Panorama Task Manager *to see details about commits that are pending (optionally, you can cancel these), in progress, completed, or failed.*

STEP 8 | Verify the configuration push from Panorama was successful.

1. Log in to the firewall CLI.
2. Run one of the following commands.

```
admin> show config pushed-template
```

```
admin> show config merged
```

The `show` commands for specific configuration objects are designed to show only the local firewall configuration and not the Panorama-pushed configuration.

For example, the `show network virtual router` command ran on the firewall CLI shows only the virtual router configuration local to the firewall and does not show the Panorama-pushed virtual router configuration.

 ©2025 Palo Alto Networks, Inc.

Commit Selective Configuration Changes for Managed Devices

On the Panorama™ management server, configuration changes occur often and are typically made by multiple administrators who are not aware of what other configuration changes were made on Panorama. It is vital to be able to control which configuration objects are committed to Panorama and prevent incomplete configurations from being pushed from Panorama to your managed firewalls. Rather than committing all pending configuration changes to Panorama, you can instead select specific device group and template stack objects to commit. A system log is generated after a successful selective commit.

The ability to select specific objects to commit allows multiple administrators to effectively make configuration changes without disrupting other administrators who make configuration changes that are not ready to be committed. Leveraging the ability to selectively commit configuration changes to Panorama allows you to maintain your defined operational procedure while still being able to successfully make independent configuration changes that are not defined within your operational scope.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Perform device group and template stack configuration changes on Panorama.

STEP 3 | **Commit** and **Commit to Panorama**.

STEP 4 | Change the commit scope to **Commit Changes Made By** to select specific device group and template stack configuration changes to commit to Panorama.

The push scope displays the name of the admin currently logged in. Click the admin name to view a list of admins who have made configuration changes that have not been committed to Panorama.

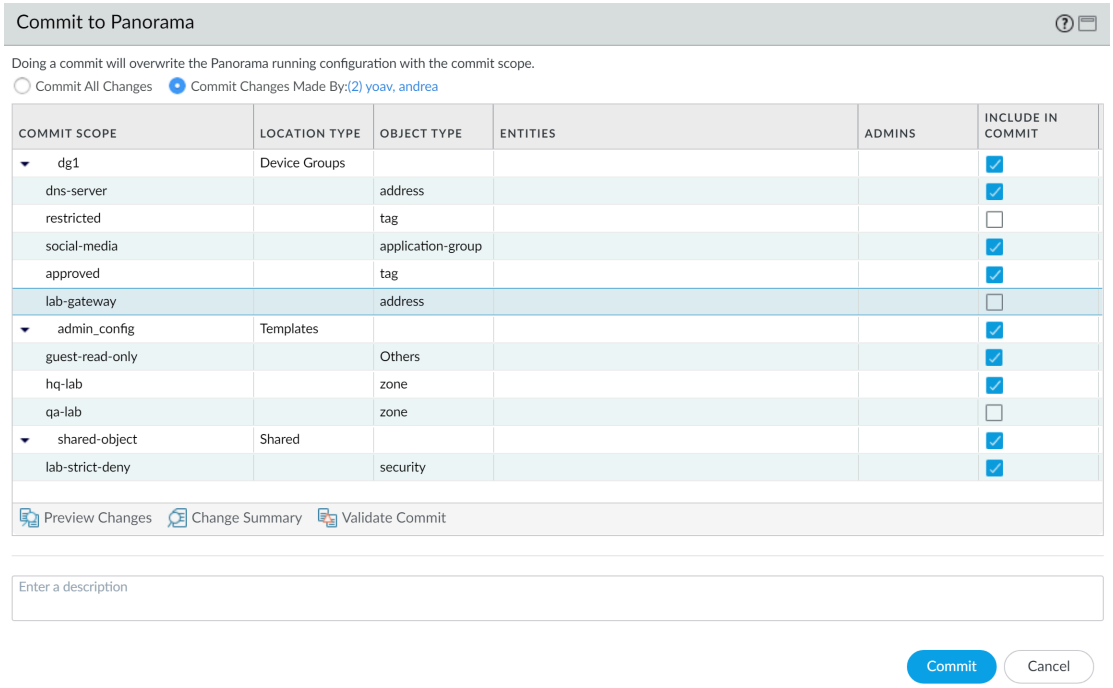
STEP 5 | In the Include in Commit column, check (enable) the configuration objects you want to include in the commit.

STEP 6 | (Optional) Preview and validate your pending configuration changes to ensure you want to commit the selective configuration changes to Panorama.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | **Commit**.

The Commit Status page displays the administrators that made configuration changes that were committed and the location of the committed configuration changes.



STEP 8 | Push Selective Configuration Changes to Managed Devices.

After you commit the selected configuration objects to Panorama, you can push those configuration objects to your managed firewalls.

 ©2025 Palo Alto Networks, Inc.

Push Selective Configuration Changes to Managed Devices

You can include the configuration changes committed by one or more Panorama administrators to push to your managed firewalls. This allows for a greater degree of control when making configuration changes and reduces the risk of pushing an incomplete configuration to your managed firewalls. To allow a Panorama administrator to selectively push configuration changes, you must configure an admin role profile that allows selective push and assign the admin role profile to the Panorama administrator. A system log is generated for a successful selective push to managed firewalls.

You can also leverage selective commit of configuration changes for further selectivity when pushing configuration changes to your managed firewalls. Selective commit allows you to select and commit specific configuration objects. After you commit, you can leverage selective push to review and push all committed configuration changes made by other Panorama administrators.

The ability specify which Panorama administrator configuration changes to include in a push to managed firewalls allows multiple administrators to effectively manage firewall configurations without disrupting other administrators and reduces the risk of pushing an incomplete configuration to your managed firewalls that could result in an outage. Leveraging the ability to selectively push configuration changes allows you to maintain your defined operational procedure while still being able to successfully make independent configuration changes that are not defined within your operational scope.

Selective push is supported for managed firewalls only and is supported for managed firewalls running any supported PAN-OS release. Selective push is not supported for Log Collectors, collector groups, WildFire appliances, and WildFire clusters. For Panorama in an active/passive high availability (HA) configuration, selective push is supported from the active HA peer only.

STEP 1 | Log in to the Panorama Web Interface.

The Panorama administrator must be configured with an admin role profile that allows the push of configuration changes made by other admins to managed firewalls. The default Superuser or Panorama admin role privileges support full object level configuration privileges.

STEP 2 | Select **Commit** and **Push to Devices**.

*You can also select **Commit and Push** to commit selective configuration changes to Panorama and push already committed changes in one operation.*

You cannot selectively push a configuration change that has not been committed.

STEP 3 | Change the push scope to **Push Changes Made By** and filter the push scope by Panorama admin to select specific device group and template stack configuration changes to push to your managed firewalls.

The push scope displays the name of the admin currently logged in. Click the admin name to view a list of admins with committed configuration changes that have not been pushed to

 ©2025 Palo Alto Networks, Inc.

managed firewalls. The push scope automatically refreshes to display an updated list of device groups and template stacks based on the admins selected.

STEP 4 | In the Include in Push column, check (enable) the configuration objects you want to include in the commit.

The push scope displays only device groups and template stacks that are `out of sync`.

You must select and push the entire device group or template stack configuration that was committed. Object level changes displayed in the push scope are informational and cannot be excluded from the push for the device group or template stack you select.

STEP 5 | (Optional) **Edit Selections** and select the managed firewalls associated with the impacted device groups and template stacks.

Skip this step to push to all managed firewalls associated with the impacted device groups and template stacks.

*Disable the **Merge with Device Candidate Config** setting if you manage and commit local firewall configuration changes independently of the Panorama managed configuration.*

This setting is enabled by default and merges any pending local firewall configurations with the configuration push from Panorama. The local firewall configuration is merged and committed regardless of the admin pushing the changes from Panorama or the admin who made the local firewall configuration changes.

STEP 6 | **Push** the configuration changes.

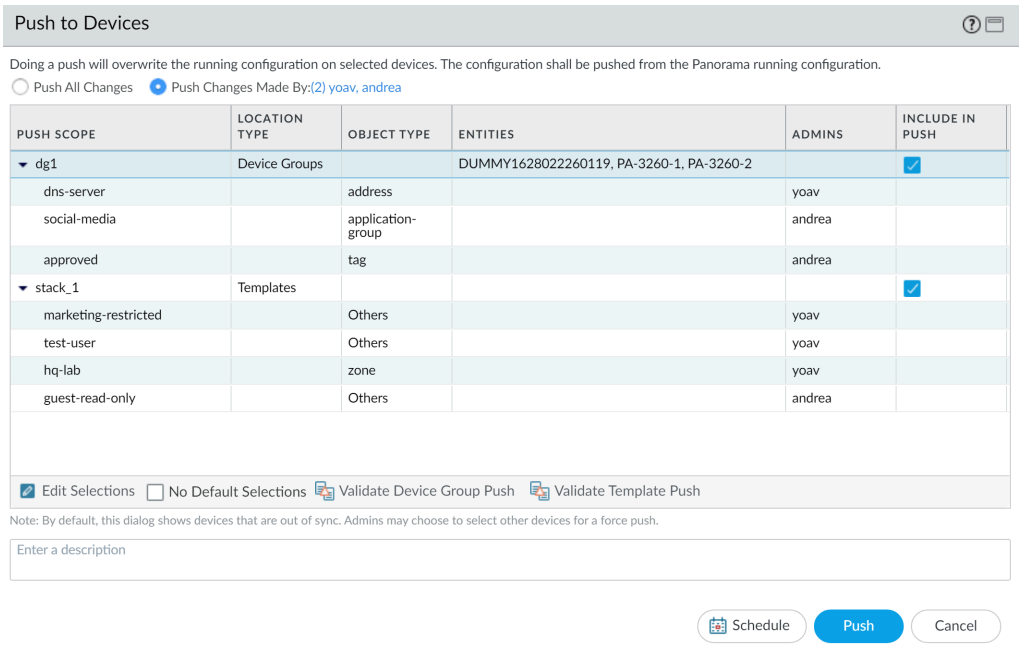
STEP 7 | If your admin role allows you to push configuration changes for other Panorama administrators, review the Confirm Push to Devices prompt and **Push**.

This warning is displayed when the administrators included in the Admin Scope make conflicting configuration changes to the same object. For example, Admin1 is allowed to push configuration changes to managed firewalls while Admin2 is not allowed. Admin1 creates `SecurityRule`, adds `ZoneA` as the source zone and commits the change. Admin2 then modifies `SecurityRule`, deletes `ZoneA`, adds `ZoneB`, and as well as making additional configuration changes. Admin2 commits the changes to Panorama. Admin1 wants to include the configuration changes made by Admin1 in the push to managed firewalls. In this scenario,

 ©2025 Palo Alto Networks, Inc.

Admin1 is prompted to confirm the push because the configuration changes made to `SecurityRule` conflict.

*If you are not confident in the configuration changes made by other Panorama administrators, **Continue push with my selected changes only** to only push your own configuration changes and overwrite any configuration object conflict with the changes you made.*



STEP 8 | Select **Panorama** > **Managed Devices** > **Summary** and click the Template Last Commit State for the impacted firewalls to review the Last Push State Details.

 ©2025 Palo Alto Networks, Inc.

Enable Automated Commit Recovery

To ensure that broken configurations caused by configuration changes pushed from the Panorama™ management server to managed firewalls, or committed locally on the firewall, enable **Automated Commit Recovery** to enable managed firewalls to test configuration changes for each commit and to verify that the changes did not break the connection between Panorama and the managed firewall. You can configure the number of tests that each managed firewall performs and the interval at which each test occurs before the managed firewall automatically reverts its configuration back to the previous running configuration. When you enable automated commit recovery, the managed firewall configuration reverts and not the Panorama configuration. Additionally, the managed firewall tests its connection to Panorama every 60 minutes to ensure continued communication in the event unrelated network configuration changed disrupted connectivity between the firewall and Panorama or if impacts from a past committed configuration affected connectivity. For high availability (HA) configurations, HA synchronization between the HA peers after a push from Panorama occurs only after a connectivity test.

Automated commit recovery is enabled by default. However, if you disabled automated commit recovery and then want to re-enable this feature in an existing production environment, first verify that there are no policy rules that will break the connection between Panorama and the managed firewall. For example, in the event where management traffic traverses the dataplane, it is possible there is a policy rule that restricts traffic from the firewall to Panorama.

The firewall generates a config log after the firewall configuration successfully reverts to the last running configuration. Additionally, the firewall generates a system log when the administrator disables this feature, when a configuration revert event begins due to a connectivity test that fails after a configuration push, and when the Panorama connectivity test that is performed every 60 minutes fails and causes the firewall configuration to revert.

*Enable **Automated Commit Recovery** independent of any other configuration change. If enabled alongside any other configuration changes that result in a connection break between Panorama and managed firewalls, the firewall configuration cannot automatically revert.*

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Select **Device** > **Setup** > **Management** and select the desired Template or Template Stack from the **Template** context drop-down.

 ©2025 Palo Alto Networks, Inc.

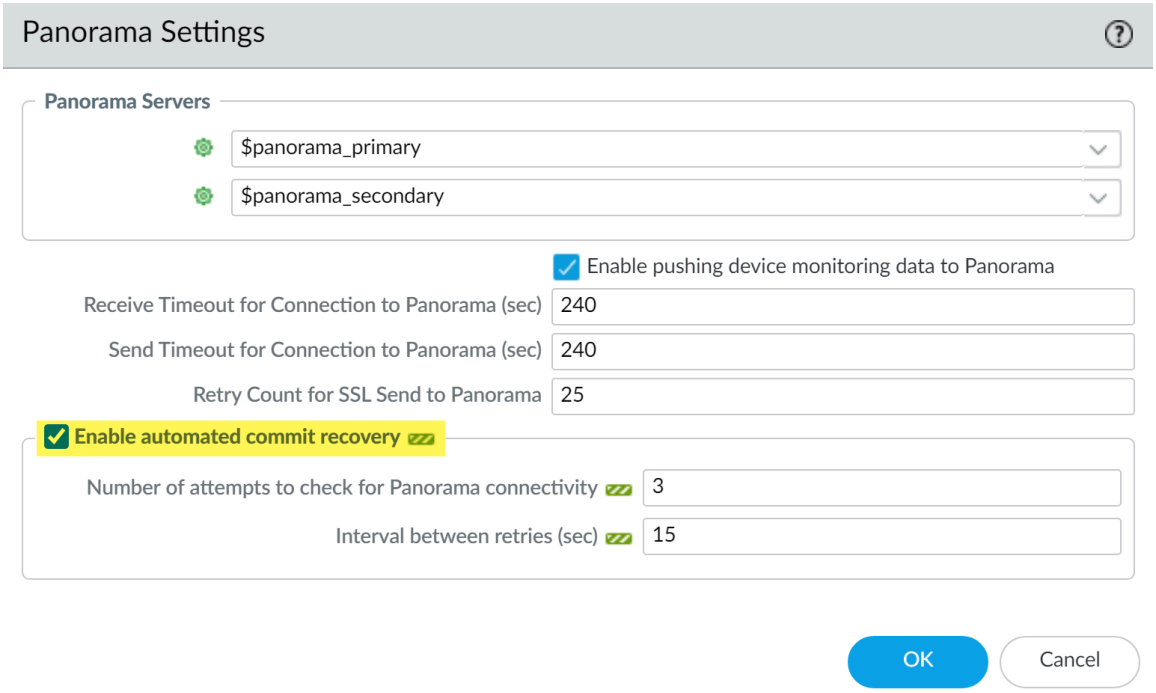
STEP 3 | Enable automated commit recovery.

*(ZTP Firewalls) Enabling automated commit recovery may cause the initial configuration push after you add ZTP firewalls to Panorama to be automatically reverted. To enable automated commit recovery for your managed ZTP firewalls, configure the **Number of attempts to check for Panorama connectivity** as **5**.*

1. **Edit** ( ) the Panorama Settings.
2. **Enable automated commit recovery**.
3. Configure the **Number of attempts to check for Panorama connectivity** (default is 1 attempt).

(ZTP Firewalls) Configure the number of attempts as **5** to avoid unintended configuration revets after the first push from Panorama.

1. Configure the **Interval between retries** (default is 10 seconds).
2. Click **OK** to save your changes.



STEP 4 | **Commit** > **Commit and Push** and **Commit and Push** your changes.

STEP 5 | Verify that the automated commit recovery feature is enabled on your managed firewalls.

1. Launch the Firewall Web Interface.
2. Select **Device** > **Setup** > **Management** and, in the Panorama Settings, verify that **Enable automated commit recovery** is enabled (checked).

 ©2025 Palo Alto Networks, Inc.

Manage Panorama and Firewall Configuration Backups

The running configuration on Panorama comprises all the settings that you have committed and that are therefore active. The candidate configuration is a copy of the running configuration plus any inactive changes that you made since the last commit. Saving backup versions of the running or candidate configuration enables you to later restore those versions. For example, if a commit validation shows that the current candidate configuration has more errors than you want to fix, you can restore a previous candidate configuration. You can also revert to the current running configuration without saving a backup first.



See Panorama Commit, Validation, and Preview Operations *for more information on committing configuration changes to Panorama and pushing the changes to managed devices.*

After a commit on a local firewall that runs PAN-OS 5.0 or later, a backup is sent of its running configuration to Panorama. Any commits performed on the local firewall will trigger the backup, including commits that an administrator performs locally on the firewall or automatic commits that PAN-OS initiates (such as an FQDN refresh). By default, Panorama stores up to 100 backups for each firewall, though this is configurable. To store Panorama and firewall configuration backups on an external host, you can schedule exports from Panorama or export on demand. You can also import configurations from firewalls into Panorama device groups and templates to Transition a Firewall to Panorama Management.

(VMware ESXi and vCloud Air only) VMware snapshot functionality is not supported for a Panorama virtual appliance deployed on VMware ESXi and vCloud Air. Taking snapshots of a Panorama virtual appliance can impact performance, result in intermittent and inconsistent packet loss, and Panorama may become unresponsive. Additionally, you may lose access to the Panorama CLI and web interface and switching to Panorama mode is not supported. Instead, save and export your named configuration snapshot to any network location.



If you are leveraging Enterprise data loss prevention (DLP)*, loading a Panorama configuration backup that does not contain the Shared Enterprise DLP configuration objects removes these Shared objects required for Enterprise DLP functionality.*

- Schedule Export of Configuration Files
- Save and Export Panorama and Firewall Configurations
- Revert Panorama Configuration Changes
- Configure the Maximum Number of Configuration Backups on Panorama
- Load a Configuration Backup on a Managed Firewall

Schedule Export of Configuration Files

Panorama saves a backup of its running configuration as well as the running configurations of all managed firewalls. The backups are in XML format with file names that are based on serial numbers (of Panorama or the firewalls). Use these instructions to schedule daily exports of the backups to a remote host. Panorama exports the backups as a single gzip file. You require superuser privileges to schedule the export.

 ©2025 Palo Alto Networks, Inc.

If Panorama has a high availability (HA) configuration, you must perform these instructions on each peer to ensure the scheduled exports continue after a failover. Panorama does not synchronize scheduled configuration exports between HA peers.

To export backups on demand, see Save and Export Panorama and Firewall Configurations.

STEP 1 | (RHEL Server version 8.3 only) Verify that for your RHEL server running version 8.3, set the `ChallengeResponseAuthentication` setting is **no** within the `sshd_config` file.

Update to **no** if needed and then restart the SSH daemon. This setting is required to export configuration files to your RHEL server running version 8.3.

STEP 2 | Select **Panorama** > **Scheduled Config Export** and click **Add**.

STEP 3 | Enter a **Name** and **Description** for the scheduled file export and **Enable** it.

STEP 4 | Using the 24-hour clock format, enter a daily **Scheduled Export Start Time** or select one from the drop-down.

If you are configuring a scheduled export to two or more servers, stagger the start time of the scheduled exports. Scheduling multiple exports at the same start time results in discrepancies between the exported configurations.

STEP 5 | Set the export **Protocol** to Secure Copy (**SCP**) or File Transfer Protocol (**FTP**).

*Only **FTP** is supported when exporting to devices running Windows.*

STEP 6 | Enter the details for accessing the server, including: **Hostname** or IP address, **Port**, **Path** for uploading the file, **Username**, and **Password**.

The **Path** supports the following characters: **.**(period), **+**, **{** and **}**, **/**, **-**, **_**, **0**-**9**, **a**-**z**, and **A**-**Z**. Spaces are not supported in the file **Path**.

*If you are exporting to an FTP server using an IPv6 address as the **Hostname**, you must enter the address enclosed in square brackets ([ ]). For example, **[2001:0db8:0000:0000:0000:8a2e:0370:7334]**.*

*If you are exporting to a BSD server, you will need to modify the SSHD password prompt to **<username>@<hostname> <password>:** .*

STEP 7 | (SCP only) Click **Test SCP server connection**. To enable the secure transfer of data, you must verify and accept the host key of the SCP server. Panorama doesn't establish the connection until you accept the host key. If Panorama has an HA configuration, perform this step on each HA peer so that each one accepts the host key of the SCP server. If Panorama can

 ©2025 Palo Alto Networks, Inc.

successfully connect to the SCP server, it creates and uploads the test file named `ssh-export-test.txt`.

*(PAN-OS 10.2.4 and later releases) A pop-up window is displayed requiring you to enter a clear text **Password**and then to **Confirm Password** in order to test the SCP server connection and enable the secure transfer of data.*

Panorama does not establish and test the SCP server connection until you enter and confirm the SCP server password. If Panorama has an HA configuration, perform this step on each HA peer so that each one can successfully connect to the SCP server. If Panorama can successfully connect to the SCP server, it creates and uploads the test file named `ssh-export-test.txt`.

STEP 8 | Click **OK** to save your changes.

STEP 9 | Select **Commit** > **Commit to Panorama** and **Commit** your changes.

Save and Export Panorama and Firewall Configurations

Saving a backup of the candidate configuration to persistent storage on Panorama enables you to later restore that backup (see Revert Panorama Configuration Changes). Additionally, Panorama allows you to save and export the device group, template, and template stack configurations that you specify. This is useful for preserving changes that would otherwise be lost if a system event or administrator action causes Panorama to reboot. After rebooting, Panorama automatically reverts to the current version of the running configuration, which Panorama stores in a file named `running-config.xml`. Saving backups is also useful if you want to revert to a Panorama configuration that is earlier than the current running configuration. Panorama does not automatically save the candidate configuration to persistent storage. You must manually save the candidate configuration as a default snapshot file (`.snapshot.xml`) or as a custom-named snapshot file. Panorama stores the snapshot file locally but you can export it to an external host.

*You don't have to save a configuration backup to revert the changes made since the last commit or reboot; just select **Config** > **Revert Changes** (see Revert Panorama Configuration Changes).*

Palo Alto Networks recommends that you back up any important configurations to an external host.

STEP 1 | Save changes to the candidate configuration.

- To overwrite the default snapshot file (`.snapshot.xml`) with all the changes that all administrators made, perform one of the following steps:
  - Select **Panorama** > **Setup** > **Operations** and **Save candidate Panorama configuration**.
  - Log in to Panorama with an administrative account that is assigned the Superuser role or an Admin Role profile with the **Save For Other Admins** privilege enabled. Then select

 ©2025 Palo Alto Networks, Inc.

Config > **Save Changes** at the top of the web interface, select **Save All Changes** and **Save**.

- To overwrite the default snapshot (`.snapshot.xml`) with changes made by administrators to specific device group, template, or template stack configurations:
  1. Select **Panorama** > **Setup** > **Operations**, **Save candidate Panorama configuration**, and **Select Device Group & Templates**.
  2. Select the specific device groups, templates, or template stacks to revert.
  3. Click **OK** to confirm the operation.
  4. (Optional) Select **Commit** > **Commit to Panorama** and **Commit** your changes to overwrite the running configuration with the snapshot.
- To create a snapshot that includes all the changes that all administrators made but without overwriting the default snapshot file:
  1. Select **Panorama** > **Setup** > **Operations** and **Save named Panorama configuration snapshot**.
  2. Specify the **Name** of a new or existing configuration file.
  3. Click **OK** and **Close**.
- To save only specific changes to the candidate configuration without overwriting any part of the default snapshot file:
  1. Log in to Panorama with an administrative account that has the role privileges required to save the desired changes.
  2. Select **Config** > **Save Changes** at the top of the web interface.
  3. Select **Save Changes Made By**.
  4. To filter the Save Scope by administrator, click ***<administrator-name>***, select the administrators, and click **OK**.
  5. To filter the Save Scope by location, clear any locations that you want to exclude. The locations can be specific device groups, templates, Collector Groups, Log Collectors, shared settings, or the Panorama management server.
  6. Click **Save**, specify the **Name** of a new or existing configuration file, and click **OK**.
- To save a specific device group, template, or template stack configuration:
  1. Select **Panorama** > **Setup** > **Operations**, **Save named Panorama configuration snapshot**, and **Select Device Group & Templates**.
  2. Select the specific device groups, templates, or template stacks to save.
  3. Click **OK** to confirm the operation.

STEP 2 | Export a candidate or running configuration to a host external to Panorama or to a firewall.

You can schedule daily exports to an SCP or FTP server (see Schedule Export of Configuration Files) or export configurations on demand. To export on demand, select **Panorama** > **Setup** > **Operations** and select one of the following options:

- **Export named Panorama configuration snapshot**—Export the current running configuration, a named candidate configuration snapshot, or a previously imported configuration (candidate or running). Panorama exports the configuration as an XML file with the **Name**

 ©2025 Palo Alto Networks, Inc.

you specify. **Select Device Group & Templates** to specify the device group, template, or template stack configurations to export.

- **Export Panorama configuration version**—Select a **Version** of the running configuration to export as an XML file. **Select Device Group & Templates** to specify the device group, template, or template stack configurations to export as an XML file.
- **Export Panorama and devices config bundle**—Generate and export the latest version of the running configuration backup of Panorama and of each managed firewall. To automate the process of creating and exporting the configuration bundle daily to a Secure Copy (SCP) or FTP server, see Schedule Export of Configuration Files.
- **Export or push device config bundle**—After you import a firewall configuration into Panorama, Panorama creates a firewall configuration bundle named <firewall_name>_import.tgz, in which all local policies and objects are removed. You can then **Export or push device config bundle** to perform one of the following actions:
  - **Push & Commit** the configuration bundle to the firewall to remove any local configuration from it, enabling you to manage the firewall from Panorama.
  - **Export** the configuration to the firewall without loading it. When you are ready to load the configuration, log in to the firewall CLI and run the configuration mode command **`load device-state`**. This command cleans the firewall in the same way as the **Push & Commit** option.

The full procedure to Transition a Firewall to Panorama Management requires additional steps.

Revert Panorama Configuration Changes

When you revert changes, you are replacing settings in the current candidate configuration with settings from another configuration. Reverting changes is useful when you want to undo changes to multiple settings as a single operation instead of manually reconfiguring each setting.

You can revert pending changes that were made to the Panorama configuration since the last commit. You can revert all pending changes on Panorama or select specific device groups, templates, or template stacks. Panorama provides the option to filter the pending changes by administrator or location. The locations can be specific device groups, templates, Collector Groups, Log Collectors, shared settings, or the Panorama management server. If you saved a snapshot file for a candidate configuration that is earlier than the current running configuration (see Save and Export Panorama and Firewall Configurations), you can also revert to that candidate configuration snapshot. Reverting to a snapshot enables you to restore a candidate configuration that existed before the last commit. Panorama automatically saves a new version of the running configuration whenever you commit changes and you can restore any of those versions.

Reverting a Panorama management server configuration requires a full commit and must be performed by a superuser. Full commits are required when performing certain Panorama operations, such as reverting and loading a Panorama configuration, and are not supported for custom Admin Role profiles.

 ©2025 Palo Alto Networks, Inc.

- Revert to the current Panorama running configuration (file named `running-config.xml`).

  This operation undoes changes you made to the candidate configuration since the last commit.

  - To revert all the changes that all administrators made, perform one of the following steps:
    - Select **Panorama** > **Setup** > **Operations**, **Revert to running Panorama configuration**, and click **Yes** to confirm the operation.
    - Log in to Panorama with an administrative account that is assigned the Superuser role or an Admin Role profile with the **Commit For Other Admins** privilege enabled. Then select **Config** > **Revert Changes**, select **Revert All Changes**, and **Revert**.
  - To revert only specific changes to the candidate configuration:
    1. Log in to Panorama with an administrative account that has the role privileges required to revert the desired changes.

       The privileges that control commit operations also control revert operations.
    2. Select **Config** > **Revert Changes**.
    3. Select **Revert Changes Made By**.
    4. To filter the Revert Scope by administrator, click <*administrator-name*>, select the administrators, and click **OK**.
    5. To filter the Revert Scope by location, clear any locations that you want to exclude.
    6. **Revert** the changes.
  - To revert specific device group, template, or template stack changes to the running configuration:
    1. Select **Panorama** > **Setup** > **Operations**, **Revert to running Panorama configuration**, and **Select Device Group & Templates**.
    2. Select the specific device groups, templates, or template stacks to revert.
    3. Click **OK** to confirm the operation.
    4. (Optional) Select **Commit** > **Commit to Panorama** and **Commit** your changes to overwrite the running configuration.

 ©2025 Palo Alto Networks, Inc.

- Revert to the default snapshot (`.snapshots.xml`) of the Panorama candidate configuration.
  - To revert all the changes that all administrators made:
    1. Select **Panorama** > **Setup** > **Operations** and **Revert to last saved Panorama configuration**.
    2. Click **Yes** to confirm the operation.
    3. (Optional) Select **Commit** > **Commit to Panorama** and **Commit** your changes to overwrite the running configuration with the snapshot.
  - To revert specific device group, template, or template stack changes to the running configuration:
    1. Select **Panorama** > **Setup** > **Operations**, **Revert to last saved Panorama configuration**, and **Select Device Group & Templates**.
    2. Select the specific device groups, templates, or template stacks to revert.
    3. Click **OK** to confirm the operation.
    4. (Optional) To overwrite the running configuration, select **Commit** > **Commit to Panorama** and **Commit** your changes with the snapshot.

- Revert to a previous version of the running configuration that is stored on Panorama.
  - To revert all changes that administrators made:
    1. Select **Panorama** > **Setup** > **Operations**, **Load Panorama configuration version**, and **Select Device Group & Templates**.
    2. Select a configuration **Version** and click **OK**.
    3. (Optional) To overwrite the running configuration with the version you just restored, select **Commit** > **Commit to Panorama** and **Commit** your changes.
  - To revert specific device group, template, or template changes to the running configuration:
    1. Select **Panorama** > **Setup** > **Operations**, **Load Panorama configuration version**, and select a configuration version **Name**.
    2. **Select Device Group & Templates** and select the specific device groups, templates, or template stacks to revert.
    3. Click **OK** to confirm the operation.
    4. (Optional) To overwrite the running configuration with the snapshot, select **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

- Revert to one of the following:
  - Custom-named version of the Panorama running configuration that you previously imported.
  - Custom-named Panorama candidate configuration snapshot (instead of the default snapshot).
    1. Select **Panorama** > **Setup** > **Operations**, **Load named Panorama configuration snapshot**, and select the **Name** of the configuration file you just imported.
    2. (Optional) **Load Shared Objects** or **Load Shared Policies** to load all shared objects or policies. You can load all shared objects and policies, as well as load all objects and policies configured in the device groups and templates you specify in the next step.
    3. (Optional) **Select Device Group & Templates**, and select the specific device group, template, or template stack configurations to load. Skip this step if you want to revert the entire Panorama configuration.
    4. Click **OK** to confirm the operation.
    5. (Optional) To overwrite the running configuration with the snapshot, select **Commit** > **Commit to Panorama** and **Commit** your changes.

- Restore a Panorama running or candidate configuration that you previously exported to an external host.
  1. Select **Panorama** > **Setup** > **Operations**, **Import named Panorama configuration snapshot**, **Browse** to the configuration file on the external host, and click **OK**.
  2. **Load named Panorama configuration snapshot** and select the **Name** of the configuration file you just imported.
  3. (Optional) **Load Shared Objects** or **Load Shared Policies** to load all shared objects or policies. You can load all shared objects and policies, as well as load all objects and or policies configured in the device groups and templates you specify in the next step.
  4. (Optional) **Select Device Group & Templates** and select the specific device group, template, or template stack configurations to load. Skip this step if you want to revert the entire Panorama configuration.
  5. Click **OK** to confirm the operation.
  6. (Optional) To overwrite the running configuration with the snapshot you just imported, select **Commit** > **Commit to Panorama** and **Commit** your changes.

Configure the Maximum Number of Configuration Backups on Panorama

STEP 1 | Select **Panorama** > **Setup** > **Management** and edit the Logging and Reporting Settings.

STEP 2 | Select **Log Export and Reporting** and enter the **Number of Versions for Config Backups** (default is 100; range is 1 to 1,048,576).

STEP 3 | Click **OK** to save your changes.

STEP 4 | Select **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

Load a Configuration Backup on a Managed Firewall

Use Panorama to load a configuration backup on a managed firewall. You can choose to revert to a previously saved or committed configuration on the firewall. Panorama pushes the selected version to the managed firewall, thereby overwriting the current candidate configuration on the firewall.

STEP 1 | Select **Panorama** > **Managed Devices** > **Summary**.

STEP 2 | Select **Manage** in the Backups column.

STEP 3 | Select from the Saved Configurations or Committed Configurations.

- Click a version number to view the contents of that version.
- **Load** a configuration version.

STEP 4 | Log in to the firewall web interface and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

Compare Changes in Panorama Configurations

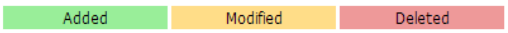
To compare configuration changes on Panorama, you can select any two sets of configuration files: the candidate configuration, the running configuration, or any other configuration version that has been previously saved or committed on Panorama. The side-by-side comparison enables you to:

- Preview the configuration changes before committing them to Panorama. You can, for example, preview the changes between the candidate configuration and the running configuration. As a best practice, select the older version on the left pane and the newer version on the right pane, to easily compare and identify modifications.
- Perform a *configuration audit* to review and compare the changes between two sets of configuration files.

Device Group and Template admins can only compare configurations for device groups and templates within their access domains.

- Compare changes in Panorama configurations.
  1. Select **Panorama** > **Config Audit**.
  2. In each drop-down, select a configuration for the comparison.
  3. Select the number of lines that you want to include for **Context** and click **Go**.

     Panorama uses color shading to highlight items you added (green), modified (yellow), or deleted (red).



- Configure the number of versions Panorama stores for configuration audits.
  1. Select **Panorama** > **Setup** > **Management** and edit the Logging and Reporting Settings.
  2. Enter the **Number of Versions for Config Audit** (range is 1–1,048,576; default is 100).
  3. Click **OK** to save your changes.
  4. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

- View and compare Panorama configuration files before committing.
  1. Select **Commit** > **Commit to Panorama** and **Preview Changes**.
  2. Select the number of **Lines of Context** you want to see, and click **OK**.

 ©2025 Palo Alto Networks, Inc.

Manage Locks for Restricting Configuration Changes

Locking the candidate or running configuration prevents other administrators from changing the configuration until you manually remove the lock or Panorama removes it automatically (after a commit). Locks ensure that administrators don't make conflicting changes to the same settings or interdependent settings during concurrent login sessions.

If you are changing settings that are unrelated to the settings other administrators are changing in concurrent sessions, you don't need configuration locks to prevent commit conflicts. Panorama queues commit operations and performs them in the order that administrators initiate the commits. For details, see Panorama Commit, Validation, and Preview Operations.

A template or device group configuration push will fail if a firewall assigned to the template or device group has a commit or config lock that an administrator set locally on that firewall.

- View details about current locks.

  For example, you can check whether other administrators have set locks and read comments they entered to explain the locks.

  Click the locked padlock ( ) at the top of the web interface. The adjacent number indicates the number of current locks.

 ©2025 Palo Alto Networks, Inc.

- Lock a configuration.

  Read-only administrators who cannot modify firewall or Panorama configurations cannot set locks.

  1. Click the padlock icon at the top of the web interface.

     The icon varies based on whether existing locks are ( ) or are not ( ) set.
  2. **Take a Lock** and select the lock **Type**:
     - **Config**—Blocks other administrators from changing the candidate configuration.

       *A custom role administrator who cannot commit changes can set a **Config** lock and save the changes to the candidate configuration. However, because that administrator cannot commit the changes, Panorama does not automatically release the lock after a commit; the administrator must manually remove the **Config** lock after making the required changes.*

     - **Commit**—Blocks other administrators from changing the running configuration.
  3. Select the **Location** to determine the scope of the lock:
     - **Shared**—Restricts changes to the entire Panorama configuration, including all device groups and templates.
     - **Template**—Restricts changes to the firewalls included in the selected template. (You can't take a lock for a template stack, only for individual templates within the stack.)
     - **Device group**—Restricts changes to the selected device group but not its descendant device groups.
  4. (Optional) As a best practice, enter a **Comment** to describe your reason for setting the lock.
  5. Click **OK** and **Close**.

- Unlock a configuration.

  Only a superuser or the administrator who locked the configuration can manually unlock it. However, Panorama automatically removes a lock after completing the commit operation that the administrator who set the lock initiated.

  1. Click the locked padlock ( ) at the top of the web interface.
  2. Select the lock entry in the list.
  3. Click **Remove Lock**, **OK**, and **Close**.

- Configure Panorama to automatically lock the running configuration when you change the candidate configuration. This setting applies to all Panorama administrators.

  1. Select **Panorama** > **Setup** > **Management** and edit the General Settings.
  2. Select **Automatically Acquire Commit Lock** and click **OK**.
  3. Select **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

Add Custom Logos to Panorama

You can upload image files to customize the following areas on Panorama:

- Background image on the login screen
- Header on the top left corner of the web interface; you can also hide the Panorama default background
- Title page and footer image in PDF reports

Supported image types include .jpg and .png. Image files for use in PDF reports cannot contain an alpha channel. The size of the image must be less than 128 Kilobytes (131,072 bytes); the recommended dimensions are displayed on screen. If the dimension is larger than the recommended size, the image will be automatically cropped.

Only non-interlaced images are supported. The emailed Scheduled reports and Run Now custom reports do not contain the PDF attachment if a custom interlaced images are included in the PDF report title or PDF report header.

STEP 1 | Select **Panorama** > **Setup** > **Operations**.

STEP 2 | In the Miscellaneous section, click **Custom Logos**.

STEP 3 | Click the Upload logo icon and select an image for any of the following options: the login screen, the left corner of the main user interface, the PDF report title page and the PDF report footer.

STEP 4 | Click **Open** to add the image. To preview the image, click the preview logo icon.

STEP 5 | (Optional) To clear the green background header on the Panorama web interface, select the check box for **Remove Panorama background header**.

STEP 6 | Click **Close** to save your changes.

STEP 7 | Select **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

Use the Panorama Task Manager

Click **Tasks** (Tasks) at the bottom of the web interface to open the Task Manager, which displays details about all the operations that administrators initiated (for example, manual commits) or that Panorama or a managed firewall initiated (for example, scheduled report generation) since the last Panorama or firewall reboot. You can use the Task Manager to troubleshoot failed operations, investigate warnings associated with completed commits, or cancel pending commits.

Device Group and Template admins can only view tasks for tasks within their access domains.

STEP 1 | Click **Tasks**.

STEP 2 | **Show** the **Running** (in progress) tasks or **All** tasks (the default), optionally filter by type (**Reports**; **Log Requests**; or commit, download, and installation **Jobs**), and select **Panorama** (default) or the firewall for which you want to see the tasks.

STEP 3 | Perform any of the following actions:

- **Display or hide task details**—By default, the Task Manager displays the Type, Status, Start Time, and Messages for each task. To see the End Time and Job ID for a task, you must manually display those columns. To display or hide a column, open the drop-down in any column header, select **Columns**, and select or clear the columns as desired.
- **Investigate warnings or failures**—Read the entries in the Messages column for task details. If the column says `Toomany messages`, click the entry in the Type column to see more information.
- **Display a commit description**—If an administrator entered a description for a commit, click **Commit Description** in the Messages column to display it.
- **Check the position of a commit in the queue**—The Messages column indicates the queue position of commits that are in progress.
- **Cancel pending commits**—**Clear Commit Queue** to cancel all pending commits (available only to predefined administrative roles). To cancel an individual commit, click **x** in the Action column (the commit remains in the queue until Panorama dequeues it). You cannot cancel commits that are in progress.

 ©2025 Palo Alto Networks, Inc.

Manage Storage Quotas and Expiration Periods for Logs and Reports

- Log and Report Storage
- Log and Report Expiration Periods
- Configure Storage Quotas and Expiration Periods for Logs and Reports
- Configure the Run Time for Panorama Reports

Log and Report Storage

You can edit the default storage quotas for each log type. When a log quota reaches the maximum size, Panorama starts overwriting the oldest log entries with the new log entries.The storage capacity for reports is not configurable. The Log storage locations and report storage capacities vary by Panorama model:

- **Panorama virtual appliance in Panorama mode**—The storage space for reports is 200MB. The appliance uses its virtual system disk to store the System and Config logs that Panorama and Log Collectors generate. The virtual system disk also stores the Application Statistics (App Stats) logs that Panorama automatically receives at 15-minute intervals from all managed firewalls. Panorama stores all other log types to its virtual logging disks (1 to 12).
- **Panorama virtual appliance in Management Only mode**—The storage space for reports is 500MB. The appliance uses its virtual system disk to store the System and Config logs that Panorama and Log Collectors generate. The virtual system disk also stores the Application Statistics (App Stats) logs that Panorama automatically receives at 15-minute intervals from all managed firewalls. You must Configure a Managed Collector to forward logs from managed firewalls as Panorama in Management Only mode cannot store any other log type.
- **Panorama virtual appliance in Legacy mode**—The storage space for reports is 200MB for Panorama 8.0 or earlier releases and 500MB for Panorama 8.0.1 and later releases. Panorama writes all logs to its assigned storage space, which can be any of one the following:
  - **Virtual system disk**—By default, approximately 11GB is allocated for log storage on the virtual system disk that you created when installing Panorama. If you add a virtual logging disk or NFS partition, Panorama still uses the system disk to store the System and Config logs that Panorama and Log Collectors generate and to store the App Stats logs collected from firewalls.
  - **Dedicated virtual logging disk**—Stores all log types except those that reside on the system disk.
  - **NFS partition**—This option is available only to Panorama running on a VMware ESXi server. The NFS partition stores all log types except those that reside on the system disk.
- **M-700, M-600, M-500, M-300, or M-200 appliance**—The storage space for reports is 500MB for Panorama 6.1 or later releases and 200MB for earlier releases. The M-Series appliances use their internal SSD to store the Config logs and System logs that Panorama and Log Collectors generate and to store the App Stats logs collected from firewalls. Panorama saves all other log types to its RAID-enabled disks. The RAID disks are either local to the M-Series appliance in Panorama mode or are in a Dedicated Log Collector (M-Series appliance in Log Collector

 ©2025 Palo Alto Networks, Inc.

mode). You edit the log storage quotas on the RAID disks when you Configure a Collector Group.

For details on the log storage options and capacities, see Panorama Models. You can Expand Log Storage Capacity on the Panorama Virtual Appliance by adding virtual logging disks or NFS storage. You can Increase Storage on the M-Series Appliance by adding RAID drives or by upgrading from 1TB drives to 2TB drives.

Log and Report Expiration Periods

You can configure automatic deletion based on time for the logs that the Panorama management server and Log Collectors collect from firewalls, as well as the logs and reports that Panorama and the Log Collectors generate locally. This is useful in deployments where periodically deleting monitored information is desired or necessary. For example, deleting user information after a certain period might be mandatory in your organization for legal reasons. You configure separate expiration periods for:

- **Reports**—Panorama deletes expired reports at the same it generates new reports (see Configure the Run Time for Panorama Reports).
- **Each log type**—Panorama evaluates logs as it receives them, and deletes logs that exceed the configured expiration period.
- *Panorama synchronizes 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.*

  Even if you don't set expiration periods, when a log quota reaches the maximum size, Panorama starts overwriting the oldest log entries with the new log entries.

Configure Storage Quotas and Expiration Periods for Logs and Reports

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Configure the storage quotas and expiration periods for:

- Logs of all types that a Panorama virtual appliance in Legacy mode receives from firewalls.
- App Stats logs that Panorama receives from firewalls.
- System and Config logs that Panorama and Log Collectors generate locally.

The Panorama management server stores these logs locally.

If you reduce a storage quota such that the current logs exceed it, after you commit the change, Panorama removes the oldest logs to fit the quota.

1. Select **Panorama** > **Setup** > **Management** and edit the Logging and Reporting Settings.
2. In the **Log Storage** settings, enter the storage **Quota (%)** for each log type.

   When you change a percentage value, the page refreshes to display the corresponding absolute value (Quota GB/MB column) based on the total allotted storage on Panorama.
3. Enter the **Max Days** (expiration period) for each log type (range is 1 to 2,000).

   By default, the fields are blank, which means the logs never expire.

   Restore Defaults *if you want to reset the quotas and expiration periods to the factory defaults.*

STEP 2 | Configure the expiration period for reports that Panorama generates.

1. Select **Log Export and Reporting** and enter the **Report Expiration Period** in days (range is 1 to 2,000).

   By default, the field is blank, which means reports never expire.
2. Click **OK** to save your changes.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Configure the storage quotas and expiration periods for logs of all types (except App Stats logs) that M-700, M-600, M-500, M-300, M-200 appliances, or Panorama virtual appliance in Panorama mode receives from firewalls.

The local or Dedicated Log Collectors store these logs.

You configure these storage quotas at the Collector Group level, not for individual Log Collectors.

1. Select **Panorama** > **Collector Groups** and edit the Collector Group.
2. In the **General** settings, click the **Log Storage** value.

   A value doesn't display unless you assigned Log Collectors to the Collector Group. If the field displays 0MB after you assign Log Collectors, verify that you enable the disk pairs when you Configure a Managed Collector *and that you committed the changes (**Panorama** > **Managed Collectors** > **Disks**).*
3. Enter the storage **Quota(%)** for each log type.

   When you change a percentage value, the page refreshes to display the corresponding absolute value (Quota GB/MB column) based on the total storage allotted to the Collector Group.
4. Enter the **Max Days** (expiration period) for each log type (range is 1 to 2,000).

   By default, the fields are blank, which means the logs never expire.

   ***Restore Defaults** if you want to reset the quotas and expiration periods to the factory defaults.*
5. Click **OK** to save your changes.

STEP 4 | Commit the changes to Panorama and push the changes to the Collector Group.

1. Select **Commit** > **Commit and Push** and **Edit Selections** in the Push Scope.
2. Select **Collector Groups**, select the Collector Group you modified, and click **OK**.
3. **Commit and Push** your changes.

STEP 5 | Verify that Panorama applied the storage quota changes.

1. Select **Panorama** > **Setup** > **Management** and, in the Logging and Reporting Settings, verify that the **Log Storage** values are correct for the logs that the Panorama management server stores.
2. Select **Panorama** > **Collector Groups**, select the Collector Group you modified, and verify that the **Log Storage** values in the **General** tab are correct for the logs that the Log Collectors store.

   *You can also verify the Collector Group storage quotas by logging in to a Log Collector CLI and entering the operational command **`show log-diskquota-pct`**.*

 ©2025 Palo Alto Networks, Inc.

Configure the Run Time for Panorama Reports

Panorama generates reports daily at the time you specify. Panorama deletes any expired reports after generating the new reports.

STEP 1 | Select **Panorama** > **Setup** > **Management** and edit the Logging and Reporting Settings.

STEP 2 | Select **Log Export and Reporting** and set the **Report Runtime** to an hour in the 24-hour clock schedule (default is 02:00; range is 00:00 [midnight] to 23:00).

STEP 3 | Select **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

Monitor Panorama

To monitor Panorama and its managed collectors, you can periodically view their System and Config logs (filter logs by type), configure an SNMP manager to collect (GET) Panorama statistics on a regular basis, or configure SNMP traps or email alerts that notify you when a monitored metric changes state or reaches a threshold on Panorama. Email alerts and SNMP traps are useful for immediate notification about critical system events that need your attention. To configure email alerts or SNMP traps, see Configure Log Forwarding from Panorama to External Destinations.

- Panorama System and Configuration Logs
- Monitor Panorama and Log Collector Statistics Using SNMP

Panorama System and Configuration Logs

You can configure Panorama to send notifications when a system event or configuration change occurs. By default, Panorama records every configuration change in the Config logs. In the System logs, each event has a severity level to indicate its urgency and impact. When you Configure Log Forwarding from Panorama to External Destinations, you can forward all System and Config logs or filter the logs based on attributes such as the receive time or severity level (System logs only). The following table summarizes the severity levels for System logs.

Panorama regularly connects to the IoT Edge Service to download policy recommendations for IoT based policies. This connection is attempted by Panorama regardless of whether the IoT license is active on any managed firewalls..

A high severity gRPC connection failure system log is generated in the event of connection failure or if Panorama manages no IoT licensed firewall. No action is needed regarding these system logs if you are not leveraging the policy recommendation capabilities of IoT or if you are not managing any IoT licensed firewalls.

If you are leveraging the policy recommendation capabilities of IoT, review the gRPC connection failure system log to understand what is causing the connection issue between Panorama and the IoT Edge Service.

*Panorama does not support querying configuration logs in the **ACC** or when monitoring configuration logs (**Monitor** > **Logs**) using the filters:*

`before-change-preview-contains`

`after-change-preview-contains`

Severity	Description
Critical	Indicates a failure and the need for immediate attention, such as a hardware failure, including high availability (HA) failover and link failures.
High	Serious issues that will impair the operation of the system, including disconnection of a Log Collector or a commit failure.

 ©2025 Palo Alto Networks, Inc.

Severity	Description
Medium	Mid-level notifications, such as Antivirus package upgrades, or a Collector Group configuration push.
Low	Minor severity notifications, such as user password changes.
Informational	Notification events such as log in or log out, any configuration change, authentication success and failure notifications, commit success, and all other events that the other severity levels don't cover.

Panorama stores the System and Config logs locally; the exact location and storage capacity varies by Panorama model (see Log and Report Storage). Upon reaching the capacity limit, Panorama deletes the oldest logs to create space for new logs. If you need to store the logs for longer periods than what the local storage allows, you can Configure Log Forwarding from Panorama to External Destinations.

For information on using Panorama to monitor firewall logs, see Monitor Network Activity.

Monitor Panorama and Log Collector Statistics Using SNMP

You can configure an SNMP manager to request information from a Panorama management server and configure Panorama to respond. For example, the SNMP manager can request the high availability (HA) mode, Panorama state, and Panorama version. If the Panorama management server has a local Log Collector, then Panorama can also provide logging statistics: average logs per second, storage duration, retention periods, log disk usage, log forwarding status from individual firewalls to Panorama and external servers, and the status of firewall-to-Log Collector connections. Panorama doesn't synchronize SNMP configurations between HA peers; you must enable SNMP requests and responses on each peer.

You can also configure a Dedicated Log Collector to respond to requests for the same logging statistics as the Panorama management server. This information is useful when evaluating whether you need to expand log storage capacity.

You can't configure an SNMP manager to control Panorama or Log Collectors (using SET messages); an SNMP manager can only collect statistics (using GET messages).

For details on how Panorama implements SNMP, see SNMP Support.

 ©2025 Palo Alto Networks, Inc.

STEP 1 | Configure the SNMP Manager to get statistics from Panorama and the Log Collectors.

The following steps are an overview of the tasks you perform on the SNMP manager. For the specific steps, refer to the documentation of your SNMP manager.

1. To enable the SNMP manager to interpret statistics, load the Supported MIBs and, if necessary, compile them.
2. For each Panorama appliance that the SNMP manager will monitor, define its connection settings (IP address and port) and authentication settings (SNMPv2c community string or SNMPv3 username and password). All Panorama appliances use port 161.

   The SNMP manager can use the same or different connection and authentication settings for multiple Panorama management servers and Log Collectors. The settings must match those you define when you configure SNMP on Panorama (see Configure the Panorama management server to respond to statistics requests from an SNMP manager. and Configure the Panorama management server to respond to statistics requests from an SNMP manager.). For example, if you use SNMPv2c, the community string you define when configuring Panorama must match the community string you define in the SNMP manager for Panorama.
3. Determine the object identifiers (OIDs) of the statistics you will monitor. For example, to monitor the logging rate, a MIB browser shows that this statistic corresponds to OID 1.3.6.1.4.1.25461.2.3.30.1.1 in PAN-PRODUCT-MIB.my. For details, see Use an SNMP Manager to Explore MIBs and Objects.
4. Configure the SNMP manager to monitor the desired OIDs.

STEP 2 | Enable SNMP traffic on the management (MGT) interface of the Panorama management server.

1. Select **Panorama** > **Setup** > **Management** and edit the Management Interface Settings.
2. In the Services section, select the **SNMP** check box and click **OK**.

STEP 3 | Enable SNMP traffic on the management (MGT) interface of any M-Series appliances in Log Collector mode:

1. Select **Panorama** > **Managed Collectors** and select the Log Collector.
2. Select the **Management** tab, select the **SNMP** check box, and click **OK**.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Configure the Panorama management server to respond to statistics requests from an SNMP manager.

1. Select **Panorama** > **Setup** > **Operations** and, in the Miscellaneous section, click **SNMP Setup**.
2. Select the SNMP **Version** and configure the authentication values as follows. For version details, see SNMP Support.

   - **V2c**—Enter the **SNMP Community String**, which identifies a community of SNMP managers and monitored devices (Panorama, in this case), and serves as a password to authenticate the community members to each other.

     Don't use the default community string ***public****; it is well known and therefore not secure.*

   - **V3**—Create at least one SNMP view group and one user. User accounts and views provide authentication, privacy, and access control when SNMP managers get statistics.

     Views—Each view is a paired OID and bitwise mask: the OID specifies a MIB, and the mask (in hexadecimal format) specifies which objects are accessible inside (include matching) or outside (exclude matching) that MIB. Click **Add** in the first list and enter a **Name** for the group of views. For each view in the group, click **Add** and configure the view **Name**, **OID**, matching **Option** (**include** or **exclude**), and **Mask**.

     Users—Click **Add** in the second list, enter a username in the Users column, select the **View** group from the drop-down, enter the authentication password (**Auth Password**) used to authenticate to the SNMP manager, and enter the privacy password (**Priv Password**) used to encrypt SNMP messages to the SNMP manager.
3. Click **OK** to save the settings.

STEP 5 | Configure the Dedicated Log Collectors (if any) to respond to SNMP requests.

For each Collector Group:

1. Select **Panorama** > **Collector Groups** and select the Collector Group.
2. Select the **Monitoring** tab, configure the same settings as in Step Configure the Panorama management server to respond to statistics requests from an SNMP manager., and click **OK**.

STEP 6 | Commit the changes to Panorama and push the changes to Collector Groups.

1. Select **Commit** > **Commit and Push** and **Edit Selections** in the Push Scope.
2. Select **Collector Groups** you, select the Collector Groups you edited, and click **OK**.
3. **Commit and Push** your changes.

STEP 7 | Monitor the Panorama and Log Collector statistics in an SNMP manager.

Refer to the documentation of your SNMP manager.

 ©2025 Palo Alto Networks, Inc.

Reboot or Shut Down Panorama

The reboot option initiates a graceful restart of Panorama. A shutdown halts the system and powers it off. To restart Panorama, after a shutdown, manually disconnect and re-cable the power cord on the system.

STEP 1 | Select **Panorama** > **Setup** > **Operations**.

STEP 2 | In the Device Operations section, select **Reboot Panorama** or **Shutdown Panorama**.

 ©2025 Palo Alto Networks, Inc.

Configure Panorama Password Profiles and Complexity

To secure the local administrator account, you can define password complexity requirements that are enforced when administrators change or create new passwords. Unlike password profiles, which can be applied to individual accounts, the password complexity rules are firewall-wide and apply to all passwords.

To enforce periodic password updates, create a password profile that defines a validity period for passwords.

STEP 1 | Configure minimum password complexity settings.

1. Select **Panorama** > **Setup** > **Management** and edit the Minimum Password Complexity section.
2. Select **Enabled**.
3. Define the **Password Format Requirements**. You can enforce the requirements for uppercase, lowercase, numeric, and special characters that a password must contain.
4. To prevent the account username (or reversed version of the name) from being used in the password, select **Block Username Inclusion (including reversed)**.
5. Define the password **Functionality Requirements**.

   If you have configured a password profile for an administrator, the values defined in the password profile will override the values that you have defined in this section.

STEP 2 | Create password profiles.

You can create multiple password profiles and apply them to administrator accounts as required to enforce security.

1. Select **Panorama** > **Password Profiles** and click **Add**.
2. Enter a **Name** for the password profile and define the following:
   1. **Required Password Change Period**—Frequency, in days, at which the passwords must be changed.
   2. **Expiration Warning Period**—Number of days before expiration that the administrator will receive a password reminder.
   3. **Post Expiration Grace Period**—Number of days that the administrator can still log in to the system after the password expires.
   4. **Post Expiration Admin Login Count**—Number of times that the administrator can log in to the system after the password has expired.

 ©2025 Palo Alto Networks, Inc.

 ©2025 Palo Alto Networks, Inc.

Panorama Plugins

The Panorama extensible plugin architecture enables support for third-party integration plugins, such as VMware NSX, and other Palo Alto Networks products, such as the GlobalProtect cloud service. With this modular architecture, you can take advantage of new capabilities without waiting for a new PAN-OS version.

You can also configure the VM-Series plugin from Panorama. The VM-Series plugin is a single plugin that enables integration with public cloud environments such as Google Cloud Platform (GCP), Azure, AWS and private cloud hypervisors such as KVM, ESXi and others. The VM-Series plugin enables you to publish metrics from VM-Series firewalls deployed in public clouds. You can use Panorama to configure the VM-Series plugin settings for public clouds and push your configuration to your managed firewalls.

- About Panorama Plugins
- VM-Series Plugin and Panorama Plugins
- Endpoint Monitoring for Cisco TrustSec

About Panorama Plugins

Panorama supports an extensible plugin architecture that enables the integration and configuration of the following capabilities:

- **AIOps**—The Cloud Connector plugin (formerly, AIOps Plugin for Panorama) enables you to proactively enforce best practice checks by validating your commits and letting you know if a policy needs work before you push it to Panorama.
- **AWS**—The AWS plugin enables you to monitor your EC2 workloads on AWS. With the plugin, you can enable communication between Panorama (running PAN-OS 8.1.3 or a later release) and your AWS VPCs so that Panorama can collect a predefined set of attributes (or metadata elements) as tags for your EC2 instances and register the information to your Palo Alto Networks firewalls. When you reference these tags in Dynamic Address Groups and match against them in Security policy rules, you can consistently enforce policy across all assets deployed within your VPCs.
- **Azure**—The Azure plugin enables you to monitor your virtual machines on the Azure public cloud. With the plugin, you can enable communication between Panorama (running PAN-OS 8.1.6 or a later release) and your Azure subscriptions so that Panorama can collect a predefined set of attributes (or metadata elements) as tags for your Azure virtual machines and register the information to your Palo Alto Networks firewalls. When you reference these tags in Dynamic Address Groups and match against them in Security policy rules, you can consistently enforce policy across all assets deployed within VNets in your subscriptions.
- **Cisco ACI**—The Cisco ACI plugin enables you to monitor endpoints in your Cisco ACI fabric. With the plugin, you enable communication between Panorama (8.1.6 and later) and your Cisco APIC so that Panorama can collect endpoint information as tags for your Endpoint Groups and register the information to you Palo Alto Networks firewalls. When you reference these tags in Dynamic Address Groups and match against them in Security policy rules, you can consistently enforce policy across all assets deployed within your Cisco ACI fabric.
- **Cisco TrustSec**—The Cisco TrustSec Plugin enables monitoring of endpoints in your Cisco TrustSec environment. With the plugin, you enable communication between Panorama and your Cisco pxGrid server so that Panorama can collect endpoint information as tags for your endpoints and register the information to you Palo Alto Networks firewalls. When you reference these tags in Dynamic Address Groups and match against them in security policy rules, you can consistently enforce policy across all assets deployed within your Cisco TrustSec environment.

  Not supported on Panorama running PAN-OS 10.2.2 or earlier releases in FIPS-CC mode.
- **Cloud Services**—The Cloud Services plugin enables the use of the Strata Logging Service and Prisma Access. The Strata Logging Service solves operational logging challenges and the Prisma Access cloud service extends your security infrastructure to your remote network locations and mobile workforce.
- **Enterprise Data Loss Prevention (DLP)**— Enterprise DLP is a set of tools and processes that allow you to protect sensitive information against unauthorized access, misuse, extraction, or sharing. Enterprise DLP is enabled through a cloud service to help you inspect content and analyze the data in the correct context so that you can accurately identify sensitive data and

 ©2025 Palo Alto Networks, Inc.

secure it to prevent incidents. Enterprise DLP is supported on Panorama and managed firewalls running PAN-OS 10.0.2 and later releases.

Not supported on Panorama running PAN-OS 10.2.2 or earlier releases in FIPS-CC mode.

- **GCP**—Enables you to secure Kubernetes services in a Google Kubernetes Engine (GKE) cluster. Configure the Panorama plugin for Google Cloud Platform (GCP) to connect to your GKE cluster and learn about the services that are exposed to the internet.
- **Panorama Interconnect**—The Panorama Interconnect plugin enables you to manage large-scale firewall deployments. Use the Interconnect plugin to set up a two-tier Panorama deployment (on Panorama running PAN-OS 8.1.3 or a later release) for a horizontal scale-out architecture. With the Interconnect plugin, you can deploy a Panorama Controller with up to 64 Panorama Nodes or 32 Panorama HA pairs to centrally manage a large number of firewalls.
- **Nutanix**—The Panorama plugin for Nutanix enables VM monitoring in your Nutanix environment. It allows you to track the virtual machine inventory within your Nutanix Prism Central so that you can consistently enforce security policy that automatically adapts to changes within your Nutanix environment. As virtual machines are provisioned, de-provisioned or moved, this solution allows you to collect the IP addresses and associated sets of attributes (or metadata elements) as tags. You can then use the tags to define Dynamic Address Groups and use them in Security policy. The Panorama plugin for Nutanix requires Panorama 9.0.4 or later.
- **SD-WAN**—The Software-Defined Wide Area Network (SD-WAN) plugin allows you to use multiple internet and private services to create an intelligent and dynamic WAN, which helps lower costs and maximize application quality and usability. Instead of using costly and time-consuming MPLS with components such as routers, firewalls, WAN path controllers, and WAN optimizers to connect your WAN to the internet, SD-WAN on a Palo Alto Networks firewall allows you to use less expensive internet services and fewer pieces of equipment.
- **VMware NSX**—The VMware NSX plugin enables integration between the VM-Series firewall on VMware NSX with VMware NSX Manager. This integration allows you to deploy the VM-Series firewall as a service on a cluster of ESXi servers.
- **VMware vCenter**—The Panorama plugin for VMware vCenter allows you to monitor the virtual machines in your vCenter environment. The plugin retrieves IP addresses of virtual machines in your vCenter environment and converts them to tags that you can use to build policy using dynamic address groups.
- **Zero Touch Provisioning**—Zero Touch Provisioning (ZTP) is designed to simplify and automate the on-boarding of new firewalls to Panorama. ZTP streamlines the initial firewall deployment process by allowing network administrators to ship managed firewalls directly to their branches and automatically add the firewall to Panorama, allowing business to save on time and resources when deploying new firewalls. ZTP is supported on PAN-OS 9.1.3 and later releases.

  Not supported on Panorama in FIPS-CC mode.

- **IPS Signature Converter**—The IPS Signature Converter plugin for Panorama provides an automated solution for converting rules from third-party intrusion prevention systems—Snort and Suricata—into custom Palo Alto Networks threat signatures. You can then register these signatures on firewalls that belong to device groups you specify and use them to enforce policy in Vulnerability Protection and Anti-Spyware Security Profiles.

 ©2025 Palo Alto Networks, Inc.

You can install multiple plugins and retrieve IP address updates from multiple sources on a single Panorama instance. This allows you to create and enforce consistent security policy to secure applications and workloads across multiple cloud environments. Retrieved IP addresses are used in security policy through dynamic address groups; when a workload is added or removed from your environment, Panorama registers the change and pushes the update to the firewalls. When deploying multiple plugins on Panorama, you must carefully plan your device group hierarchy to ensure that updates are passed to your firewalls correctly.

Refer to the Palo Alto Networks Compatibility Matrix for details on the different plugin versions and compatibility information.

Install Panorama Plugins

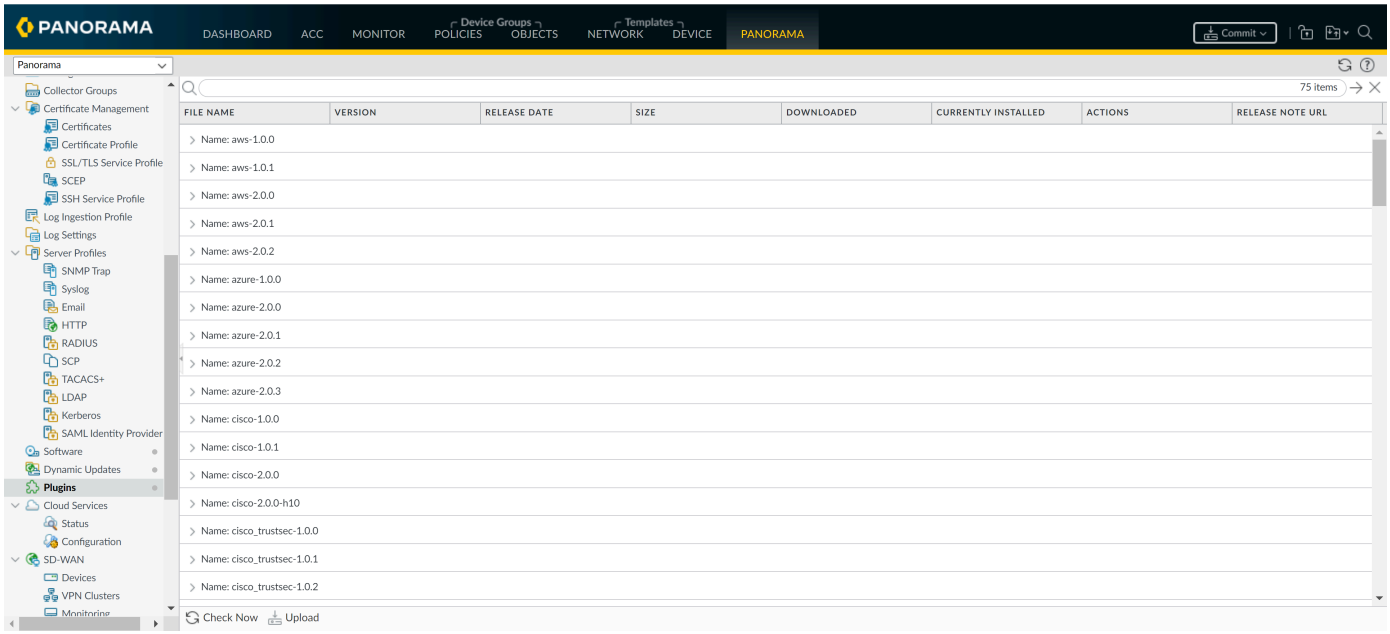
You can install one or more of the available plugins on Panorama to enable the integration of Cloud Services, VMware NSX, or for monitoring your virtual machines on AWS or Azure public cloud.

For the cloud services plugin, you must activate a valid auth code on the Customer Support Portal and select the region—Americas or Europe—to which you want to send logs.

*If you have a version of a plugin currently installed and you **Install** a new version of the plugin, Panorama replaces the currently installed version.*

STEP 1 | Download the plugin.

1. Select **Panorama** > **Plugins**.



2. Select **Check Now** to retrieve a list of available updates.
3. Select **Download** in the Action column to download the plugin.

   Refer to the Compatibility Matrix for the minimum supported PAN-OS version for each Panorama plugin.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Install the plugin.

Select the version of the plugin and click **Install** in the Action column to install the plugin. Panorama will alert you when the installation is complete. For more details on installing plugins, refer to VMware NSX plugin or the Cloud Services plugin documentation.

When installing the plugin for the first time on a Panorama HA pair, install the plugin on the passive peer before the active peer. On installing the plugin on the passive peer, it transitions to a non-functional state. Then, after you successfully install the plugin on the active peer, the passive peer returns to a functional state.

 ©2025 Palo Alto Networks, Inc.

VM-Series Plugin and Panorama Plugins

What is the difference between the VM-Series Plugin and various plugins for Panorama?

The VM-Series Plugin is for the VM-Series firewalls, and is a single plugin that enables integration with public cloud environments such as Google Cloud Platform (GCP), Azure and AWS, and private cloud hypervisors such as KVM, ESXi and others. When you deploy the firewall, the built-in plugin automatically detects the virtual environment on which the firewall is deployed and loads up the plugin components that enable you to manage interactions with that cloud environment. For example, when you deploy the VM-Series firewall on GCP, the VM-Series firewall loads the plugin components that enable the integration with GCP. You can then use the VM-Series plugin to configure the VM-Series firewall on GCP to publish metrics to Google Stackdriver Monitoring. Similarly, the VM-Series plugin on the VM-Series firewall on Azure enables you to configure the firewall to publish metrics Azure Application Insights or set up the details that the firewalls need to function as an HA pair. The VM-Series Plugin is pre-installed on the VM-Series firewall, and you can upgrade or downgrade but cannot delete it. On Panorama the VM-Series plugin is available but it is not pre-installed. If you choose to use Panorama to manage the integrations on your firewalls, install the VM-Series plugin on Panorama to establish communication with the VM-Series plugin on your firewalls.

The Panorama plugins are for both hardware-based firewalls and the VM-Series firewalls. Since Panorama plugins are optional, you can add, remove, reinstall, or upgrade them on Panorama. The Panorama plugin is not built-in, and you must install the plugin to enable communication with the managing the environment you need. For example, you use the Cloud Services plugin on Panorama to enable the set up between the Panorama/firewalls and the Cortex Data Lake. The GCP plugin on Panorama enables communication between Panorama and your GCP deployment so that you can secure traffic entering or exiting a service deployed in a Google Kubernetes Engine (GKE) cluster.

Install the VM-Series Plugin on Panorama

To view and configure cloud integrations deployed on your VM-Series firewalls, the VM-Series plugin must be installed on both Panorama and the VM-Series firewall. The plugin is automatically installed on the firewall, but you must manually install the plugin on Panorama before you can push configurations to your device groups.



The VM-Series plugin supports all clouds, so an upgrade might not apply to your VM-Series firewalls. Before upgrading the plugin, consult the release notes. Update the plugin only when there are changes relevant to your cloud.

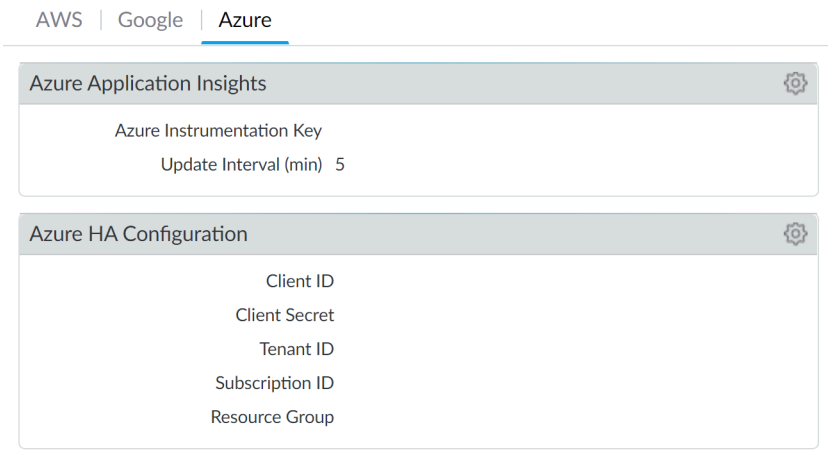
STEP 1 | Download the VM-Series plugin.

1. Select **Panorama** > **Plugins** and use **Check Now** to look for new plugin packages. The VM-Series plugin name is `vm_series`.
2. Consult the plugin release notes to determine which version provides upgrades useful to you.
3. Select a version of the plugin and select **Download** in the Action column.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Install the VM-Series plugin.

1. Click **Install** in the Action column. Panorama alerts you when the installation is complete.
2. To view the plugin, select **Device** > **VM-Series**.
   - If your firewall is installed on a private cloud and the hypervisor or service does not have an integration, you see a tab named VM-Series and the default message, `VM Series plugin infrastructure support is installed to allow the firewall's functionality to be enhanced in response to new features launched by hypervisor, or to meet new security needs.`
   - If your firewall is deployed on a public cloud, Panorama displays tabs for all supported clouds.



STEP 3 | (Optional) Save your configuration and push it to your managed firewalls.

STEP 4 | (Optional) On the VM-Series firewall, select **Device** > **VM-Series**. If you have configured the integration for your platform, you see a single tab for the cloud in which the firewall is deployed. If you have not configured an integration, you see the default message about the VM-Series plugin infrastructure.

 ©2025 Palo Alto Networks, Inc.

Endpoint Monitoring for Cisco TrustSec

Install and configure the Panorama plugin for Cisco TrustSec to retrieve the IP addresses of endpoints in your environment and build security policy for those endpoints using Dynamic Address Groups.

- Panorama Plugin for Cisco TrustSec
- Install the Panorama Plugin for Cisco TrustSec
- Configure the Panorama Plugin for Cisco TrustSec
- Troubleshoot the Panorama Plugin for Cisco TrustSec

Panorama Plugin for Cisco TrustSec

The Panorama plugin for Cisco TrustSec enables you to create security policy for your TrustSec environment using dynamic address groups. The plugin monitors for changes in TrustSec security groups and registers that information with Panorama and forwards IP information to the firewall, so Panorama can apply the correct policy to corresponding endpoints. The Panorama plugin for Cisco TrustSec supports up to 16 pxGrid (Cisco ISE) servers.

The Panorama plugin for Cisco TrustSec currently supports dynamic objects but not static objects.

The Panorama plugin processes the endpoint information and converts it to a set of tags that you can use as match criteria for placing IP addresses in dynamic address groups. Panorama creates a tag for each security group tag (SGT) on your pxGrid servers. The tags are constructed in the following format:

cts.svr_<pxgrid-server-name>.sgt_<SGT-name>

To retrieve endpoint IP-address-to-tag mapping information, you must configure a Monitoring Definition for each pxGrid server in your environment. The pxGrid server configuration specifies the username and password and is referenced by the monitoring definition that allows Panorama to connect to the pxGrid. Additionally, you can configure the plugin to verify the pxGrid server identity with a certificate profile on Panorama. It also specifies the device groups and corresponding notify groups containing the firewalls to which Panorama pushes the tags. After you configure the Monitoring Definition and the plugin retrieves the tags, you can create dynamic address groups and add the tags as match criteria.

The Panorama Plugin for Cisco TrustSec version 1.0.2 and later supports Bulk Sync and PubSub monitoring modes. The plugin selects a mode based on the Panorama version—Bulk Sync mode if the Panorama version is earlier than 10.0.0, and PubSub mode on Panorama 10.0.0 and later. The user interface displays the configuration options for the default monitoring mode.

- Bulk Sync
- PubSub

Bulk Sync

Bulk Sync mode uses two intervals to retrieve information from your pxGrid servers—the monitoring interval and full-sync interval. This mode is the default when the Panorama Plugin

 ©2025 Palo Alto Networks, Inc.

for Cisco TrustSec version 1.0.2 or later is installed on a Panorama version earlier than 10.0.0. Panorama versions earlier than 10.0.0 support IP-tab updates to configd every 10 seconds.

- **Monitoring interval**—The monitoring interval is the amount of time that the plugin waits before querying for changes. If no changes have occurred, the monitoring interval resets. If there are changes, the plugin processes the changes before resetting the monitoring interval. The default monitoring interval is 60 seconds. You can set the monitoring interval from 10 seconds to one day (86,400 seconds).

  The minimum monitoring interval is 30 seconds when the Panorama plugin for Cisco TrustSec 1.0.0 is installed.

- **Full-sync interval**—The full-sync interval is the amount of time that the plugin waits before updating the dynamic objects from all pxGrid servers regardless of any changes occurred. This ensures that the plugin is synchronized with the pxGrid server even if a change event is missed by the monitoring interval. You can set the full-sync interval from 600 seconds (10 minutes) to 86,400 seconds (one day). You must configure the full-sync interval from the Panorama CLI.

If the monitoring interval is greater than the full-sync interval, the full-sync interval is ignored and a full synchronization is performed at every monitoring interval.

PubSub

PubSub mode monitors notifications directly from the Cisco ISE server (the subscription daemon), parses for IP tags, and sends relevant information to the tag processing daemon (tag-proc). PubSub is the default mode when the Panorama Plugin for Cisco TrustSec version 1.0.2 or later is installed on Panorama version 10.0.0 or later. Panorama versions 10.0.0 or later support IP-tab updates to configd every 100 milliseconds.

- **Push interval**—The push interval is the amount of time between pushes. If the previous push takes too much time, the next push is triggered as soon as it finishes. The minimum push interval is 100 milliseconds (0 seconds) and the maximum is 60 seconds. The default push interval is 0 seconds.
- **Enable Full Sync**—Enable this option to trigger a complete update. If you enable full sync, you can set the full-sync interval. Default is no.
- **Full-sync interval**—The full-sync interval is the amount of time that the plugin waits before updating the dynamic objects from all pxGrid servers regardless of any changes occurred. The default full-sync interval is 10 minutes. You can set the full-sync interval from 600 seconds (10 minutes) to 86,400 seconds (one day). You must configure the full-sync interval from the Panorama CLI.
- **Reconnection interval**—The initial reconnection interval is 1 second, and it is doubled if the previous reconnection failed. The maximum reconnection interval is 64 sec. There is no limit to the number of reconnection attempts.

Install the Panorama Plugin for Cisco TrustSec

To get started with endpoint monitoring with Cisco TrustSec, download and install the Cisco TrustSec plugin on Panorama. To correlate the plugin version with the Panorama version, see Panorama Plugins in the Compatibility Matrix.

 ©2025 Palo Alto Networks, Inc.

Cisco TrustSec plugin upgrade or downgrade requires a commit.

If you have a Panorama HA configuration, repeat this installation process on each Panorama peer. When installing the plugin on Panorama appliances in an HA pair, install the plugin on the passive peer before the active peer. After installing the plugin on the passive peer, it will transition to a non-functional state. Installing the plugin on the active peer returns the passive peer to a functional state.

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away
- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  ```
  request plugins dau plugin-name <plugin-name> unblock-device-push yes
  ```

  You can cancel this command by executing:

  ```
  request plugins dau plugin-name <plugin-name> unblock-device-push no
  ```

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Select **Panorama** > **Plugins**.

STEP 2 | Click **Check Now** to get the latest version of the plugin.

STEP 3 | Select **Download** in the Action column to download the plugin.

STEP 4 | Select the version of the plugin and click **Install** in the Action column to install the plugin. Panorama will alert you when the installation is complete.

Configure the Panorama Plugin for Cisco TrustSec

After you install the plugin, you must also assign a notify group to your Cisco TrustSec plugin configuration. A notify group is a list of device groups that includes the firewalls to which Panorama should push all the tags it retrieves from the pxGrid server.

Each Panorama with the Cisco TrustSec plugin installed can support up to 16 pxGrid servers and 16 monitoring definitions. And each monitoring definition has one pxGrid server and one notify group.

 ©2025 Palo Alto Networks, Inc.

The Panorama plugin for Cisco TrustSec currently supports dynamic objects but not static objects.

The following configuration instructions cover both Bulk Sync and PubSub monitoring modes; some user interface features are enabled, or visible based on the monitoring mode.

STEP 1 | (Optional) Bypass proxy server settings, configured on Panorama under **Panorama** > **Setup** > **Services** > **Proxy Server**, for communication between Panorama and Cisco TrustSec. This command allows Panorama to communicate directly with Cisco TrustSec by bypassing the Panorama proxy configuration while maintaining proxied communication for other services.

1. Log in to the Panorama CLI.
2. Execute the following command to enable or disable proxy bypass.

   ```
   admin@Panorama> set plugins cisco_trustsec bypass-proxy {yes|no}
   ```

   Select **yes** to enable proxy bypass and **no** to disable proxy bypass. This is set to **no** by default.

STEP 2 | Configure the full-sync interval if you want to change it from the default 600 seconds (10 minutes).

1. Log in to the Panorama CLI.
2. Enter configure mode.

   ```
   admin@Panorama> configure
   ```

3. Use the following command to set the full-sync interval. The range is 600 seconds to 86,400 seconds (one day).

   ```
   admin@Panorama# set plugins cisco_trustsec full-sync-interval
   <interval-in-seconds>
   ```

STEP 3 | Log in to the Panorama web interface.

STEP 4 | You must add the firewalls as managed devices on Panorama and create Device Groups so that you can configure Panorama to notify these groups with the VM information it retrieves. Device groups can include VM-Series firewalls or virtual systems on the hardware firewalls.

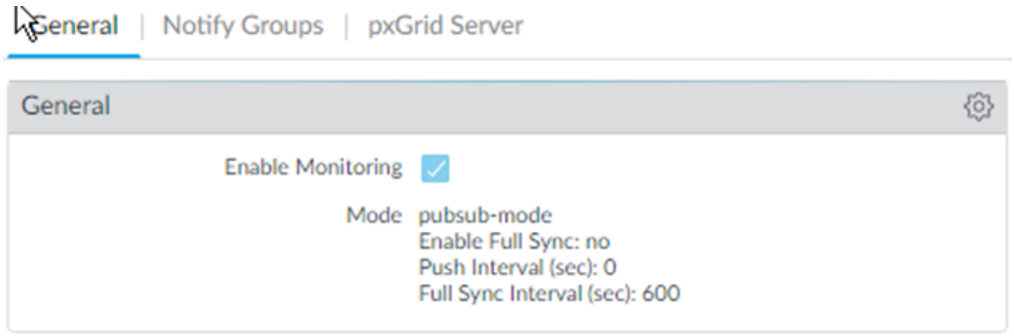
 ©2025 Palo Alto Networks, Inc.

STEP 5 | Configure Cisco TrustSec monitoring.

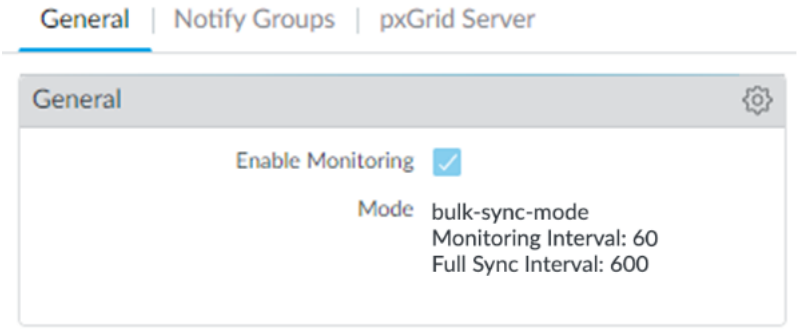
1. Select **Panorama** > **Cisco TrustSec** > **Setup** > **General**.

   Enable Cisco TrustSec Monitoring is enabled by default. This enables monitoring for all clusters in your deployment.

   The user interface selects the PubSub monitoring mode if the Panorama plugin for Cisco TrustSec is 1.0.2 or later is installed on Panorama 10.0.0 or later:



   The plugin selects Bulk Sync mode when it is installed on a Panorama version earlier than 10.0.0:

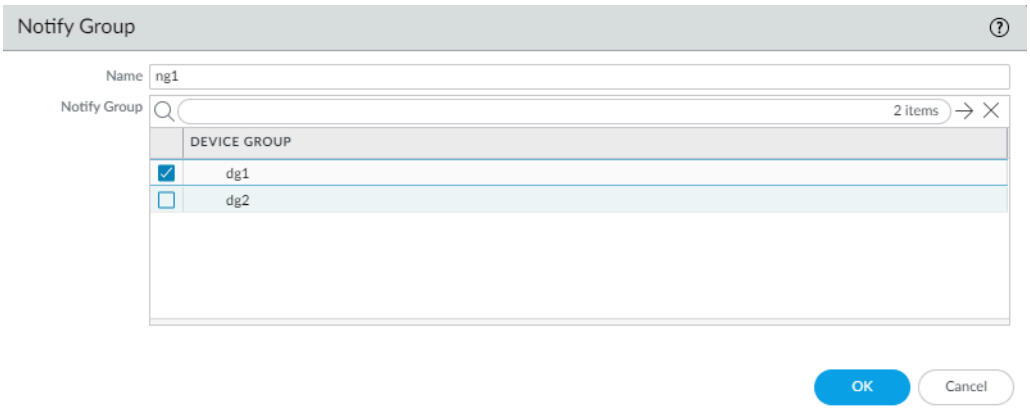


2. Click the gear to edit the setup parameters.
   - **Push Interval** (PubSub only)—Minimum 0, maximum 60 seconds, default is 0 (100 milliseconds).
   - **Enable Full Sync** (PubSub only, optional)—Select this optional to enable a full sync. Default is no.
   - **Full Sync Interval**.
     - PubSub—If Enable Full Sync is selected, you can set the full sync interval in seconds. Range is 600 seconds to 86400 seconds (one day), and the default value is 600
     - Bulk Sync—Enabled by default in Bulk Sync mode. Range is 600 seconds to 86400 seconds (one day), and the default value is 600.
   - **Monitoring Interval** (Bulk Sync only)—10 to 86400 seconds, default is 60—Set the polling interval at which Panorama queries the pxGrid server for endpoint address information. This is the time period between the end of a monitoring event and the start of the next event.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Create a notify group.

1. Select **Panorama** > **Cisco TrustSec** > **Setup** > **Notify Groups**.
2. Click **Add**.
3. Enter a descriptive **Name** for your notify group.
4. Select the device groups you created in previously.



STEP 7 | (Optional) If enabling server identity verification of the pxGrid server, configure a certificate profile on Panorama.

STEP 8 | Create, activate, and approve the pxGrid client name and client password.

1. Log in to the Panorama CLI.
2. Execute the following command to create the client name.
   - If you have a certificate profile, create the client name as follows:

     admin@Panorama> request plugins cisco_trustsec create-account client-name <client-name> host <ise-server-ip>
   - If you skipped Step 6 and you do not have a certificate, enter:

     request plugins cisco_trustsec create-account server-cert-verification-enabled no client-name <client-name>host <host-name>
3. Execute the following command to create the client name.

```
admin@Panorama> request plugins cisco_trustsec create-account
 client-name test host 10.10.10.15

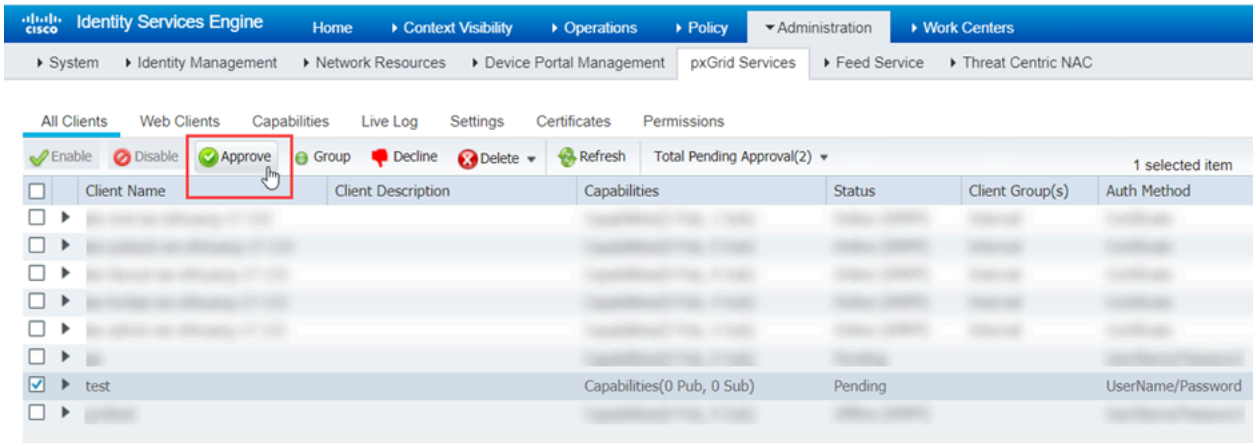
AccountCreate in progress...
AccountCreate successful.
  client nodename: test
  client password: <xxxxxxxx>

AccountActivate in progress...
AccountActivate successful.
```

 ©2025 Palo Alto Networks, Inc.

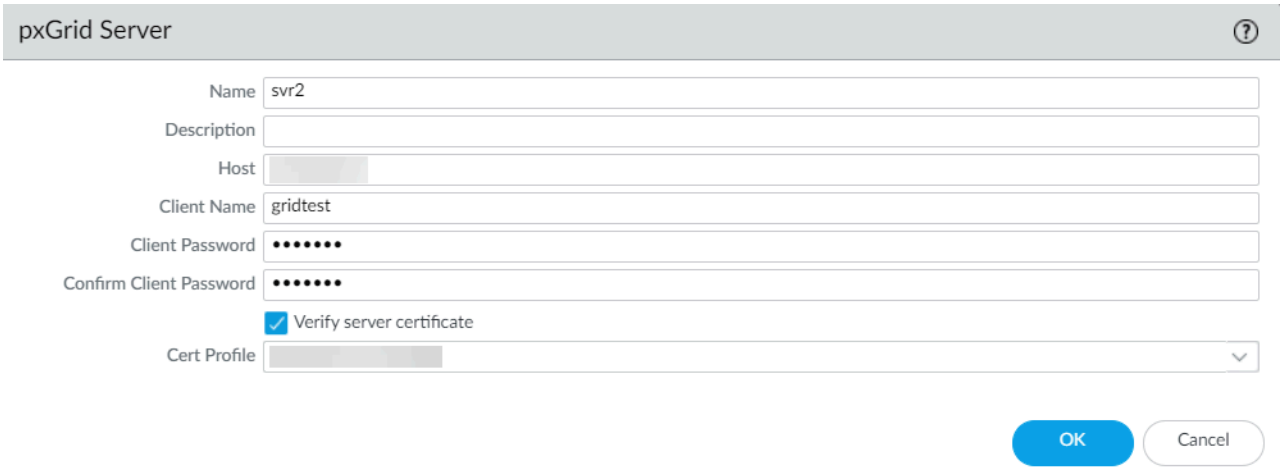
```
Please approve the account on the server.
```

4. Log in to your Cisco ISE server to approve the account.
5. Select **Administration** > **pxGrid Services** > **All Clients**.
6. Select the client name you create on Panorama.
7. Click **Approve**.



STEP 9 | Add pxGrid server information. The Panorama plugin for Cisco TrustSec supports up to 16 pxGrid (Cisco ISE) servers.

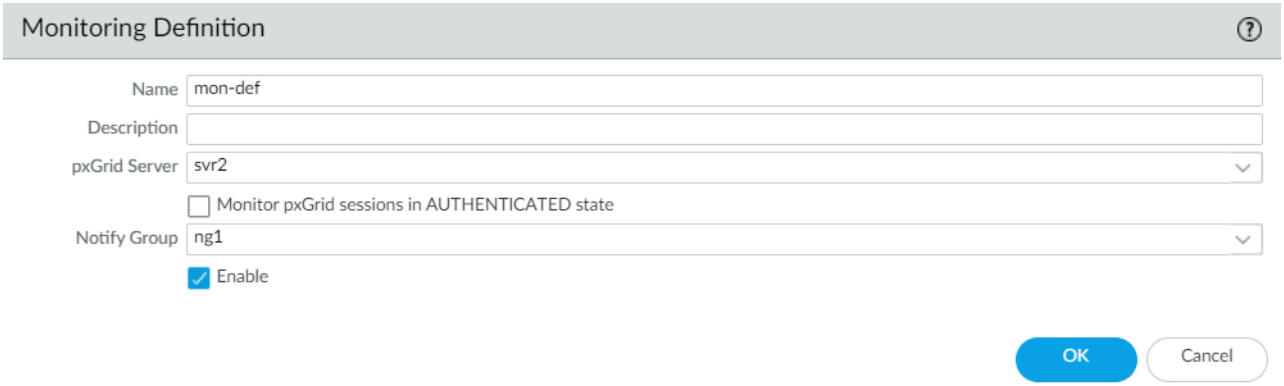
1. Select **Panorama** > **Cisco TrustSec** > **Setup** > **pxGrid Server**.
2. Enter a descriptive **Name** for your pxGrid server.
3. In the **Host** field, enter the IP address or FQDN for your pxGrid server.
4. Enter the client name you created in the previous step.
5. Enter and confirm the client password you generated in the previous step.
6. Verify the pxGrid server identity.
   1. Select **Verify server certificate**.
   2. Select your certificate profile from the **Cert Profile** drop-down.
7. Click **OK**.



 ©2025 Palo Alto Networks, Inc.

STEP 10 | Configure the Monitoring Definition.

1. Select **Panorama** > **Cisco TrustSec** > **Monitoring Definition** and click **Add**.
2. Enter a descriptive **Name** and optionally a **Description** to identify the monitoring definition.
3. Select the **pxGrid Server**.
4. (Optional) Set Panorama to **Monitor pxGrid sessions in AUTHENTICATED state**. By default, Panorama retrieves IP-Tag mappings from sessions in the “STARTED” state. ISE sessions have the “STARTED” state when there is a corresponding accounting start packet. If no accounting start packet is present for the session, then the session state is “AUTHENTICATED”.
5. Select the **Notify Group**.
6. Click **OK**.



STEP 11 | **Commit** your changes.

STEP 12 | Create active ISE sessions so that Panorama can learn SGT tags for dynamic address group definition. To create active sessions, use ISE to authenticate devices.

Panorama does not collect default SGT tags on ISE.

 ©2025 Palo Alto Networks, Inc.

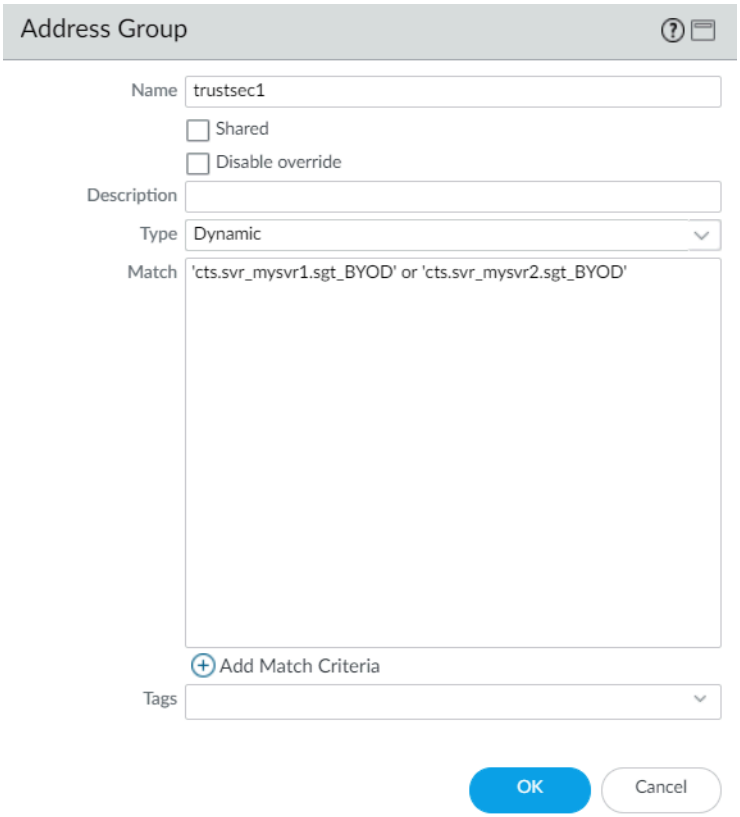
STEP 13 | Create dynamic address groups and verify that addresses are added to dynamic address groups.

1. Select **Objects** > **Address Groups**.
2. Select the Device Group you created for monitoring endpoints in your Cisco TrustSec environment from the **Device Group** drop-down.
3. Click **Add** and enter a **Name** and **Description** for the dynamic address group.

   The dynamic address group naming convention is: **`cts.svr_<server-name>.sgt_<SGT-name>`**
4. Select **Type** as **Dynamic**.
5. Click **Add Match Criteria**.
6. Select the **And** or **Or** operator and click the plus (+) icon next to the security group name to add it to the dynamic address group.

   Panorama can only display security group tags it has learned from active sessions. Security group tags in live sessions appear in the match criteria list.
7. Select **Panorama** > **Objects** > **Address Groups**.
8. Click **More** in the Addresses column of a dynamic address group.

   Panorama displays a list of IP addresses added to that dynamic address groups based on the match criteria you specified.



 ©2025 Palo Alto Networks, Inc.

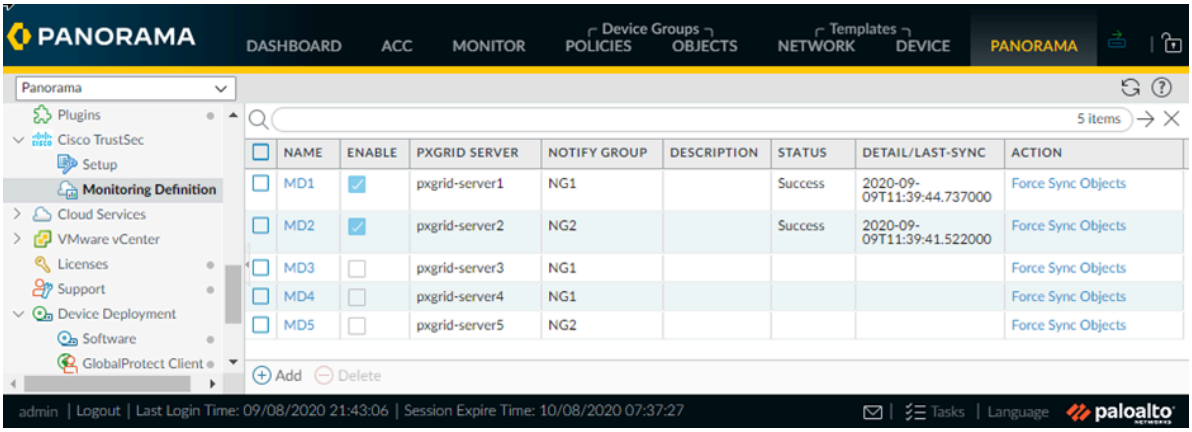
STEP 14 | Use dynamic address groups in policy.

Dynamic address groups are empty until you attach them to a policy. You won't see any IP addresses in your dynamic address group unless a policy is using it.

1. Select **Policies** > **Security**.
2. Click **Add** and enter a **Name** and a **Description** for the policy.
3. Add the **Source Zone** to specify the zone from which the traffic originates.
4. Add the **Destination Zone** at which the traffic is terminating.
5. For the **Destination Address**, select the dynamic address group you just created.
6. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles to the rule.
7. Repeats Steps 1 through 6 to create another policy rule.
8. Click **Commit**.

STEP 15 | (Optional) Update the dynamic objects from the pxGrid server at any time by synchronizing dynamic objects. Synchronizing dynamic objects enables you to maintain context on changes in the virtual environment and allows you to enable applications by automatically updating the dynamic address groups used in policy rules.

1. Select **Panorama** > **Cisco TrustSec** > **Monitoring Definition**.
2. Click **Synchronize Dynamic Objects**.



Troubleshoot the Panorama Plugin for Cisco TrustSec

- Plugin Status Commands
- Debug Commands
- Debug Logs

Plugin Status Commands

- Clear counters:

```
clear plugins cisco_trustsec counters
```

 ©2025 Palo Alto Networks, Inc.

- Display monitor status:

```
show plugins cisco_trustsec status
```

- Display counters:

```
show plugins cisco_trustsec counters
```

Debug Commands

- Check IP addresses in dynamic address groups.

```
show object registered-ip tag <tag>
```

```
show object registered-ip all
```

- Fetch the tags of an IP address from a server. The fetched ip-tag mappings are logged in `plugin_cisco_trustsec.log`. No ip-tag mappings are pushed to the notify group associated with the server. No retry if failed.

```
debug plugins cisco_trustsec query pxgrid-server $server-name ip
 $ip-address
```

- Force synchronize with a server and push the mappings to the configd process. No retry if failed.

```
request plugins cisco_trustsec synchronize-dynamic-objects name
 $server-name
```

- Force synchronize with all servers and push the mappings to the configd process. No retry if failed.

```
 request plugins cisco_trustsec synchronize-dynamic-objects all
```

- Force synchronize the mappings from configd process to VM-Series firewalls. No retry if failed.

```
request plugins cisco_trustsec sync
```

Debug Logs

The logs are in the following locations on the disk:

```
/opt/plugins/var/log/pan/plugin_cisco_trustsec.log
/opt/plugins/var/log/pan/plugin_cisco_trustsec_sub.log
/opt/plugins/var/log/pan/plugin_cisco_trustsec_ret.log
/opt/plugins/var/log/pan/plugin_cisco_trustsec_proc.log
```

 ©2025 Palo Alto Networks, Inc.

The size limit for a log file (shared by all plugins installed on your Panorama device) is 10 million bytes. A log file can accept 93,000 session logins. If you configure log rotation, a backup log can support 186,000 session logins.

- Change the plugin debug level.

```
request plugins debug level $level plugin-name cisco_trustsec
```

  - **off**: No debug log.
  - **low**: Dump only basic debug logs.
  - **medium**: Dump detailed debug logs.
  - **high**: Dump everything including request/response messages with servers.
- Merge the logs into a single log file:

```
request plugins cisco_trustsec merge-logs
```

- Show the debug log in the CLI:
  - Cisco TrustSec plugin version 1.0.2 or later installed on a Panorama version earlier than 10.0.0:

```
tail mp-log plugin_cisco_trustsec_merged.log
```

  - Cisco TrustSec plugin version 1.0.2 or later installed on Panorama version 10.0.0 or later:

```
tail follow yes plugins-log
```

 ©2025 Palo Alto Networks, Inc.

 ©2025 Palo Alto Networks, Inc.

Troubleshooting

The following topics address issues for the Panorama™ management server and Dedicated Log Collectors:

- Troubleshoot Panorama System Issues
- Troubleshoot Log Storage and Connection Issues
- Replace an RMA Firewall
- Troubleshoot Commit Failures
- Troubleshoot Registration or Serial Number Errors
- Troubleshoot Reporting Errors
- Troubleshoot Device Management License Errors
- Troubleshoot Automatically Reverted Firewall Configurations
- View Task Success or Failure Status
- Test Policy Match and Connectivity for Managed Devices
- Generate a Stats Dump File for a Managed Firewall
- Recover Managed Device Connectivity to Panorama
- Restore an Expired Device Certificate

Troubleshoot Panorama System Issues

- Generate Diagnostic Files for Panorama
- Diagnose Panorama Suspended State
- Monitor the File System Integrity Check
- Manage Panorama Storage for Software and Content Updates
- Recover from Split Brain in Panorama HA Deployments

Generate Diagnostic Files for Panorama

Diagnostic files aid in monitoring system activity and in discerning potential causes for issues on Panorama. To assist Palo Alto Networks Technical Support in troubleshooting an issue, the support representative might request a tech support file. The following procedure describes how to download a tech support file and upload it to your support case.

STEP 1 | Select **Panorama** > **Support** and click **Generate Tech Support File**.

STEP 2 | Download and save the file to your computer.

STEP 3 | Upload the file to your case on the Palo Alto Networks Customer Support web site.

Diagnose Panorama Suspended State

If Panorama is in a suspended state, check for the following conditions:

- **Serial numbers**—Verify that the serial number on each Panorama virtual appliance is unique. If the same serial number is used to create two or more instances of Panorama, all instances using the same serial number will be suspended.
- **Mode**—If you deploy the Panorama virtual appliance in a high availability (HA) configuration, verify that both HA peers are in the same mode: Panorama mode or Legacy mode.
- **HA priority**—Verify that you have set the HA priority setting on one peer as *Primary* and the other as *Secondary*. If the priority setting is identical on both peers, the Panorama peer with a higher numerical value in serial number is placed in a suspended state.
- **Panorama software version**—Verify that both Panorama HA peers are running the same Panorama software version (major and minor version number).

Monitor the File System Integrity Check

Panorama periodically performs a file system integrity check (FSCK) to prevent corruption of the Panorama system files. This check occurs after eight reboots or at a reboot that occurs 90 days after the last FSCK was executed. If Panorama is running a FSCK, the web interface and Secure Shell (SSH) login screens will display a warning to indicate that an FSCK is in progress. You cannot log in until this process completes. The time to complete this process varies by the size of the storage system; depending on the size, it can take several hours before you can log back in to Panorama.

After you successfully download and install a PAN-OS software update on Panorama or a managed firewall, the software update is validated after Panorama or the managed firewall

 ©2025 Palo Alto Networks, Inc.

reboots as part of the software installation process to ensure the PAN-OS software integrity. This ensures that the now running software update is known good and that the Panorama or managed firewall are not compromised to due remote or physical exploitation.

To view the progress on the FSCK, set up console access to Panorama and view the status.

Manage Panorama Storage for Software and Content Updates

You can Install Content and Software Updates for Panorama, upgrade firewalls, and upgrade upgrade Log Collectors using the Panorama™ management server. You cannot configure the amount of space available on Panorama to store updates. When the alloted storage capacity reaches 90%, Panorama alerts you to free up space (delete stored updates) for new downloads or uploads. The maximum number of updates is a global setting that applies to all the updates that Panorama stores. You must access the CLI to configure this setting. The default value is two updates of each type.

- Modify the maximum number of updates of each type.

  Access the Panorama CLI and enter the following, where <*number*> can be between 2 and 64:

  ```
  > set max-num-images count <number>
  ```

- View the number of updates that Panorama currently stores.

  Enter:

  ```
  > show max-num-images
  ```

- Use the web interface to delete updates to free up space on Panorama.
  1. Select the type of update to delete:
     - Firewall or Log Collector updates:

       PAN-OS/Panorama software images—Select **Panorama** > **Device Deployment** > **Software**.

       GlobalProtect™ agent/app software updates—Select **Panorama** > **Device Deployment** > **GlobalProtect Client**.

       Content updates—Select **Panorama** > **Device Deployment** > **Dynamic Updates**.
     - Panorama software images—Select **Panorama** > **Software**.
     - Panorama content updates—Select **Panorama** > **Dynamic Updates**.
  2. Click the **X** icon in the far right column for the image or update.

 ©2025 Palo Alto Networks, Inc.

- Use the CLI to delete updates to free up space on Panorama.

  Delete software images by version:

  ```
  > delete software version <version_number>
  ```

  Delete content updates:

  ```
  > delete content update <filename>
  ```

Recover from Split Brain in Panorama HA Deployments

When Panorama is configured in a high availability (HA) setup, the managed firewalls are connected to both the active and passive Panorama HA peers. When the connection between the active and the passive Panorama peers fails, before the passive Panorama takes over as the active peer it checks whether any firewall is connected to both the active and the passive peer. If even one firewall is connected to both peers, the failover is not triggered.

In the rare event that a failover is triggered when a set of firewalls are connected to the active peer and a set of firewalls are connected to the passive peer, but none of the firewalls are connected to both peers, it is called a split brain. When a split brain occurs, the following conditions occur:

- Neither Panorama peer is aware of the state nor the HA role of the other peer.
- Both Panorama peers become active and manage a unique set of firewalls.

To resolve a split brain, debug your network issues and restore connectivity between the Panorama HA peers.

However, if you need to make configuration changes to your firewalls without restoring the connection between the peers, here are a couple of options:

- Manually add the same configuration changes on both Panorama peers. This ensures that when the link is reestablished the configuration is synchronized.
- If you need to add/change the configuration at only one Panorama location, make the changes and synchronize the configuration (make sure that you initiate the synchronization from the peer on which you made the changes) when the link between the Panorama peers is re-established. To synchronize the peers, select the **Dashboard** tab and click the **Sync to peer** link in the High Availability widget.
- If you need to add/change the configuration for only the connected firewalls at each location, you can make configuration changes independently on each Panorama peer. Because the peers are disconnected, there is no replication and each peer now has a completely different configuration file (they are out of sync). Therefore, to ensure that the configuration changes on each peer are not lost when the connection is restored, you cannot allow the configuration to be automatically re-synchronized. To solve this problem, export the configuration from each Panorama peer and manually merge the changes using an external diff and merge tool. After the changes are integrated, you can import the unified configuration file on the primary Panorama and then synchronize the imported configuration file with the peer.

 ©2025 Palo Alto Networks, Inc.

Troubleshoot Log Storage and Connection Issues

Migrating logs is supported only for M-Series appliance. Refer to Migrate a Panorama Virtual Appliance to a Different Hypervisor to migrate a Panorama virtual appliance.

- Verify Panorama Port Usage
- Resolve Zero Log Storage for a Collector Group
- Replace a Failed Disk on an M-Series Appliance
- Replace the Virtual Disk on an ESXi Server
- Replace the Virtual Disk on vCloud Air
- Migrate Logs to a New M-Series Appliance in Log Collector Mode
- Migrate Logs to a New M-Series Appliance in Panorama Mode
- Migrate Logs to a New M-Series Appliance Model in Panorama Mode in High Availability
- Migrate Logs to the Same M-Series Appliance Model in Panorama Mode in High Availability
- Migrate Log Collectors after Failure/RMA of Non-HA Panorama
- Regenerate Metadata for M-Series Appliance RAID Pairs
- View Log Query Jobs

Verify Panorama Port Usage

To ensure that Panorama can communicate with managed firewalls, Log Collectors, and WildFire appliances and appliance clusters, and its high availability (HA) peer, use the following table to verify the ports that you must open on your network. Panorama uses TCP protocol for port communications.

By default, Panorama uses the management (MGT) interface to manage devices (firewalls, Log Collectors, and WildFire appliances and appliance clusters), collect logs, communicate with Collector Groups, and deploy software and content updates to devices. However, you can optionally assign the log collection and Collector Group communication functions to the Eth1 or Eth2 interfaces on an M-700, M-600, M-500, M-300, or M-200 appliance running Panorama 6.1 through 7.1. If the appliance runs Panorama 8.0 or a later release, you can assign any function to the Eth1, Eth2, Eth3, Eth4, or Eth5 interfaces on the M-700, M-600, M-500, M-300, or M-200 appliance. The ports listed in the following table apply regardless of which function you assign to which interface. For example, if you assign log collection to MGT and assign Collector Group communication to Eth2, then MGT will use port 3978 and Eth2 will use port 28270. (The Panorama virtual appliance can only use the MGT interface for all these functions.)

 ©2025 Palo Alto Networks, Inc.

Communicating Systems & Direction of Connection Establishment	Ports Used in Panorama 5.x	Ports Used in Panorama 6.x to 7.x	Ports Used in Panorama 8.x and later	Description
Panorama and Panorama (HA) Direction: Each peer initiates its own connection to the other	28	28	28	For HA connectivity and synchronization if encryption is enabled. Used for communication between Log Collectors in a Collector Group for log distribution.
Panorama and Panorama (HA) Direction: Each peer initiates its own connection to the other	28769 and 28260 (5.1) 28769 and 49160 (5.0)	28260 and 28769	28260 and 28769	For HA connectivity and synchronization if encryption is not enabled.
Panorama and managed firewalls Direction: Initiated by the firewall	3978	3978	3978	A bi-directional connection where the logs are forwarded from the firewall to Panorama; and configuration changes are pushed from Panorama to the managed firewalls. Context switching commands are sent over the same connection.
Panorama and Log Collector Direction: Initiated by the Log Collector	3978	3978	3978	For management and log collection/reporting. Used for communication between the local Log Collector on a Panorama in Panorama mode, and for communicating with Log Collectors in a distributed log collection deployment.
Panorama and managed devices (firewalls, Log Collectors, and WildFire	3978	3978	28443	Devices running PAN-OS 8.x or later releases use port 28443 to

 ©2025 Palo Alto Networks, Inc.

Communicating Systems & Direction of Connection Establishment	Ports Used in Panorama 5.x	Ports Used in Panorama 6.x to 7.x	Ports Used in Panorama 8.x and later	Description
appliances and appliance clusters) Direction: • Initiated by managed devices running PAN-OS 8.x or later releases. • Initiated by Panorama for devices running PAN-OS 7.x or earlier releases.				retrieve software and content update files from Panorama. Devices running 7.x or earlier releases do not retrieve update files from Panorama; Panorama pushes the update files to the devices over port 3978. Support for Panorama management of WildFire appliances and appliance clusters requires PAN-OS 8.0.1 or later installed on the managed WildFire appliances. We recommend that Panorama runs 8.0.1 or later to manage WildFire appliances and appliance clusters.
Log Collector to Log Collector Direction: Each Log Collector initiates a connection to the other Log Collectors in the Collector Group	49190	28270	28270	For distributing blocks and all binary data between Log Collectors.
Panorama to Strata Logging Service	NA	NA	444 *Version 8.0.5 and later.*	For setting up a secure communication channel with the Strata Logging Service. The managed firewalls use port 3978 to communicate with Strata Logging Service.

 ©2025 Palo Alto Networks, Inc.

Resolve Zero Log Storage for a Collector Group

The log storage capacity for the Collector Group might display as 0MB if the disk pairs are not enabled for logging in the Log Collectors. To enable the disk pairs, perform the following steps for each Log Collector in the Collector Group.

STEP 1 | Add the RAID disk pairs.

1. Select **Panorama** > **Managed Collectors** and click the Collector Name.
2. Select **Disks**, **Add** each RAID disk pair, and click **OK**.

STEP 2 | Commit the changes to Panorama and push the changes to the Collector Group.

1. Select **Commit** > **Commit and Push** and **Edit Selections** in the Push Scope.
2. Select **Collector Groups**, select the Collector Group you modified, and click **OK**.
3. **Commit and Push** your changes.

STEP 3 | Verify the state of the Log Collectors and disk pairs.

1. Select **Panorama** > **Managed Collectors** and verify that the configuration of each Log Collector is synchronized with Panorama.

   The Configuration Status column should display **In Sync** and the Run Time Status column should display **connected**.
2. Click **Statistics** in the last column for each Log Collector and verify that the disk pairs are **Enabled** and **Available**.

Replace a Failed Disk on an M-Series Appliance

If a disk fails on the M-Series appliance, you must replace the disk and reconfigure it in a RAID 1 array. For details, refer to the M-Series appliance Hardware Reference Guides.

Replace the Virtual Disk on an ESXi Server

You cannot resize a virtual disk after adding it to the Panorama virtual appliance running on a VMware ESXi server. Because the Panorama virtual appliance in Legacy mode allows only one log storage location, you must replace the virtual disk as follows to modify the log storage capacity. In Panorama mode, you can simply add another disk (up to the maximum of 12) to Expand Log Storage Capacity on the Panorama Virtual Appliance.

On the Panorama virtual appliance in Legacy mode, you will lose the logs on the existing disk when you replace it. For the options to preserve existing logs, see Preserve Existing Logs When Adding Storage on Panorama Virtual Appliance in Legacy Mode.

STEP 1 | Remove the old virtual disk.

1. Access the VMware vSphere Client and select the **Virtual Machines** tab.
2. Right-click the Panorama virtual appliance and select **Power** > **Power Off**.
3. Right-click the Panorama virtual appliance and select **Edit Settings**.
4. Select the virtual disk in the **Hardware** tab and click **Remove**.
5. Select one of the Removal Options and click **OK**.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Add the new virtual disk.

1. Add a Virtual Disk to Panorama on an ESXi Server.

   Panorama running on ESXi 5.5 and later versions supports a virtual disk of up to 8TB. Panorama running on an earlier ESXi version supports a virtual disk of up to 2TB.
2. In the vSphere Client, right-click the Panorama virtual appliance and select **Power** > **Power On**.

   The reboot process might take several minutes and the message `cache data unavailable` will display.

STEP 3 | Verify that the modified log storage capacity is correct.

1. Log in to the Panorama virtual appliance.
2. Select **Panorama** > **Setup** > **Management** and verify that the Logging and Reporting Settings section, Log Storage field, displays the modified log storage capacity accurately.

Replace the Virtual Disk on vCloud Air

You cannot resize a virtual disk after adding it to the Panorama virtual appliance running on VMware vCloud Air. Because the Panorama virtual appliance in Legacy mode allows only one log storage location, you must replace the virtual disk as follows to modify the log storage capacity. In Panorama mode, you can simply Add a Virtual Disk to Panorama on vCloud Air (up to the maximum of 12).

On the Panorama virtual appliance in Legacy mode, you will lose the logs on the existing disk when you replace it. For the options to preserve existing logs, see Preserve Existing Logs When Adding Storage on Panorama Virtual Appliance in Legacy Mode.

STEP 1 | Remove the old virtual disk.

1. Access the vCloud Air web console and select your **Virtual Private Cloud OnDemand** region.
2. Select the Panorama virtual appliance in the **Virtual Machines** tab.
3. Select **Actions** > **Edit Resources**.
4. Click **x** for the virtual disk you are removing.

STEP 2 | Add the new virtual disk.

1. **Add another disk**.
2. Set the **Storage** to up to 8TB and set the storage tier to **Standard** or **SSD-Accelerated**.
3. **Save** your changes.

STEP 3 | Reboot Panorama.

1. Log in to the Panorama virtual appliance.
2. Select **Panorama** > **Setup** > **Operations** and **Reboot Panorama**.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Verify that the modified log storage capacity is correct.

1. Log in to the Panorama virtual appliance after it reboots.
2. Select **Panorama** > **Setup** > **Management** and verify that the Logging and Reporting Settings section, Log Storage field, displays the modified log storage capacity accurately.

Migrate Logs to a New M-Series Appliance in Log Collector Mode

If you need to replace an M-Series appliance in Log Collector mode (Dedicated Log Collector), you can migrate the logs it collected from firewalls by moving its RAID disks to a new M-Series appliance. This procedure is supported for:

- Recovering logs after a system failure on an M-Series appliance and you are migrating to the same M-Series appliance model
- Migrating from an M-100 appliance to an M-500 appliance
- Migrating from an M-200 appliance to an M-600 appliance

Migrating logs by removing the logging disks from an M-100, M-300, or M-500 and loading them into an M-600 Panorama management server is not supported. To migrate from an M-100, M-300, or M-500 to an M-600 appliance, set up the M-600 appliance, configure log forwarding to the new M-600 appliance and configure the M-Series appliance as a managed Log Collector until you no longer needed access to the logs stored on the M-Series appliance.

STEP 1 | Perform initial setup of the new M-Series appliance that will be a Dedicated Log Collector.

1. Rack mount the M-Series appliance. Refer to the M-Series Appliance Hardware Reference Guides for instructions.
2. Perform Initial Configuration of the M-Series Appliance.

   When configuring interfaces, configure only the Management (MGT) interface. Switching to Log Collector mode (later in this procedure) removes the configurations for any other interfaces. If the Log Collector will use interfaces other than MGT, add them when configuring the Log Collector (see Step 2).
3. Register Panorama.
4. Purchase and activate the Panorama support license or transfer licenses as follows only if the new M-Series appliance is the same hardware model as the old M-Series appliance.

 ©2025 Palo Alto Networks, Inc.

If the new M-Series appliance is a different model than the old M-Series appliance, you must purchase new licenses.

1. Log in to the Palo Alto Networks Customer Support web site.
2. Select the **Assets** tab and click the **Spares** link.
3. Click the Serial Number of the new M-Series appliance.
4. Click **Transfer Licenses**.
5. **Select** the old M-Series appliance and click **Submit**.

5. Activate a firewall management license. If you are migrating from an M-200 appliance to a M-600 appliance, enter the auth-code associated with the migration license.
6. Install Content and Software Updates for Panorama. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
7. Switch from Panorama mode to Log Collector mode:
   1. Access the Log Collector CLI and switch to Log Collector mode:

      ```
      > request system system-mode logger
      ```

   2. Enter **Y** to confirm the mode change. The M-Series appliance reboots. If the reboot process terminates your terminal emulation software session, reconnect to the M-Series appliance to display the Panorama login prompt.

      *If you see a **CMS Login** prompt, press Enter without typing a username or password.*
8. Use the Log Collector CLI to enable connectivity between the Log Collector and Panorama management server. <IPaddress1 is for the MGT interface of the primary Panorama and <IPaddress2> is for the MGT interface of the secondary Panorama.

   ```
   > configure
   # set deviceconfig system panorama-server <IPaddress1>
    panorama-server-2 <IPaddress2>
   # commit
   # exit
   ```

STEP 2 | On the Panorama management server, add the new Log Collector as a managed collector.

For all steps with commands that require a serial number, you must type the entire serial number; pressing the Tab key won't complete a partial serial number.

1. Configure the Log Collector as a managed collector using the Panorama web interface or using the following CLI commands:

   ```
   > configure
   # set log-collector <LC_serial_number> deviceconfig system
    hostname <LC_hostname>
   ```

 ©2025 Palo Alto Networks, Inc.

```
# exit
```

*If the old Log Collector used interfaces other than the MGT interface for log collection and Collector Group communication, you must define those interfaces on the new Log Collector when you configure it as a managed collector (**Panorama** > **Managed Collectors** > **Interfaces**).*

2. Commit your changes to Panorama. Don't commit the changes to the Collector Group just yet.

```
> configure
# commit
# exit
```

3. Verify that the Log Collector is connected to Panorama and that the status of its disk pairs is present/available.

```
> show log-collector serial-number <log-collector_SN>
```

The disk pairs will display as disabled at this stage of the restoration process.

STEP 3 | Remove the RAID disks from the old Log Collector.

1. Power off the old Log Collector by pressing the Power button until the system shuts down.
2. Remove the disk pairs. For details, refer to the disk replacement procedure in the M-Series Appliance Hardware Reference Guides.

STEP 4 | Prepare the disks for migration.

Generating the metadata for each disk pair rebuilds the indexes. Therefore, depending on the data size, this process can take a long time to complete. To expedite the process, you can launch multiple CLI sessions and run the metadata regeneration command in each session to complete the process simultaneously for every pair. For details, see Regenerate Metadata for M-Series Appliance RAID Pairs.

1. Insert the disks into the new Log Collector. For details, refer to the disk replacement procedure in the M-Series Appliance Hardware Reference Guides.

The disk carriers of the M-200 appliance are incompatible with those of the M-600 appliance. Therefore, when migrating between these hardware models, you must unscrew each disk from its old carrier and insert the disk in the new carrier before inserting the disk in the new appliance.

You must maintain the disk pair association. Although you can place a disk pair from slot A1/A2 on the old appliance into slot B1/B2 on the new appliance, you must keep

 ©2025 Palo Alto Networks, Inc.

the disks together in the same slot; otherwise, Panorama might not restore the data successfully.

2. Enable the disk pairs by running the following CLI command for each pair:

```
> request system raid add <slot> force no-format
```

For example:

```
> request system raid add A1 force no-format
> request system raid add A2 force no-format
```

The **force** and **no-format** arguments are required. The **force** argument associates the disk pair with the new Log Collector. The **no-format** argument prevents reformatting of the drives and retains the logs stored on the disks.

3. Generate the metadata for each disk pair.

```
> request metadata-regenerate slot <slot_number>
```

For example:

```
> request metadata-regenerate slot 1
```

STEP 5 | Add a Log Collector with no disks to a Collector Group.

From this point, only commits that are required to complete the migration process on Panorama and the Log Collectors. Hold off making any other changes.

1. Access the Panorama CLI.
2. Overwrite Panorama restriction to allow Log Collector with no disk to be added to a Collector Group: **request log-migration-set-start**

STEP 6 | Migrate the logs.

You must use the Panorama CLI for this step, not the web interface.

You must assign the new Log Collector to the Collector Group that contains the old Log Collector.

1. Assign the new Log Collector to the Collector Group and commit your changes to Panorama.

```
> configure
# set log-collector-group <collector_group_name> logfwd-
setting collectors <new_LC_serial_number>
# commit
```

 ©2025 Palo Alto Networks, Inc.

```
# exit
```

2. For each disk pair, migrate the logs from the old Log Collector to the new Log Collector and attach the disk pair to the new Log Collector.

```
> request log-migration from <old_LC_serial_number> old-
disk-pair <log_disk_pair> to <new_LC_serial_number> new-disk-
pair <log_disk_pair>
```

For example:

```
> request log-migration from 003001000010 old-disk-pair A to
 00300100038 new-disk-pair A
```

STEP 7 | Reconfigure the Collector Group.

1. Use the web interface to assign the new Log Collector to the firewalls that forward logs (**Panorama** > **Collector Groups** > **Device Log Forwarding**). Give the new Log Collector the same priority in the firewall preference lists as the old Log Collector.

   You cannot use the CLI change the priority assignments of firewall preference lists.

2. Delete the old Log Collector from the Collector Group.

```
> configure
# delete log-collector-group <group_name> logfwd-setting
 collectors <old_LC_serial_number>
```

For example:

```
# delete log-collector-group DC-Collector-Group logfwd-setting
 collectors 003001000010
```

3. Delete the old Log Collector from the Panorama configuration and commit your changes to Panorama.

```
# delete log-collector <old_LC_serial_number>
# commit
```

 ©2025 Palo Alto Networks, Inc.

```
# exit
```

4. Commit the Collector Group changes so that the managed firewalls can send logs to the new Log Collector.

   ```
   > commit-all log-collector-config log-collector-
   group <collector_group_name>
   ```

   For example:

   ```
   > commit-all log-collector-config log-collector-group DC-
   Collector-Group
   ```

STEP 8 | Generate new keys on the new Dedicated Log Collector.

This command is required in order to add the new Log Collector to the Collector Group and should only be run for the Collector Group of the Log Collector being replaced. This step deletes the existing RSA keys and allows Panorama to create new RSA keys.

1. Access the Panorama CLI.
2. Delete all RSA keys on new Log Collector:

   request logdb update-collector-group-after-replace collector-group <collector-group-name>

   The process can take up to 10 minutes to completed.

 ©2025 Palo Alto Networks, Inc.

STEP 9 | Confirm that SearchEngine Status is Active for all Log Collectors in the Collector Group.

Do not continue until SearchEngine Status is Active for all Log Collectors in the Collector Group. This will result in purging of logs from the Log Collector being replaced.

1. Access the Panorama CLI.
2. Show the Log Collector details by running the following commands either:
   - On Panorama for all Log Collectors:

     show log-collector all

     Alternatively, you can run the following command on each Dedicated Log Collector:

     ```
     show log-collector detail
     ```
3. Confirm that SearchEngine Status is Active.

```
Redistribution status:       none
```

```
Last commit-all: commit succeeded, current ring version 1
```

```
SearchEngine status:     Active
```

```
md5sum 4e5055a359f7662fab8f8c4f57e24525 updated at 2017/06/14
 09:58:19
```

STEP 10 | On the new Log Collector, replace previous Log Collector serial number with the new Log Collector serial number.

You must replace the old Log Collector serial number with the new Log Collector serial number so that the new Log Collector will not run in to purging issues, resulting in the Log Collector being unable to purge old data from the migrated logs when necessary.

1. Access the Log Collector CLI.
2. Replace old Log Collector serial number with new Log Collector serial number:

   request log-migration-update-logger from *<old-log-collector-serial-number>* to *<new-log-collector-serial-number>*

Migrate Logs to a New M-Series Appliance in Panorama Mode

If you need to replace an M-Series appliance in Panorama mode (Panorama management server), you can migrate the logs it collected from firewalls by moving its RAID disks to a new M-Series appliance. This procedure is supported for:

- Recovering logs after a system failure on an M-Series appliance and you are migrating to the same M-Series appliance model

 ©2025 Palo Alto Networks, Inc.

- Migrating from an M-100 appliance to an M-500 appliance
- Migrating from an M-200 appliance to an M-600 appliance

> *Migrating logs by removing the logging disks from any M-Series appliance and loading them into an M-600 Panorama management server is not supported. To migrate to an M-600 appliance, set up the M-600 appliance, configure log forwarding to the new M-600 appliance and configure the M-Series appliance as a managed Log Collector until you no longer needed access to the logs stored on the M-Series appliance.*

This migration procedure covers the following scenarios where you are replacing a single M-Series appliance, not in a HA configuration, with a managed collector (Log Collector) in a Collector Group.

STEP 1 | Forward any logs on the SSD of the old M-Series appliance to an external destination if you want to preserve them.

The SSD stores the System and Config logs that Panorama and Log Collectors generate. You cannot move the SSD between M-Series appliances.

Configure Log Forwarding from Panorama to External Destinations.

STEP 2 | Export the Panorama configuration from the decommissioned M-Series appliance in Panorama mode.

1. Log in to the Panorama appliance and select **Panorama** > **Setup** > **Operations**.
2. Click **Save named Panorama configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**.
3. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**. Panorama exports the configuration to your client system as an XML file.

STEP 3 | Remove the RAID disks from the old M-Series appliance.

1. Power off the old M-Series appliance by pressing the Power button until the system shuts down.
2. Remove the disk pairs. For details, refer to the disk replacement procedure in the M-Series Appliance Hardware Reference Guides.

STEP 4 | Perform initial setup of the new M-Series appliance.

1. Rack mount the M-Series appliance. Refer to the M-Series Appliance Hardware Reference Guides for instructions.
2. Perform Initial Configuration of the M-Series Appliance.
3. Register Panorama.
4. Purchase and activate a Panorama support license or transfer licenses as follows only if the new M-Series appliance is the same hardware model as the old M-Series appliance.

 ©2025 Palo Alto Networks, Inc.

If the new M-Series appliance is a different model than the old M-Series appliance, you must purchase new licenses.

1. Log in to the Palo Alto Networks Customer Support web site.
2. Select the **Assets** tab and click the **Spares** link.
3. Click the Serial Number of the new M-Series appliance.
4. Click **Transfer Licenses**.
5. **Select** the old M-Series appliance and click **Submit**.

5. Activate a firewall management license. If you are migrating from an M-200 appliance to an M-600 appliance, enter the auth-code associated with the migration license.
6. Install Content and Software Updates for Panorama. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.

STEP 5 | Load the Panorama configuration snapshot that you exported from the decommissioned M-Series appliance into the new M-Series appliance in Panorama mode.

1. Log in to the Panorama Web Interface of the new M-Series appliance and select **Panorama** > **Setup** > **Operations**.
2. Click **Import named Panorama configuration snapshot**, **Browse** to the configuration file you exported from the decommissioned M-Series appliance, and click **OK**.
3. Click **Load named Panorama configuration snapshot**, select the **Name** of the configuration you just imported, select a **Decryption Key** (the master key for Panorama), and click **OK**. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file. If errors occurred, save them to a local file. Resolve each error to ensure the migrated configuration is valid.

   *To replace an RMA Panorama, make sure you **Retain Rule UUIDs** when you load the named Panorama configuration snapshot. If you do not select this option, Panorama removes all previous rule UUIDs from the configuration snapshot and assigns new UUIDs to the rules on Panorama, which means it does not retain information associated with the previous UUIDs, such as the policy rule hit count.*

4. Perform any additional configuration changes as needed.

   *If the old M-Series appliance used interfaces other than the MGT interface for Panorama services (such as log collection), you must define those interfaces on the new M-Series appliance (**Panorama** > **Setup** > **Interfaces**).*

5. Select **Commit** > **Commit to Panorama** and **Validate Commit**. Resolve any errors before proceeding.
6. **Commit** your changes to the Panorama configuration.

 ©2025 Palo Alto Networks, Inc.

STEP 6 | Insert the disks into the new M-Series appliance. For details, refer to the disk replacement procedure in the M-Series Appliance Hardware Reference Guides.

The disk carriers of the M-200 appliance are incompatible with those of the M-600 appliance. Therefore, when migrating between these hardware models, you must unscrew each disk from its old carrier and insert the disk in the new carrier before inserting the disk in the new appliance.

You must maintain the disk pair association. Although you can place a disk pair from slot A1/A2 on the old appliance into slot B1/B2 on the new appliance, you must keep the disks together in the same slot; otherwise, Panorama might not restore the data successfully.

STEP 7 | Contact Palo Alto Networks Customer Support to copy log collector group metadata from the decommissioned M-Series appliance to the new M-Series appliance and restart the mgmtsrvr process.

Refer to the Palo Alto Networks Knowledge Base when working with your Palo Alto Networks TAC engineer.

STEP 8 | If the M-Series appliance was part of a Collector Group, verify that the decommissioned M-Series appliance serial number is still part of the correct Collector Group:

```
debug log-collector-group show name <Log Collector Group name>
```

If the decommissioned M-Series appliance serial number is no longer a part of the correct Collector Group, then the Tech Support folders were incorrectly copied in the previous step. Contact Palo Alto Networks Customer Support again to copy the Tech Support folders to the correct location.

STEP 9 | Prepare the disks for migration.

Generating the metadata for each disk pair rebuilds the indexes. Therefore, depending on the data size, this process can take a long time to complete. To expedite the process, you can launch multiple CLI sessions and run the metadata regeneration command in each session to complete the process simultaneously for every pair. For details, see Regenerate Metadata for M-Series Appliance RAID Pairs.

1. Insert the disks into the new M-Series appliance. For details, refer to the disk replacement procedure in the M-Series Appliance Hardware Reference Guides.

   The disk carriers of the M-200 appliance are incompatible with those of the M-600 appliance. Therefore, when migrating between these hardware models, you must unscrew each disk from its old carrier and insert the disk in the new carrier before inserting the disk in the new appliance.

   You must maintain the disk pair association. Although you can place a disk pair from slot A1/A2 on the old appliance into slot B1/B2 on the new appliance, you must keep

 ©2025 Palo Alto Networks, Inc.

the disks together in the same slot; otherwise, Panorama might not restore the data successfully.

2. Enable the disk pairs by running the following CLI command for each pair:

```
admin> request system raid add <slot> force no-format
```

For example:

```
admin> request system raid add A1 force no-format
admin> request system raid add A2 force no-format
```

The **force** and **no-format** arguments are required. The **force** argument associates the disk pair with the new appliance. The **no-format** argument prevents reformatting of the drives and retains the logs stored on the disks.

3. Generate the metadata for each disk pair.

This step may take up to 6 hours depending on the volume of log data on the disks.

```
admin> request metadata-regenerate slot <slot_number>
```

For example:

```
admin> request metadata-regenerate slot 1
```

STEP 10 | Configure the local Log Collector on the new M-Series appliance.

For all steps with commands that require a serial number, you must type the entire serial number; pressing the Tab key won't complete a partial serial number.

Don't enable the disks on the new M-Series appliance at this point. When you successfully migrate the logs, Panorama automatically enables the disks.

1. Configure the local Log Collector as a managed collector using the Panorama web interface or using the following CLI commands:

```
admin> configure
admin# set log-collector <log-collector_SN> deviceconfig
 system hostname <log-collector-hostname>
```

 ©2025 Palo Alto Networks, Inc.

```
admin# exit
```

2. Verify that the local Log Collector is connected to Panorama and that the status of its disk pairs is present/available.

```
admin> show log-collector serial-number <log-collector_SN>
```

The disk pairs will display as disabled at this stage of the restoration process.

3. Commit your changes to Panorama. Don't commit the changes to the Collector Group just yet.

```
admin> configure
admin# commit
```

STEP 11 | Add a Log Collector with no disks to a Collector Group.

From this point, only commits that are required to complete the migration process on Panorama and the Log Collectors. Hold off making any other changes.

1. Access the Panorama CLI of the new M-Series appliance.
2. Overwrite Panorama restriction to allow Log Collector with no disk to be added to a Collector Group: **request log-migration-set-start**
3. Commit the overwritten restriction:

```
admin> configure
admin# commit force
```

STEP 12 | Migrate the logs.

1. Access the Panorama CLI of the new M-Series appliance.
2. Add the new local Log Collector as a member of the Collector Group and commit your changes to Panorama.

```
admin# set log-collector-group <collector_group_name> logfwd-
setting collectors <SN_managed_collector>
admin# commit
```

 ©2025 Palo Alto Networks, Inc.

```
admin# exit
```

The old local Log Collector still appears in the list of members, because you haven't deleted it from the configuration.

3. For each disk pair, migrate the logs to the new appliance.

```
admin> request log-migration from <old_LC_serial_number> old-disk-pair <log_disk_pair> to <new_LC_serial_number> new-disk-pair <log_disk_pair>
```

For example:

```
admin> request log-migration from 003001000010 old-disk-pair A
 to 00300100038 new-disk-pair A
```

4. Commit the changes to Panorama.

```
admin> configure
admin# commit
```

STEP 13 | Reconfigure the Collector Group.

1. Log in to the Panorama Web Interface of the new M-Series appliance to assign the new Log Collector to the firewalls that forward logs (**Panorama** > **Collector Groups** > **Device Log Forwarding**). Give the new Log Collector the same priority in the firewall preference lists as the old Log Collector.

   You cannot use the CLI change the priority assignments of firewall preference lists.

2. Access the Panorama CLI of the new M-Series appliance.
3. Delete the old Log Collector from the Collector Group.

```
admin# delete log-collector-group <group_name> logfwd-setting
 collectors <old_LC_serial_number>
```

For example:

```
admin# delete log-collector-group DC-Collector-Group logfwd-setting collectors 003001000010
```

4. Delete the old Log Collector from the Panorama configuration and commit your changes to Panorama.

```
admin# delete log-collector <old_LC_serial_number>
admin# commit
```

 ©2025 Palo Alto Networks, Inc.

```
admin# exit
```

5. Commit the Collector Group changes so that the managed firewalls can send logs to the new Log Collector.

   ```
   admin> commit-all log-collector-config log-collector-group <collector_group_name>
   ```

   For example:

   ```
   admin> commit-all log-collector-config log-collector-group DC-Collector-Group
   ```

STEP 14 | Generate new keys on the new Log Collector.

This command is required in order to add the new Log Collector to the Collector Group and should only be run for the Collector Group of the Log Collector being replaced. This step deletes the existing RSA keys and allows Panorama to create new RSA keys.

1. Access the Panorama CLI of the new M-Series appliance.
2. Delete all RSA keys on the new Log Collector:

   ```
   request logdb update-collector-group-after-replace collector-group <collector-group-name>
   ```

   The process can take up to 10 minutes to completed.

 ©2025 Palo Alto Networks, Inc.

STEP 15 | Confirm that SearchEngine Status is Active for all Log Collectors in the Collector Group.

Do not continue until SearchEngine Status is Active for all Log Collectors in the Collector Group. This will result in purging of logs from the Log Collector being replaced.

1. Access the Panorama CLI of the new M-Series appliance.
2. Show the Log Collector details by running the following commands either:
   - On Panorama for all Log Collectors:

     show log-collector all

     Alternatively, you can run the following command on each Dedicated Log Collector:

     ```
     show log-collector detail
     ```
3. Confirm that SearchEngine Status is Active.

```
Redistribution status:       none
```

```
Last commit-all: commit succeeded, current ring version 1
```

```
SearchEngine status:     Active
```

```
md5sum 4e5055a359f7662fab8f8c4f57e24525 updated at 2017/06/14
 09:58:19
```

STEP 16 | On the new Log Collector, replace previous Log Collector serial number with the new Log Collector serial number.

You must replace the old Log Collector serial number with the new Log Collector serial number so that the new Log Collector will not run in to purging issues, resulting in the Log Collector being unable to purge old data from the migrated logs when necessary.

1. Access the Log Collector CLI.
2. Replace old Log Collector serial number with new Log Collector serial number:

   request log-migration-update-logger from *<old-log-collector-serial-number>* to *<new-log-collector-serial-number>*

Migrate Logs to a New M-Series Appliance Model in Panorama Mode in High Availability

If you need to replace an M-700, M-600, M-500, M-300, M-200 or M-100 appliance in Panorama mode (Panorama management server) with a different M-Series appliance than the M-Series appliance being replaced, you can migrate the logs it collected from firewalls by moving its RAID disks to the new M-Series appliance. Moving the disks enables you to migrate logs as part of a

 ©2025 Palo Alto Networks, Inc.

hardware upgrade (from an M-100 appliance to an M-500 appliance). You can migrate an M-100 appliance to and from an M-500 appliance. M-100 and M-500 appliances cannot be migrated to or from M-200 or M-600 appliances.

Migrating logs by removing the logging disks from any M-Series appliance and loading them into an M-600 Panorama management server is not supported. To migrate to an M-600 appliance, set up the M-600 appliance, configure log forwarding to the new M-600 appliance and configure the M-Series appliance as a managed Log Collector until you no longer needed access to the logs stored on the M-Series appliance.

This migration procedure covers the following scenarios:

- One Panorama HA peer has a managed collector (Log Collector) in a Collector Group.

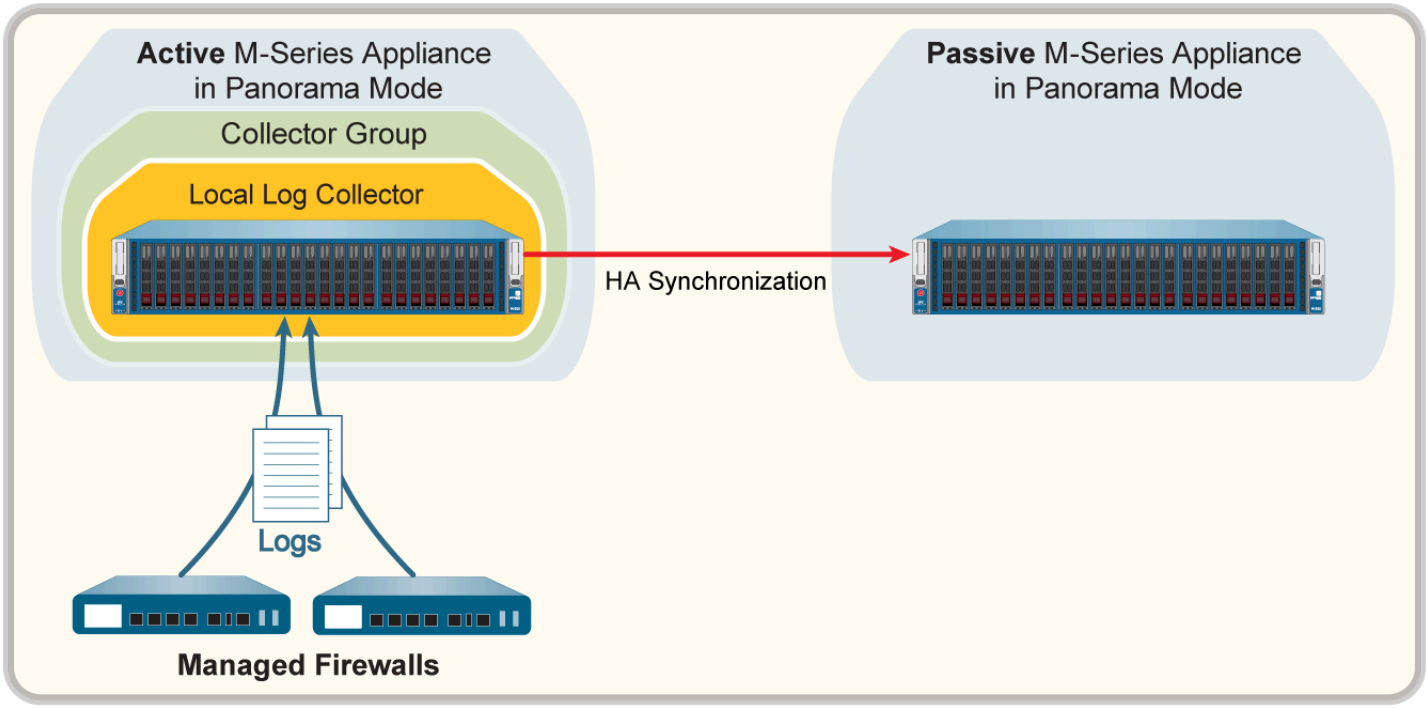


Figure 29: Panorama HA Peer with Collector Group

- Both Panorama HA peers have managed collectors that belong to a single Collector Group. For details, see Multiple Local Log Collectors Per Collector Group.
- Both Panorama HA peers have a managed collector and each is assigned to a separate Collector Group. For details, see Single Local Log Collector Per Collector Group.

STEP 1 | Forward any logs on the SSD of the old M-Series appliance to an external destination if you want to preserve them.

The SSD stores the System and Config logs that Panorama and Log Collectors generate. You cannot move the SSD between M-Series appliances.

Configure Log Forwarding from Panorama to External Destinations.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Export the Panorama configuration from the Primary decommissioned M-Series appliance in Panorama mode.

1. Log in to the Panorama Web Interface of the M-Series appliance you are replacing and select **Panorama** > **Setup** > **Operations**.
2. Click **Save named Panorama configuration snapshot**, enter a **Name** to identify the configuration, and click **OK**.
3. Click **Export named Panorama configuration snapshot**, select the **Name** of the configuration you just saved, and click **OK**. Panorama exports the configuration to your client system as an XML file.

STEP 3 | Remove the RAID disks from the old M-Series appliance.

1. Power off the old M-Series appliance by pressing the Power button until the system shuts down.
2. Remove the disk pairs. For details, refer to the disk replacement procedure in the M-Series Appliance Hardware Reference Guides.

STEP 4 | Perform initial setup of the new M-Series appliance.

Repeat this step for each of the new M-Series appliances in the HA configuration.

1. Rack mount the M-Series appliance. Refer to the M-Series Appliance Hardware Reference Guides for instructions.
2. Perform Initial Configuration of the M-Series Appliance.
3. Register Panorama.
4. Purchase and activate a Panorama support license or transfer licenses as follows only if the new M-Series appliance is the same hardware model as the old M-Series appliance. If the new M-Series appliance is a different model than the old M-Series appliance, you must purchase new licenses.
    1. Log in to the Palo Alto Networks Customer Support web site.
    2. Select the **Assets** tab and click the **Spares** link.
    3. Click the Serial Number of the new M-Series appliance.
    4. Click **Transfer Licenses**.
    5. **Select** the old M-Series appliance and click **Submit**.
5. Activate a firewall management license. If you are migrating from an M-100 appliance to an M-500 appliance, enter the auth-code associated with the migration license.
6. Install Content and Software Updates for Panorama. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
7. Set Up HA on Panorama. The new M-Series appliance must have the same priority as the HA peer you are replacing.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Load the Panorama configuration snapshot that you exported from the Primary decommissioned M-Series appliance into the new Primary M-Series appliance in Panorama mode.

1. Log in to the Panorama Web Interface of the new M-Series appliance and select **Panorama** > **Setup** > **Operations**.
2. Click **Import named Panorama configuration snapshot**, **Browse** to the configuration file you exported from the decommissioned M-Series appliance, and click **OK**.
3. Click **Load named Panorama configuration snapshot**, select the **Name** of the configuration you just imported, select a **Decryption Key** (the master key for Panorama), and click **OK**. Panorama overwrites its current candidate configuration with the loaded configuration. Panorama displays any errors that occur when loading the configuration file. If errors occurred, save them to a local file. Resolve each error to ensure the migrated configuration is valid.

   *To replace an RMA Panorama, make sure you **Retain Rule UUIDs** when you load the named Panorama configuration snapshot. If you do not select this option, Panorama removes all previous rule UUIDs from the configuration snapshot and assigns new UUIDs to the rules on Panorama, which means it does not retain information associated with the previous UUIDs, such as the policy rule hit count.*
4. Perform any additional configuration changes as needed.

   *If the old M-Series appliance used interfaces other than the MGT interface for Panorama services (such as log collection), you must define those interfaces on the new M-Series appliance (**Panorama** > **Setup** > **Interfaces**).*
5. Select **Commit** > **Commit to Panorama** and **Validate Commit**. Resolve any errors before proceeding.
6. **Commit** your changes to the Panorama configuration. Once committed, the Panorama configuration is synced across the HA peers.

STEP 6 | Insert the disks into the new M-Series appliance. For details, refer to the disk replacement procedure in the M-Series Appliance Hardware Reference Guides.

Repeat this step for each of the new M-Series appliances in the HA configuration.

The disk carriers of the M-100 appliance are incompatible with those of the M-500 appliance. Therefore, when migrating between these hardware models, you must unscrew each disk from its old carrier and insert the disk in the new carrier before inserting the disk in the new appliance.

You must maintain the disk pair association. Although you can place a disk pair from slot A1/A2 on the old appliance into slot B1/B2 on the new appliance, you must keep the disks together in the same slot; otherwise, Panorama might not restore the data successfully.

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Contact Palo Alto Networks Customer Support to copy log collector group metadata from the decommissioned M-Series appliance to the new M-Series appliance and restart the mgmtsrvr process.

Refer to the Palo Alto Networks Knowledge Base when working with your Palo Alto Networks TAC engineer.

STEP 8 | If the M-Series appliance was part of a Collector Group, verify that the decommissioned M-Series appliance serial number is still part of the correct Collector Group:

```
debug log-collector-group show name <Log CollectorGroup name>
```

If the decommissioned M-Series appliance serial number is no longer a part of the correct Collector Group, then the Tech Support folders were incorrectly copied in the previous step. Contact Palo Alto Networks Customer Support again to copy the Tech Support folders to the correct location.

STEP 9 | Prepare the disks for migration.



Generating the metadata for each disk pair rebuilds the indexes. Therefore, depending on the data size, this process can take a long time to complete. To expedite the process, you can launch multiple CLI sessions and run the metadata regeneration command in each session to complete the process simultaneously for every pair. For details, see Regenerate Metadata for M-Series Appliance RAID Pairs.

1. Enable the disk pairs by running the following CLI command for each pair:

```
admin> request system raid add <slot> force no-format
```

For example:

```
admin> request system raid add A1 force no-format
```

 ©2025 Palo Alto Networks, Inc.

```
admin> request system raid add A2 force no-format
```

The **force** and **no-format** arguments are required. The **force** argument associates the disk pair with the new appliance. The **no-format** argument prevents reformatting of the drives and retains the logs stored on the disks.

2. Generate the metadata for each disk pair.

   This step may take up to 60 hours to complete depending on the volume of log data on the disks. The new M-Series appliance must be in an `active` state. Panorama may experience issues getting the assigned Log Collector ID if Panorama is in a `suspended` state.

   If the new M-Series appliance is the `passive` HA peer, log in to the Panorama web interface of the currently `active` HA peer and select ***Panorama > High Availability*** *and* ***Suspend local Panorama for high availability****.*

   After you successfully migrate the logs, ***Make local Panorama functional for high availability****.*

```
admin> request metadata-regenerate slot <slot_number>
```

For example:

```
admin> request metadata-regenerate slot 1
```

STEP 10 | Configure the local Log Collector on the new M-Series appliance.

For all steps with commands that require a serial number, you must type the entire serial number; pressing the Tab key won't complete a partial serial number.

Don't enable the disks on the new M-Series appliance at this point. When you successfully migrate the logs, Panorama automatically enables the disks.

1. Configure the local Log Collector as a managed collector using the Panorama web interface or using the following CLI commands:

```
admin> configure
admin# set log-collector <log-collector_SN> deviceconfig
 system hostname <log-collector-hostname>
admin# exit
```

2. Commit your changes to Panorama. Don't commit the changes to the Collector Group just yet.

```
admin> configure
```

 ©2025 Palo Alto Networks, Inc.

```
admin# commit
```

3. Verify that the local Log Collector is connected to Panorama and that the status of its disk pairs is present/available.

```
admin> show log-collector serial-number <log-collector_SN>
```

The disk pairs will display as disabled at this stage of the restoration process.

STEP 11 | Add a Log Collector with no disks to a Collector Group.

From this point, only commits that are required to complete the migration process on Panorama and the Log Collectors. Hold off making any other changes.

1. Access the Panorama CLI of the new M-Series appliance.
2. Overwrite Panorama restriction to allow Log Collector with no disk to be added to a Collector Group: **request log-migration-set-start**
3. Commit the changes to Panorama.

```
admin> configure
admin# commit force
```

STEP 12 | Migrate the logs.

The new M-Series appliance must be the `active` HA peer before you can begin migrating logs. If the new M-Series appliance, log in to the Panorama web interface *of the `active` HA peer and select* ***Panorama > High Availability*** *and* ***Suspend local Panorama for high availability****.*

After you successfully migrate the logs, ***Make local Panorama functional for high availability****.*

1. Access the Panorama CLI of the new M-Series appliance.
2. Add the new local Log Collector as a member of the Collector Group and commit your changes to Panorama.

```
admin# set log-collector-group <collector_group_name> logfwd-
setting collectors <SN_managed_collector>
admin# commit
```

 ©2025 Palo Alto Networks, Inc.

```
admin# exit
```

The old local Log Collector still appears in the list of members, because you haven't deleted it from the configuration.

3. For each disk pair, migrate the logs to the new appliance.

```
admin> request log-migration from <old_LC_serial_number> old-disk-pair <log_disk_pair> to <new_LC_serial_number> new-disk-pair <log_disk_pair>
```

For example:

```
admin> request log-migration from 003001000010 old-disk-pair A to 00300100038 new-disk-pair A
```

4. Commit the changes to Panorama.

```
admin> configure
admin# commit
```

STEP 13 | Reconfigure the Collector Group.

1. Log in to the Panorama Web Interface of the new M-Series appliance to assign the new Log Collector to the firewalls that forward logs (**Panorama** > **Collector Groups** > **Device Log Forwarding**). Give the new Log Collector the same priority in the firewall preference lists as the old Log Collector.

   You cannot use the CLI change the priority assignments of firewall preference lists.

2. Access the Panorama CLI of the new M-Series appliance.
3. Delete the old Log Collector from the Collector Group.

```
admin# delete log-collector-group <group_name> logfwd-setting collectors <old_LC_serial_number>
```

For example:

```
admin# delete log-collector-group DC-Collector-Group logfwd-setting collectors 003001000010
```

4. Delete the old Log Collector from the Panorama configuration and commit your changes to Panorama.

```
admin# delete log-collector <old_LC_serial_number>
admin# commit
```

 ©2025 Palo Alto Networks, Inc.

```
admin# exit
```

5. Commit the Collector Group changes so that the managed firewalls can send logs to the new Log Collector.

   ```
   admin> commit-all log-collector-config log-collector-group <collector_group_name>
   ```

   For example:

   ```
   admin> commit-all log-collector-config log-collector-group DC-Collector-Group
   ```

STEP 14 | Generate new keys on the new Log Collector.

This command is required in order to add the new Log Collector to the Collector Group and should only be run for the Collector Group of the Log Collector being replaced. This step deletes the existing RSA keys and allows Panorama to create new RSA keys.

1. Access the Panorama CLI of the new M-Series appliance.
2. Delete all RSA keys on the new Log Collector:

   ```
   request logdb update-collector-group-after-replace collector-group <collector-group-name>
   ```

   The process can take up to 10 minutes to completed.

 ©2025 Palo Alto Networks, Inc.

STEP 15 | Confirm that SearchEngine Status is Active for all Log Collectors in the Collector Group.

Do not continue until SearchEngine Status is Active for all Log Collectors in the Collector Group. This will result in purging of logs from the Log Collector being replaced.

1. Access the Panorama CLI of the new M-Series appliance.
2. Show the Log Collector details by running the following commands either:
   - On Panorama for all Log Collectors:

     show log-collector all

     Alternatively, you can run the following command on each Dedicated Log Collector:

     ```
     show log-collector detail
     ```
3. Confirm that SearchEngine Status is Active.

```
Redistribution status:       none
```

```
Last commit-all: commit succeeded, current ring version 1
```

```
SearchEngine status:     Active
```

```
md5sum 4e5055a359f7662fab8f8c4f57e24525 updated at 2017/06/14
 09:58:19
```

STEP 16 | On the new Log Collector, replace previous Log Collector serial number with the new Log Collector serial number.

You must replace the old Log Collector serial number with the new Log Collector serial number so that the new Log Collector will not run in to purging issues, resulting in the Log Collector being unable to purge old data from the migrated logs when necessary.

1. Access the Log Collector CLI.
2. Replace old Log Collector serial number with new Log Collector serial number:

   request log-migration-update-logger from <*old-log-collector-serial-number*> to <*new-log-collector-serial-number*>

 ©2025 Palo Alto Networks, Inc.

STEP 17 | Set up the new secondary Panorama high availability peer.

1. Forward any logs on the SSD of the old M-Series appliance to an external destination if you want to preserve them.
2. Remove the RAID disks from the old M-Series appliance.
3. Perform initial setup of the new M-Series appliance.
4. Insert the disks into the new M-Series appliance.
5. Repeat Steps 7 through 16 to migrate the logs from the old M-Series appliance to the new M-Series appliance.
6. Set Up HA on Panorama. The new M-Series appliance must have the same priority as the HA peer you are replacing.
7. Log in to the Panorama Web Interface of the primary HA peer and click **Dashboard** > **High Availability** > **Sync to peer** to synchronize the configuration of the M-Series appliance HA peers.

Migrate Logs to the Same M-Series Appliance Model in Panorama Mode in High Availability

If you need to replace an M-700, M-600, M-500, M-300, M-200, or M-100 appliance deployed in high availability (HA) configuration in Panorama mode (Panorama management server) with the same M-Series appliance as the M-Series appliance being replaced, you can migrate the logs it collected from firewalls by moving its RAID disks to the new M-Series appliance. Moving the disks enables you to recover logs after a system failure on the M-Series appliance.

This migration procedure covers the following scenarios:

- One Panorama HA peer has a managed collector (Log Collector) in a Collector Group.

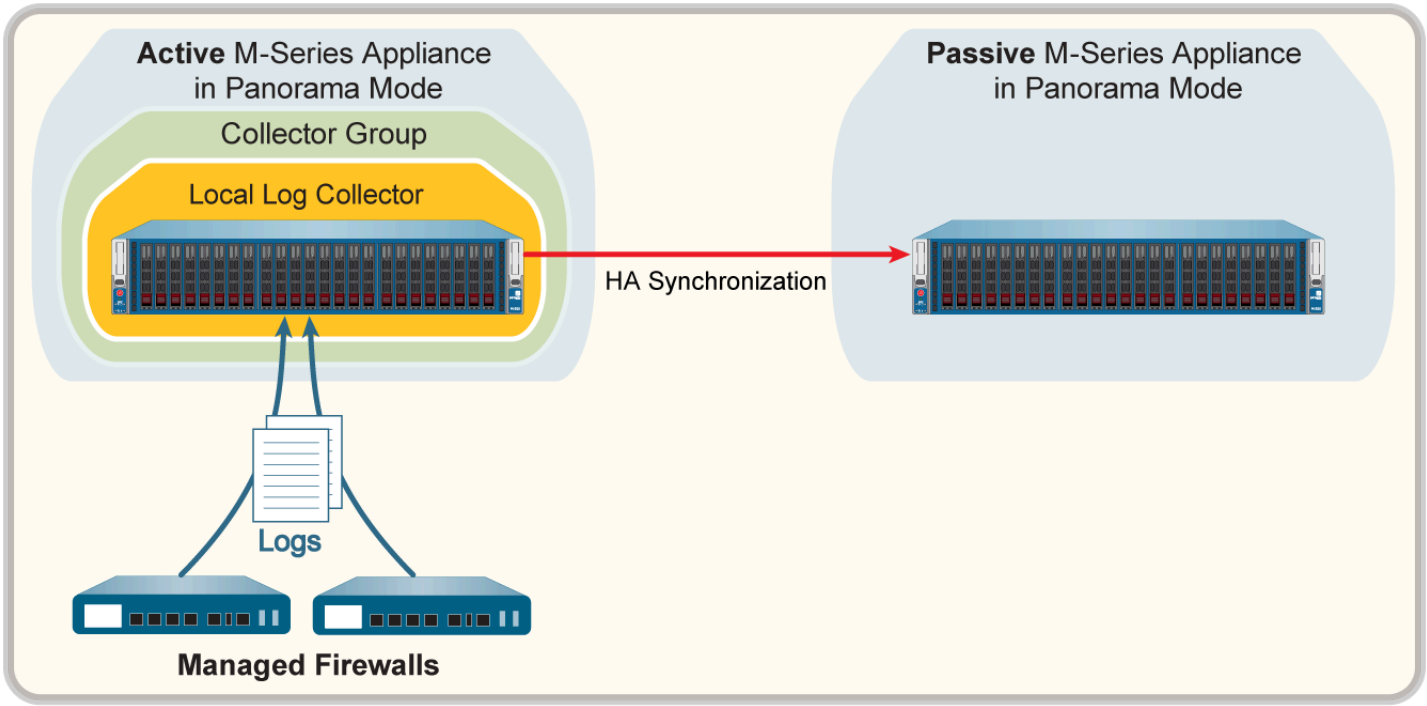


Figure 30: Panorama HA Peer with Collector Group

 ©2025 Palo Alto Networks, Inc.

- Both Panorama HA peers have managed collectors that belong to a single Collector Group. For details, see Multiple Local Log Collectors Per Collector Group.
- Both Panorama HA peers have a managed collector and each is assigned to a separate Collector Group. For details, see Single Local Log Collector Per Collector Group.

STEP 1 | Forward any logs on the SSD of the old M-Series appliance to an external destination if you want to preserve them.

The SSD stores the System and Config logs that Panorama and Log Collectors generate. You cannot move the SSD between M-Series appliances.

Configure Log Forwarding from Panorama to External Destinations.

STEP 2 | (RMA of Active Primary HA peer only) Reconfigure the high availability configuration of the Panorama HA peers to make the `Secondary` HA peer the `Primary` HA peer during the RMA process.

This step is required if you are replacing the `Primary` HA peer in an A/P HA configuration to help ensure the reset secure communication status and the certificate authority (CA) of the new M-Series appliance are not unintentionally synchronized to the existing peer in the HA configuration. When replacing the `Primary` HA peer, reconfiguring the HA election settings ensures Panorama management of devices remains uninterrupted during the RMA process.

Skip this step if you are replacing the `Secondary` HA peer in an A/P HA configuration.

1. Log in to the Panorama Web Interface of the `Primary` HA peer.
2. Select **Panorama** > **High Availability** and edit the Election Settings.
3. For the Priority, select **secondary** and click **OK**.
4. Select **Commit** and **Commit to Panorama**.
   The Panorama HA peer you are replacing is now the `Secondary` HA peer.
5. Log in to the Panorama Web Interface of the `Secondary` HA peer.
6. Select **Panorama** > **High Availability** and edit the Election Settings.
7. For the Priority, select **primary** and click **OK**.
8. Select **Commit** and **Commit to Panorama**.
   The previously `Secondary` Panorama is now the `Primary` HA peer.

STEP 3 | Suspend HA functionality on the Panorama HA peer you are replacing.

Step is required to help ensure the reset secure communication status and the certificate authority (CA) of the new M-Series appliance are not unintentionally synchronized to the existing peer in the HA configuration during the RMA process. This puts the Panorama HA peer in a `suspended` state.

1. Log in to the Panorama Web Interface of the Panorama HA peer you are replacing.
2. Select **Panorama** > **High Availability** and **Suspend local Panorama for high availability**.
3. Click **OK** to confirm suspending HA on Panorama HA peer.

 ©2025 Palo Alto Networks, Inc.

STEP 4 | Reset the secure connection setting on the Panorama HA peer you are replacing.

1. Log in to the Panorama CLI of the Panorama HA peer you are replacing.
2. Reset the secure connection state.

   This command resets all managed device connections and is irreversible.

   ```
   admin> request sc3 reset
   ```

3. Restart the management server on the Panorama HA peer you are replacing.

   ```
   admin> debug software restart process management-server
   ```

STEP 5 | Remove the RAID disks from the old M-Series appliance.

1. Power off the old M-Series appliance by pressing the Power button until the system shuts down.
2. Remove the disk pairs. For details, refer to the disk replacement procedure in the M-Series Appliance Hardware Reference Guides.

STEP 6 | Perform initial setup of the new M-Series appliance.

1. Rack mount the M-Series appliance. Refer to the M-Series Appliance Hardware Reference Guides for instructions.
2. Perform Initial Configuration of the M-Series Appliance.

   *If the old M-Series appliance used interfaces other than the MGT interface for Panorama services (such as log collection), you must define those interfaces during initial configuration of the new M-Series appliance (**Panorama** > **Setup** > **Interfaces**).*

3. Register Panorama.
4. Purchase and activate a Panorama support license or transfer licenses as follows only if the new M-Series appliance is the same hardware model as the old M-Series appliance.

 ©2025 Palo Alto Networks, Inc.

If the new M-Series appliance is a different model than the old M-Series appliance, you must purchase new licenses.

1. Log in to the Palo Alto Networks Customer Support web site.
2. Select the **Assets** tab and click the **Spares** link.
3. Click the Serial Number of the new M-Series appliance.
4. Click **Transfer Licenses**.
5. **Select** the old M-Series appliance and click **Submit**.

5. Activate a firewall management license. If you are migrating from an M-100 appliance to an M-500 appliance, enter the auth-code associated with the migration license.
6. Install Content and Software Updates for Panorama. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
7. Perform any additional configuration changes as needed.

   *If the old M-Series appliance used interfaces other than the MGT interface for Panorama services (such as log collection), you must define those interfaces on the new M-Series appliance (**Panorama** > **Setup** > **Interfaces**).*
8. Set Up HA on Panorama. The new M-Series appliance must have the same priority as the HA peer you are replacing.

   The new M-Series appliance must be added to the HA configuration as the `Secondary` HA peer. Adding the new M-series as the `Primary` HA peer forces synchronization of the reset secure communication setting status to the existing M-Series appliance, resulting in the interruption of Panorama management of devices.

STEP 7 | Insert the disks into the new M-Series appliance. For details, refer to the disk replacement procedure in the M-Series Appliance Hardware Reference Guides.

The disk carriers of the M-100 appliance are incompatible with those of the M-500 appliance. Therefore, when migrating between these hardware models, you must unscrew each disk from its old carrier and insert the disk in the new carrier before inserting the disk in the new appliance.

You must maintain the disk pair association. Although you can place a disk pair from slot A1/A2 on the old appliance into slot B1/B2 on the new appliance, you must keep the disks together in the same slot; otherwise, Panorama might not restore the data successfully.

STEP 8 | If the M-Series appliance was part of a Collector Group, verify that the decommissioned M-Series appliance serial number is still part of the correct Collector Group:

```
debug log-collector-group show name <Log CollectorGroup name>
```

 ©2025 Palo Alto Networks, Inc.

STEP 9 | Prepare the disks for migration.

Generating the metadata for each disk pair rebuilds the indexes. Therefore, depending on the data size, this process can take a long time to complete. To expedite the process, you can launch multiple CLI sessions and run the metadata regeneration command in each session to complete the process simultaneously for every pair. For details, see Regenerate Metadata for M-Series Appliance RAID Pairs.

1. Enable the disk pairs by running the following CLI command for each pair:

```
admin> request system raid add <slot> force no-format
```

For example:

```
admin> request system raid add A1 force no-format
admin> request system raid add A2 force no-format
```

The **force** and **no-format** arguments are required. The **force** argument associates the disk pair with the new appliance. The **no-format** argument prevents reformatting of the drives and retains the logs stored on the disks.

2. Generate the metadata for each disk pair.

This step may take up to 60 hours to complete depending on the volume of log data on the disks. The new M-Series appliance must be in an `active` state. Panorama may experience issues getting the assigned Log Collector ID if Panorama is in a `suspended` state.

If the new M-Series appliance is the `passive` HA peer, log in to the Panorama web interface *of the currently `active` HA peer and select* ***Panorama > High Availability*** *and* ***Suspend local Panorama for high availability****.*

After you successfully migrate the logs, ***Make local Panorama functional for high availability****.*

```
admin> request metadata-regenerate slot <slot_number>
```

For example:

```
admin> request metadata-regenerate slot 1
```

 ©2025 Palo Alto Networks, Inc.

STEP 10 | Configure the local Log Collector on the new M-Series appliance.

For all steps with commands that require a serial number, you must type the entire serial number; pressing the Tab key won't complete a partial serial number.

Don't enable the disks on the new M-Series appliance at this point. When you successfully migrate the logs, Panorama automatically enables the disks.

1. Configure the local Log Collector as a managed collector using the Panorama web interface or using the following CLI commands:

```
admin> configure
admin# set log-collector <log-collector_SN> deviceconfig
 system hostname <log-collector-hostname>
admin# exit
```

2. Commit your changes to Panorama. Don't commit the changes to the Collector Group just yet.

```
admin> configure
admin# commit
```

3. Verify that the local Log Collector is connected to Panorama and that the status of its disk pairs is present/available.

```
admin> show log-collector serial-number <log-collector_SN>
```

The disk pairs will display as disabled at this stage of the restoration process.

STEP 11 | Add a Log Collector with no disks to a Collector Group.

From this point, only commits that are required to complete the migration process on Panorama and the Log Collectors. Hold off making any other changes.

1. Access the Panorama CLI.
2. Overwrite Panorama restriction to allow Log Collector with no disk to be added to a Collector Group: **request log-migration-set-start**
3. Commit the overwritten restriction:

```
admin> configure
admin# commit force
```

 ©2025 Palo Alto Networks, Inc.

STEP 12 | Migrate the logs.

The existing M-Series appliance must be the `active` *HA peer before you can begin migrating logs. If the new M-Series appliance is the* `passive` *HA peer,* log in to the Panorama web interface *of the currently* `active` *HA peer and select* ***Panorama > High Availability*** *and* ***Suspend local Panorama for high availability****.*

After you successfully migrate the logs, ***Make local Panorama functional for high availability****.*

1. Access the Panorama CLI.
2. Add the new local Log Collector as a member of the Collector Group and commit your changes to Panorama.

```
admin# set log-collector-group <collector_group_name> logfwd-
setting collectors <SN_managed_collector>
admin# commit
admin# exit
```

   The old local Log Collector still appears in the list of members, because you haven't deleted it from the configuration.
3. For each disk pair, migrate the logs to the new appliance.

```
admin> request log-migration from <old_LC_serial_number> old-
disk-pair <log_disk_pair> to <new_LC_serial_number> new-disk-
pair <log_disk_pair>
```

   For example:

```
admin> request log-migration from 003001000010 old-disk-pair A
 to 00300100038 new-disk-pair A
```

4. Commit the changes to Panorama.

```
admin> configure
admin# commit
```

 ©2025 Palo Alto Networks, Inc.

STEP 13 | Reconfigure the Collector Group.

1. Use the web interface to assign the new Log Collector to the firewalls that forward logs (**Panorama** > **Collector Groups** > **Device Log Forwarding**). Give the new Log Collector the same priority in the firewall preference lists as the old Log Collector.

   You cannot use the CLI change the priority assignments of firewall preference lists.

2. Delete the old Log Collector from the Collector Group.

   ```
   admin# delete log-collector-group <group_name> logfwd-setting
    collectors <old_LC_serial_number>
   ```

   For example:

   ```
   admin# delete log-collector-group DC-Collector-Group logfwd-
   setting collectors 003001000010
   ```

3. Delete the old Log Collector from the Panorama configuration and commit your changes to Panorama.

   ```
   admin# delete log-collector <old_LC_serial_number>
   admin# commit
   admin# exit
   ```

4. Synchronize the configuration of the M-Series appliance HA peers.

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

5. Commit the Collector Group changes so that the managed firewalls can send logs to the new Log Collector.

   ```
   admin> commit-all log-collector-config log-collector-
   group <collector_group_name>
   ```

   For example:

   ```
   admin> commit-all log-collector-config log-collector-group DC-
   Collector-Group
   ```

 ©2025 Palo Alto Networks, Inc.

STEP 14 | Generate new keys on the new Log Collector.

This command is required in order to add the new Log Collector to the Collector Group and should only be run for the Collector Group of the Log Collector being replaced. This step deletes the existing RSA keys and allows Panorama to create new RSA keys.

1. Access the Panorama CLI.
2. Delete all RSA keys on the new Log Collector:

   request logdb update-collector-group-after-replace collector-group *<collector-group-name>*

   The process can take up to 10 minutes to completed.

STEP 15 | Confirm that SearchEngine Status is Active for all Log Collectors in the Collector Group.

Do not continue until SearchEngine Status is Active for all Log Collectors in the Collector Group. This will result in purging of logs from the Log Collector being replaced.

1. Access the Panorama CLI.
2. Show the Log Collector details by running the following commands either:
   - On Panorama for all Log Collectors:

     show log-collector all

     Alternatively, you can run the following command on each Dedicated Log Collector:

     ```
     show log-collector detail
     ```
3. Confirm that SearchEngine Status is Active.

```
Redistribution status:       none
```

```
Last commit-all: commit succeeded, current ring version 1
```

```
SearchEngine status:     Active
```

```
md5sum 4e5055a359f7662fab8f8c4f57e24525 updated at 2017/06/14
 09:58:19
```

 ©2025 Palo Alto Networks, Inc.

STEP 16 | On the new Log Collector, replace previous Log Collector serial number with the new Log Collector serial number.

You must replace the old Log Collector serial number with the new Log Collector serial number so that the new Log Collector will not run in to purging issues, resulting in the Log Collector being unable to purge old data from the migrated logs when necessary.

1. Access the Log Collector CLI.
2. Replace old Log Collector serial number with new Log Collector serial number:

   request log-migration-update-logger from *<old-log-collector-serial-number>* to *<new-log-collector-serial-number>*

STEP 17 | Restore HA functionality for the `suspended` Panorama HA peer to force the reset secure communication state changes to the Secondary HA peer.

1. Log in to the Panorama Web Interface of the new (`Secondary`) Panorama HA peer.
2. Select **Panorama** > **High Availability** and **Make local Panorama functional for high availability**.
3. Click **OK** to confirm restoring HA functionality on the new Panorama HA peer.

STEP 18 | Restart the management server on the new Panorama HA peer.

1. Log in to the Panorama CLI of the new (`Secondary`) Panorama HA peer.
2. Restart the management server.

   ```
   admin> debug software restart process management-server
   ```

STEP 19 | (RMA of Active Primary HA peer only) Restore the high availability configuration of the Panorama HA peers.

Skip this step if you are replaced the `Secondary` HA peer in an A/P HA configuration.

1. Log in to the Panorama Web Interface of the `Primary` HA peer.
2. Select **Panorama** > **High Availability** and edit the Election Settings.
3. For the Priority, select **secondary** and click **OK**.
4. Select **Commit** and **Commit to Panorama**.
5. Log in to the Panorama Web Interface of the `Secondary` HA peer.
6. Select **Panorama** > **High Availability** and edit the Election Settings.
7. For the Priority, select **primary** and click **OK**.
8. Select **Commit** and **Commit to Panorama**.

Migrate Log Collectors after Failure/RMA of Non-HA Panorama

If a system failure occurs on a Panorama management server that is not deployed in a high availability (HA) configuration, use this procedure to restore the configuration on the replacement Panorama and restore access to the logs on the Dedicated Log Collectors that it manages. The allowed migration scenarios vary by Panorama management server model:

 ©2025 Palo Alto Networks, Inc.

Old/Failed Panorama	New/Replacement Panorama
Panorama virtual appliance	• Panorama virtual appliance • M-200 appliance • M-500 appliance • M-600 appliance
M-100 appliance	• Panorama virtual appliance • M-200 appliance • M-500 appliance • M-600 appliance
M-500 appliance	• Panorama virtual appliance • M-200 appliance • M-500 appliance • M-600 appliance

Panorama maintains a ring file that maps the segments and partitions that Dedicated Log Collectors use to store logs. An M-Series appliance in Panorama mode stores the ring file on its internal SSD; a Panorama virtual appliance stores the ring file on its internal disk. When a system failure occurs, a non-HA Panorama cannot automatically recover the ring file. Therefore, when you replace Panorama, you must restore the ring file to access the logs on the Dedicated Log Collectors.

This procedure requires that you backed up and exported your Panorama configuration before the system failure occurred.

Palo Alto Networks recommends deploying Panorama in an HA configuration. The active Panorama peer automatically synchronizes the ring file to the passive peer in an HA configuration, thereby maintaining access to logs on the Dedicated Log Collectors even if you must replace one of the peers.

STEP 1 | Perform initial setup of the new Panorama appliance.

1. Set Up the M-Series Appliance or Set Up the Panorama Virtual Appliance based on your needs. If you are setting up a new M-Series appliance, refer to the M-Series Appliance

 ©2025 Palo Alto Networks, Inc.

Hardware Reference Guides for instructions on how to rack mount the new M-Series appliance.

2. Perform Initial Configuration of the M-Series Appliance or Perform Initial Configuration of the Panorama Virtual Appliance.

   *If the old M-Series appliance used interfaces other than the MGT interface for Panorama services (such as log collection), you must define those interfaces during initial configuration of the new M-Series appliance (**Panorama** > **Setup** > **Interfaces**). The Panorama virtual appliance does not support interfaces other than MGT.*

3. Register Panorama.
4. Transfer licenses as follows only if the new Panorama appliance is the same model as the old appliance. Otherwise, you must purchase new licenses.
   1. Log in to the Palo Alto Networks Customer Support web site.
   2. Select the **Assets** tab and click the **Spares** link.
   3. Click the Serial Number of the new M-Series appliance.
   4. Click **Transfer Licenses**.
   5. **Select** the old appliance and click **Submit**.
5. Activate a Panorama Support License.
6. Activate a firewall management license.
7. Install Content and Software Updates for Panorama.

   The M-500 appliance requires Panorama 7.0 or a later release. M-200 and M-600 appliances require Panorama 8.1. For important details about software versions, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Restore the configuration from the old Panorama to the replacement Panorama.

1. Log in to the new Panorama and select **Panorama** > **Setup** > **Operations**.
2. Click **Import named Panorama configuration snapshot**, **Browse** to the backup configuration file, and click **OK**.
3. Click **Load named Panorama configuration snapshot**, select the **Name** of the file you just imported, and click **OK**.

   *To replace an RMA Panorama, make sure you **Retain Rule UUIDs** when you load the named Panorama configuration snapshot. If you do not select this option, Panorama removes all previous rule UUIDs from the configuration snapshot and assigns new UUIDs to the rules on Panorama, which means it does not retain information associated with the previous UUIDs, such as the policy rule hit count.*
4. Select **Commit** > **Commit to Panorama** and **Commit** your changes.
5. Select **Panorama** > **Managed Collectors** and verify that the Connected column displays a check mark for the Dedicated Log Collector.

   If the Dedicated Log Collector doesn't appear, you must reconfigure it and its Collector Group as described in the next step. Otherwise, skip the following step to Fetch the ring file to restore access to the logs stored on the Dedicated Log Collector.

STEP 3 | Reconfigure the Dedicated Log Collector and Collector Group if they are missing on Panorama.

1. Access the CLI of the Dedicated Log Collector and enter the following commands to display the name of its Collector Group.

   1. Enter the command:

   ```
   > request fetch ring from log-collector <serial_number>
   ```

   The following error will display:

   ```
   Server error: Failed to fetch ring info from <serial_number>
   ```

   2. Enter the command:

   ```
   > less mp-log ms.log
   ```

   The following error will display:

   ```
   Dec04 11:07:08 Error:
    pan_cms_convert_resp_ring_to_file(pan_ops_cms.c:3719):
   ```

 ©2025 Palo Alto Networks, Inc.

```
 Current configuration does not contain group CA-Collector-
Group
```

In this example, the error message indicates that the missing Collector Group has the name CA-Collector-Group.

2. Configure the Collector Group and assign the Dedicated Log Collector to it.

```
> configure
# set log-collector-group <collector-group-name>
# set log-collector-group <collector-group-name> logfwd-
setting
collector <serial-number>
```

3. Commit the changes to Panorama but not to the Collector Group.

```
# commit
# exit
```

STEP 4 | Fetch the ring file to restore access to the logs stored on the Dedicated Log Collector.

1. Access the CLI of the new Panorama.
2. Fetch the ring file:

```
> request fetch ring from log-collector <serial-number>
```

For example:

```
> request fetch ring from log-collector 123456789012
```

If you don't know the serial number of the Dedicated Log Collector, log in to its CLI and enter the ***show system info*** *operational command.*

3. Commit your changes to the Collector Group.

```
> commit-all log-collector-config log-collector-
group <collector-group-name>
```

Regenerate Metadata for M-Series Appliance RAID Pairs

When a system failure occurs on the M-700, M-600, M-500, M-300, or M-200 appliance and you need to physically move the disks from one appliance to another, regenerating the metadata is necessary. The metadata is required to locate logs on the disk; when a user issues a log query, the query consults this metadata to access the requested log data.

For each configured RAID disk pair in the M-Series appliance, you must access the appliance CLI and run the following command to regenerate the metadata:

```
> request metadata-regenerate slot <slot_number>
```

 ©2025 Palo Alto Networks, Inc.

For example:

```
> request metadata-regenerate slot 1
```

The size of the RAID disks determines how long metadata regeneration takes. On average, it takes an hour for every 100GB. When you run the command, the CLI session is locked until the command is fully executed. You can use multiple CLI sessions to save time. For example, to replace four RAID pairs of 1TB drives with a total of 4TB of log data, launch four CLI sessions and run the command in each session to regenerate metadata simultaneously for all the pairs/slots in about 10 hours.

During metadata regeneration, the Collector Group to which these disks belong is not available and the disk pair is not available for any logging or reporting operations (writes/queries). However, you can perform other tasks such as handling new firewall connections or managing configuration changes on the managed firewalls. All other Collector Groups that Panorama manages and that aren't part of this RMA process can perform the assigned logging and reporting functionality as normal.

View Log Query Jobs

You can view your log query jobs to investigate and better understand why querying log data is taking a longer than expected. To begin, you must first show all log query jobs run on the Panorama. After you identify the log query job that you need to investigate, use the job ID to view detailed information about the query to better understand why your log query is running into issues. When querying log data on Panorama the detailed job ID information is overwritten as new log query jobs are executed.

STEP 1 | Log in to the Panorama CLI.

STEP 2 | View the log query jobs executed on Panorama.

The CLI output includes general information about each executed log query such as the jod ID, when the query was run, the query state, the log database that was queried, the number of

 ©2025 Palo Alto Networks, Inc.

logs queried, how long (in ms) it took for the query to return results, the admin that executed the query. and any filters applied to the query.

```
admin@Panorama> show query jobs
```

```
admin@bingdot34> show query jobs

ID        Enqueue Time         State       Database  nlogs     Runtime(ms)    Us
er            Filter
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
42        2020/01/02 14:35:46  COMPLETE    threat    110       166.27         ad
min           ((((receive_time leq 'now')) and (((subtype eq 'file')) or ((subty
pe eq 'data')))) and ((receive_time in 'last-hour')))


41        2020/01/02 14:35:46  COMPLETE    system    110       163.84         ad
min           (((receive_time leq now)) and (receive_time in last-hour))

40        2020/01/02 14:35:46  COMPLETE    config    110       158.23         ad
min           (((receive_time leq now)) and (receive_time in last-hour))

39        2020/01/02 14:35:36  COMPLETE    config    110       162.58         ad
min           (((receive_time leq now)) and (receive_time in last-hour))

38        2020/01/02 14:35:36  COMPLETE    system    110       172.68         ad
min           (((receive_time leq now)) and (receive_time in last-hour))

37        2020/01/02 14:35:36  COMPLETE    threat    110       188.80         ad
min           ((((receive_time leq 'now')) and (((subtype eq 'file')) or ((subty
pe eq 'data')))) and ((receive_time in 'last-hour')))
```

STEP 3 | View details log query information about a specific job using the job ID.

```
admin@Panorama> show query jobid <Job ID>
```

```
admin@bingdot34> show query jobid 42

Serial          ID        State     Num Req   Num Proc  RTT(Max)      Avg Recs/R
TTSoftware Ver    CG                              Last Update Time
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
LOGDB           42        DONE      110       0         0.00      0.00
9.2.0           LOCAL                                 2020/01/02 14:35:46

PODABCD12       42        FAILED    110       0         0.00      0.00
9.2.0           PODABCD12                             2020/01/02 14:35:46
```

 ©2025 Palo Alto Networks, Inc.

Replace an RMA Firewall

To minimize the effort required to restore the configuration on a managed firewall involving a Return Merchandise Authorization (RMA), replace the serial number of the old firewall with that of the new firewall on Panorama. To then restore the configuration on the replacement firewall, either import a firewall state that you previously generated and exported from the firewall or use Panorama to generate a *partial device state* for managed firewalls running PAN-OS 5.0 and later versions. By replacing the serial number and importing the firewall state, you can resume using Panorama to manage the firewall.

- Partial Device State Generation for Firewalls
- Before Starting RMA Firewall Replacement
- Restore the Firewall Configuration after Replacement

Partial Device State Generation for Firewalls

When you use Panorama to generate a partial device state, it replicates the configuration of the managed firewalls with a few exceptions for Large Scale VPN (LSVPN) setups. You create the partial device state by combining two facets of the firewall configuration:

- Centralized configuration that Panorama manages—Panorama maintains a snapshot of the shared policy rules and templates that it pushes to firewalls.
- Local configuration on the firewall—When you commit a configuration change on a firewall, it sends a copy of its local configuration file to Panorama. Panorama stores this file and uses it to compile the partial device state bundle.

> *In an LSVPN setup, the partial device state bundle that you generate on Panorama is not the same as the version that you export from a firewall (by selecting **Device** > **Setup** > **Operations** and clicking **Export device state**). If you manually ran the device state export or scheduled an XML API script to export the file to a remote server, you can use the exported device state in your firewall replacement workflow.*
>
> *If you did not export the device state, the device state that you generate in the replacement workflow will not include the dynamic configuration information, such as the certificate details and registered firewalls, that is required to restore the complete configuration of a firewall functioning as an LSVPN portal. See* Before Starting RMA Firewall Replacement *for more information.*

Panorama does not store the device state; you generate it on request using the CLI commands listed in Restore the Firewall Configuration after Replacement.

Before Starting RMA Firewall Replacement

- If the firewall belongs to an SD-WAN cluster, you must follow the workflow to replace an SD-WAN device when there is an RMA.
- The firewall you will replace must have PAN-OS 5.0.4 or a later version. Panorama cannot generate the *device state* for firewalls running older PAN-OS versions.

 ©2025 Palo Alto Networks, Inc.

- Record the following details about the firewall you will replace:
  - **Serial number**—You must enter the serial number on the Palo Alto Networks Customer Support web site to transfer the licenses from the old firewall to the replacement firewall. You will also enter this information on Panorama, to replace all references to the old serial number with the new serial number of the replacement firewall.
  - (Recommended) **PAN-OS version and the content database version**—Installing the same software and content database versions, including the URL database vendor, enables you to create the same state on the replacement firewall. If you decide to install the latest version of the content database, you might notice differences because of updates and additions to the database. To determine the versions installed on the firewall, access the firewall System logs stored on Panorama.
- Prepare the replacement firewall for deployment. Before you import the device state bundle and restore the configuration, you must:
  - Verify that the replacement firewall is the same model as the old firewall and is enabled for similar operational capability. Consider the following operational features: must the replacement firewall have multiple virtual systems, support jumbo frames support, or operate in CC or FIPS mode?
  - Configure network access, transfer the licenses, and install the appropriate PAN-OS and content database versions.
- You must use the Panorama CLI to complete this firewall replacement process, and therefore your administrator account must have the superuser or panorama-admin user role.
- If you have an LSVPN configuration, and are replacing a Palo Alto Networks firewall deployed as a satellite or as an LSVPN portal, the dynamic configuration information that is required to restore LSVPN connectivity will not be available when you restore the partial device state generated on Panorama. If you followed the recommendation to frequently generate and export the device state for firewalls in an LSVPN configuration, use the device state that you previously exported from the firewall itself instead of generating one on Panorama.

  If you have not manually exported the device state from the firewall, and need to generate a partial device state on Panorama, the missing dynamic configuration impacts the firewall replacement process as follows:
  - **If the firewall you are replacing is a GlobalProtect portal** that is explicitly configured with the serial number of the satellites (**Network** > **GlobalProtect** > **Portals** > **Satellite Configuration**), when restoring the firewall configuration, although the dynamic configuration is lost, the portal firewall will be able to authenticate the satellites successfully. The successful authentication will populate the dynamic configuration information and LSVPN connectivity will be reinstated.
  - **If you are replacing a satellite firewall**, it will not be able to connect and authenticate to the portal. This failure occurs either because the serial number was not explicitly configured on the firewall (**Network** > **GlobalProtect** > **Portals** > **Satellite Configuration**) or, if the serial number was explicitly configured, because the serial number of the replaced firewall does not match that of the old firewall. To restore connectivity after importing the device state bundle, the satellite administrator must log in to the firewall and enter the credentials

 ©2025 Palo Alto Networks, Inc.

(username and password) for authenticating to the portal. After authentication, the dynamic configuration required for LSVPN connectivity is generated on the portal.

However, if the firewall was configured in a high availability configuration, after restoring the configuration, the firewall will automatically synchronize the running configuration with its peer and attain the latest dynamic configuration required to function seamlessly.

Restore the Firewall Configuration after Replacement

To restore the firewall configuration on the new firewall, you will first perform initial configuration on the new firewall, including setting the operational mode, upgrading the PAN-OS software and content release version to match what was installed on the old firewall. You will then export the device state of the old firewall from Panorama and import it onto the new firewall. Finally, you will go back to Panorama to validate that the new firewall has connected and then sync it with Panorama.

If the firewall belongs to an SD-WAN cluster, you must follow the workflow to replace an SD-WAN device when there is an RMA.

STEP 1 | Perform initial configuration on the new firewall and verify network connectivity.

Use a serial port connection or a Secure Shell (SSH) connection to add an IP address, a DNS server IP address, and to verify that the new firewall can access the Palo Alto Networks updates server.

STEP 2 | (Optional) Set the Operational mode on the new firewall to match that on the old firewall.

A serial port connection is required for this task.

1. Enter the following CLI command to access maintenance mode on the firewall:

```
> debug system maintenance-mode
```

2. For Operational mode, select **Set FIPS Mode** or **Set CCEAL 4 Mode** from the main menu.

STEP 3 | Retrieve the license(s) on the new firewall.

Enter the following command to retrieve the licenses:

```
> request license fetch
```

STEP 4 | (Optional) Match the operational state of the new firewall with that of the old firewall. For example, enable multi-virtual system (multi-vsys) capability for a firewall that was enabled for multi-vsys capability.

Enter the commands that pertain to your firewall settings:

```
> set system setting multi-vsys on
> set system setting jumbo-frame on
```

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Upgrade the PAN-OS version on the new firewall.

You must upgrade to the same PAN-OS installed on the old firewall. You must upgrade the content release versions to the same or later version that is installed on the old firewall.

Enter the following commands:

1. To upgrade the content release version:

```
> request content upgrade download latest
> request content upgrade install version latest
```

2. To upgrade the anti-virus release version:

```
> request anti-virus upgrade download latest
> request anti-virus upgrade install version latest
```

3. To upgrade the PAN-OS software version:

```
> request system software download version <version>
> request system software install version <version>
```

STEP 6 | Go to the Panorama CLI and export the device state bundle from the old firewall to a computer using Secure Copy (SCP) or TFTP (you cannot do this from the web interface).

If you manually exported the device state from the firewall, you can skip this step.

The export command generates the device state bundle as a tar zipped file and exports it to the specified location. This device state will not include the LSVPN dynamic configuration (satellite information and certificate details).

Enter one of the following commands:

```
> scp export device-state device <old serial#> to <login>
 @ <serverIP>: <path>
```

or

```
> tftp export device-state device <old serial#> to <serverIP>
```

 ©2025 Palo Alto Networks, Inc.

STEP 7 | Replace the serial number of the old firewall with that of the new replacement firewall on Panorama.

By replacing the serial number on Panorama you allow the new firewall to connect to Panorama after you restore the configuration on the firewall.

1. Enter the following command in Operational mode:

```
> replace device old <old SN#> new <new SN#>
```

2. Enter Configuration mode and commit your changes.

```
> configure
# commit
```

3. Exit Configuration mode.

```
# exit
```

 ©2025 Palo Alto Networks, Inc.

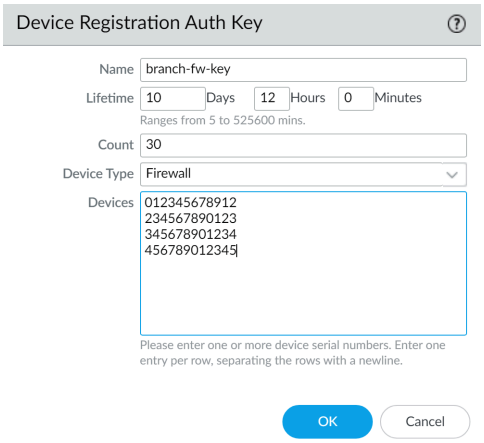
STEP 8 | (Optional) Create a device registration auth key on Panorama.

This step is required if no valid device registration auth key is created on Panorama. Skip this step if a valid device registration auth key is already created on Panorama.

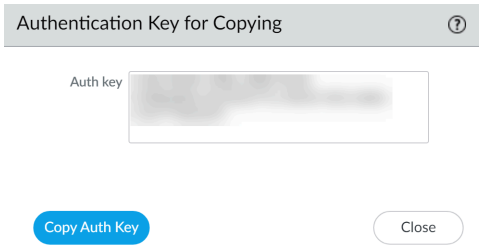
Exporting the device state bundle does not export the device registration auth key used to add the firewall to Panorama management. When you restore the firewall configuration after replacement, you must create a new device registration auth key to add the new firewall to Panorama.

1. Log in to the Panorama Web Interface.
2. Select **Panorama** > **Device Registration Auth Key** and **Add** a new authentication key.
3. 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.
   - **Count**—Specify how many times the authentication key may be used to onboard new firewalls.
   - **Device Type**—Specify that the authentication key is used to authenticate a **Firewall**.

     *Select **Any** to use the device registration auth key to onboard both firewalls and Log Collectors.*
   - (Optional) **Devices**—Enter one or more device serial numbers to specify for which firewalls the authentication key is valid.
4. Click **OK**.



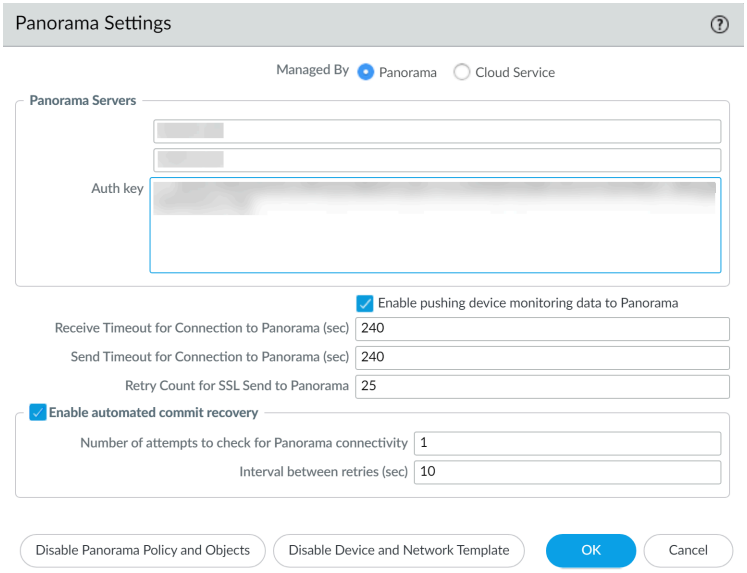
5. **Copy Auth Key** and **Close**.



 ©2025 Palo Alto Networks, Inc.

STEP 9 | On the new firewall, import the device state and add the device registration auth key.

1. Log in to the firewall web interface.
2. Select **Device** > **Setup** > **Operations** and click the **Import Device State** link in the Configuration Management section.
3. Browse to locate the file and click **OK**.
4. Select **Device** > **Setup** > **Management** and edit the Panorama Settings
5. Enter the **Auth key** you created on Panorama and click **OK**.



6. **Commit** your changes to the running configuration on the firewall.

STEP 10 | From Panorama, verify that you successfully restored the firewall configuration.

1. Access the Panorama web interface and select **Panorama** > **Managed Devices**.
2. Verify that the Connected column for the new firewall has a check mark.

STEP 11 | Synchronize the firewall with Panorama.

1. Access the Panorama web interface, select **Commit** > **Commit and Push** and **Edit Selections** in the Push Scope.
2. Select **Device Groups**, select the device group that contains the firewall, and **Include Device and Network Templates**.
3. Select **Collector Groups** and select the Collector Group that contains the firewall.
4. Click **OK** to save your changes to the Push Scope.
5. **Commit and Push** your changes.

If you need to generate reports for a period when the old firewall was still functional after you installed the new firewall, you must generate a separate query for each firewall serial number because replacing the serial number on Panorama does not overwrite the information in logs.

 ©2025 Palo Alto Networks, Inc.

Troubleshoot Commit Failures

If commit or push operation failures occur on Panorama, check for the following conditions. Review the troubleshooting steps to resolve your commit failures.

Symptom	Condition	Resolution
Panorama Commit Issues	Panorama commit lock does not release after a commit success.	Select **Panorama** > **Setup** > **Management** and edit the General Settings to disable **Automatically Acquire Commit Lock** and **Commit**.
	Panorama commit fails due to the following error: `Configured dailytrsum quota of 27 MB is less than the minimum needed 32 MB.`	Select **Panorama** > **Setup** > **Management** and edit the Logging And Reporting settings. Increase the `Quota %` value for the Daily Traffic Summary, Daily Threat Summary, Weekly Traffic Summary, and Weekly Threat Summary log storage to a value greater than 35 MB. Alternatively, you can **Restore Defaults**.
Panorama Push Issues	The Panorama management server has an earlier software version than the Dedicated Log Collectors or firewalls that it manages.	Upgrade the Panorama management server to the same or a higher software version than the managed firewalls, Log Collectors, and WildFire appliances and appliance clusters. For details, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.
	The ability to receive template and device groups configuration changes from Panorama is disabled on the firewall.	Access the firewall web interface, select **Device** > **Setup**, edit the Panorama Settings, and then click **Enable Device and Network Template** and **Enable Panorama Policy and Objects**.
	Configuration push from Panorama to managed firewalls fail due to the device registration authentication key issues.	Reset the secure connection state on the managed firewall experiencing push issues if: • A managed device disconnects from Panorama without reason and is not able to reconnect.

 ©2025 Palo Alto Networks, Inc.

Symptom	Condition	Resolution
		• You transitioned firewall management from Panorama running PAN-OS 10.1 or later release to a different Panorama running PAN-OS 10.1 or later release. • You reset Panorama or the managed firewall to factory default settings and managed firewalls are unable reconnect. In this case, you need to Recover Managed Device Connectivity to Panorama.
	Configuration push from Panorama fails due to local configuration changes pending on the firewall.	When you **Push to Devices** or **Commit to Panorama** from Panorama, **Edit Selections** and disable **Merge with Device Candidate Config**.

- Triage Commit Issues on Panorama
- Troubleshoot Template or Device Group Push Failures
- Troubleshoot Panorama Push Failure Due to Pending Local Firewall Changes

Triage Commit Issues on Panorama

Where Can I Use This?	What Do I Need?
• Panorama	☐ Device Management license ☐ Support license

Triage commit issues on the Panorama management server to identify the reason why your commit failed.

STEP 1 | Review the PAN-OS Release Notes to identify any limitations, changes to default behavior, or known issues that may cause your commits to fail.

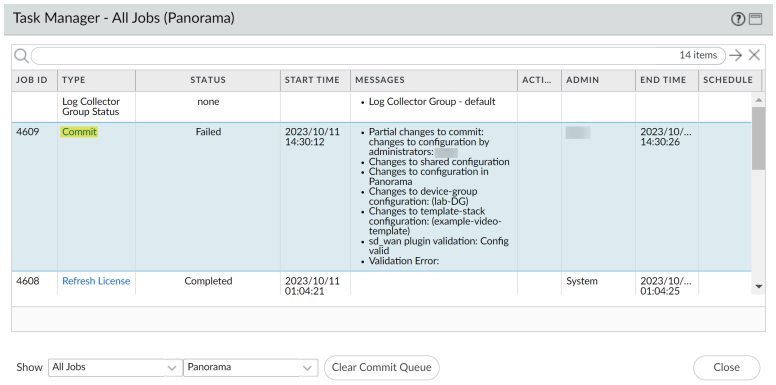
STEP 2 | Log in to the Panorama Web Interface.

 ©2025 Palo Alto Networks, Inc.

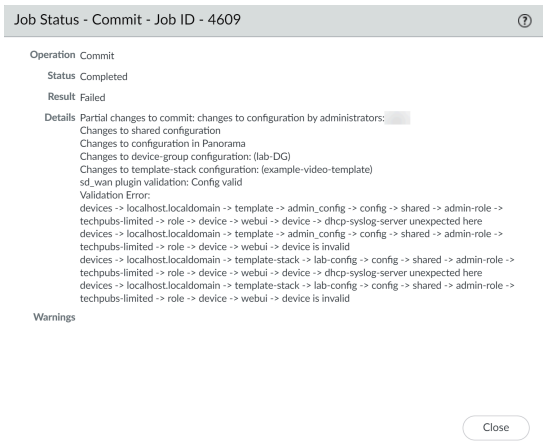
STEP 3 | Review the Panorama Task Manager.

1. Select **Tasks**.
2. Locate the commit operation and make note of the `Job ID`, and `Start Time`.

   In the `Type` column, click **Commit** to view the job details.



3. Review the `Validation Errors` to understand what is causing the commit to fail. This will help you understand if the commit is failing on Panorama or on the firewall.



 ©2025 Palo Alto Networks, Inc.

STEP 4 | Review the PAN-OS processes and process logs.

1. Log in to the Panorama CLI.
2. Enable debug logs on Panorama for more verbose log output

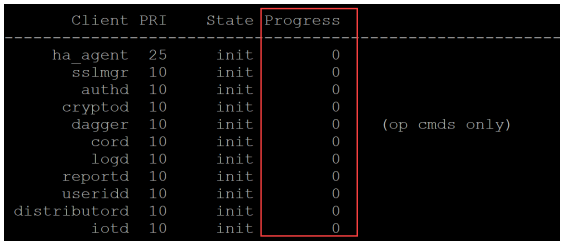
   ```
   admin> debug management-server
   ```

3. Review the management processes to see if any are in a degraded `State`.

   This tells you which management process logs are impacting the commit failure. This is denoted in the `Progress` column by an asterisk (*). The `Client` column displays the various management process related to a configuration commit.

   If this is showing no issues, then the commit failure is likely happening on the firewall. If that is the case, you will need to enter this command on the firewall CLI.

   ```
   admin> show management-clients
   ```

4. Review the Panorama log file to check for failures.

   In the below command, enter the `Client` experiencing issues.

   ```
   admin> less mp-log <client>.log
   ```

   Use the `Start Time` to locate the error causing the commit to fail. the reason the commit failed is indicated by `Commit Failed`.

5. Log in to the firewall CLI and review the device server processes.

   ```
   admin> less mp-log devsrvr.log
   ```

   This command also provides additional information about where the failure in the configuration commit process on the firewall. This will also show if External Dynamic Lists (EDL) are consuming too much device memory.

Troubleshoot Template or Device Group Push Failures

Where Can I Use This?	What Do I Need?
• Panorama	☐ Device Management license ☐ Support license

Troubleshoot the ability for your managed firewalls to receive template and device groups configuration changes from Panorama management server is disabled on the firewall.

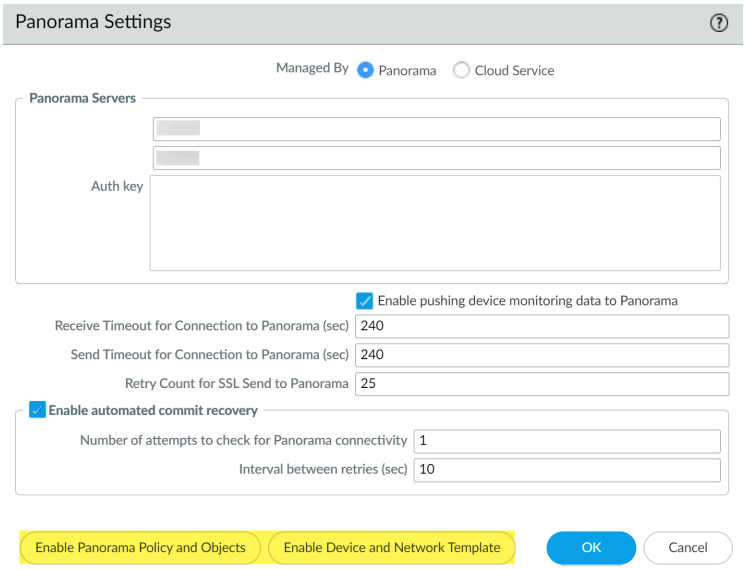
 ©2025 Palo Alto Networks, Inc.

STEP 1 | Log in to the firewall web interface.

STEP 2 | Select **Device** > **Setup** > **Management** and edit the Panorama Settings.

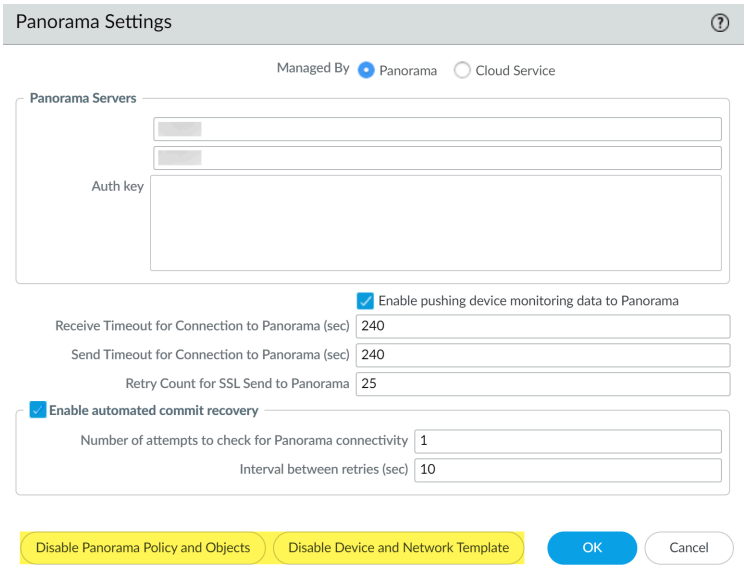
STEP 3 | Review the `Panorama Policy and Object` and `Device and Network Template` settings.

The example below depicts these Panorama Settings configured to block Panorama-pushed device group and template configurations.



STEP 4 | Click each setting to enable device group and template configuration pushes from Panorama.

Click **OK** when prompted to enable Panorama Policy and Objects and Device and Network Template settings. The example below depicts these Panorama Settings configured to allow Panorama-pushed device group and template configurations.



STEP 5 | Log in to the Panorama Web Interface and push the configuration changes from Panorama.

 ©2025 Palo Alto Networks, Inc.

Troubleshoot Panorama Push Failure Due to Pending Local Firewall Changes

Where Can I Use This?	What Do I Need?
• Panorama	❑ Device Management license ❑ Support license

By default, **Merge the Device Candidate Config** setting is enabled when you push a configuration from the Panorama management to a managed firewall. This setting commits any pending local configuration changes on the firewall alongside the configuration pushed from Panorama. In the event local configuration changes are made, the push may fail if a local candidate configuration made on the firewall is incomplete or invalid and this setting is enabled.

If you commonly make local configuration changes on your managed firewalls, you can disable this setting to prevent any local configuration changes from being committed alongside the configuration pushed from Panorama.

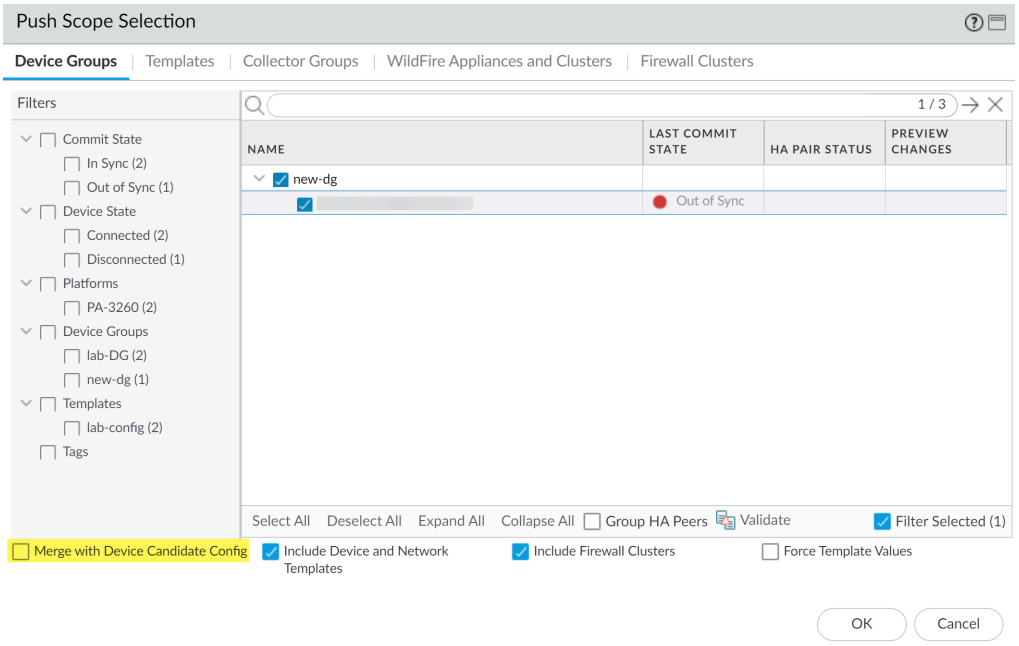
STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Select **Commit** > **Push to Devices** or **Commit and Push**.

STEP 3 | **Edit Selections**.

STEP 4 | Uncheck (disable) **Merge with Device Candidate Config**.

STEP 5 | Click **OK**.



STEP 6 | **Push**.

 ©2025 Palo Alto Networks, Inc.

Troubleshoot Registration or Serial Number Errors

On the M-700, M-600, M-500, M-300, or M-200 appliance, if the **Panorama** > **Support** page doesn't display support license details or the **Panorama** > **Setup** > **Management** page displays Unknown for the **Serial Number** even after you Register Panorama, perform the following steps:

STEP 1 | Record the Panorama serial number from the order fulfillment email that Palo Alto Networks sent when you placed your order for Panorama.

STEP 2 | Select **Panorama** > **Setup** > **Management** and edit the General Settings.

STEP 3 | Enter the **Serial Number** and click **OK**.

STEP 4 | Select **Commit** > **Commit to Panorama** and **Commit** your changes.

 ©2025 Palo Alto Networks, Inc.

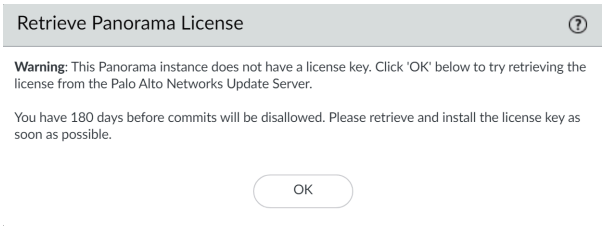
Troubleshoot Reporting Errors

If Panorama fails to generate a report, or the report is missing expected data, its content versions (such as the Applications database) might differ from those on the managed collectors and firewalls. The content versions on Panorama must be the same as or lower than the content versions on the managed collectors and firewalls. For details, see Panorama, Log Collector, Firewall, and WildFire Version Compatibility.

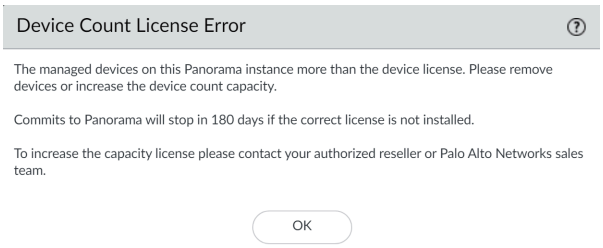
 ©2025 Palo Alto Networks, Inc.

Troubleshoot Device Management License Errors

After upgrading to PAN-OS 8.1, the Panorama virtual appliance will check if a device management licenses has been successfully installed. If a device management license has not been successfully installed, or the number of firewalls managed by the Panorama virtual appliance exceeds the device management license limit, you have 180 days to install a valid device management license. If no valid device management license has been installed, the following alert appears each time you log in to the Panorama web interface:



If the number of firewalls managed by the Panorama virtual appliance exceeds the device management license limit, the following alerts appears each time you log in to the Panorama web interface:



To resolve, install a valid device management license:

STEP 1 | Contact your Palo Alto Networks sales representative or your authorized reseller to purchase the appropriate device management license.

STEP 2 | Log in to the Panorama Web Interface.

STEP 3 | Activate/Retrieve a device management license based on whether the Panorama virtual appliance is online or offline.

- Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is Internet-connected.
- Activate/Retrieve a Firewall Management License when the Panorama Virtual Appliance is not Internet-connected.

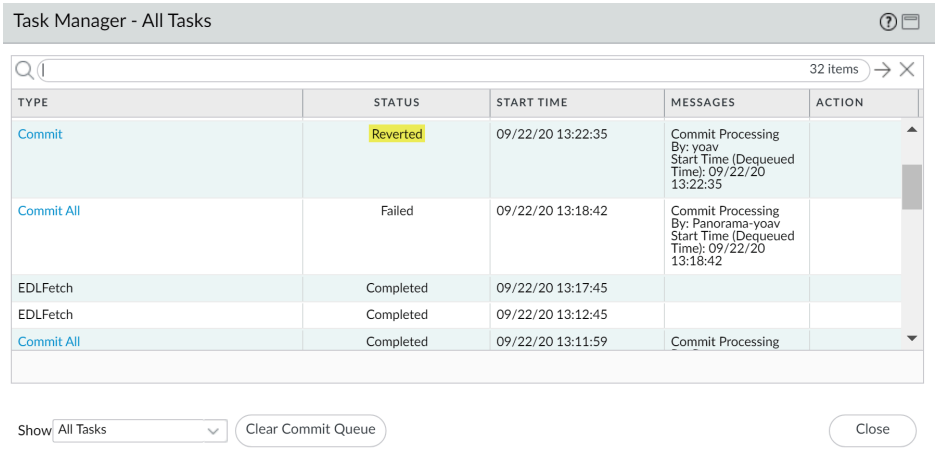
 ©2025 Palo Alto Networks, Inc.

Troubleshoot Automatically Reverted Firewall Configurations

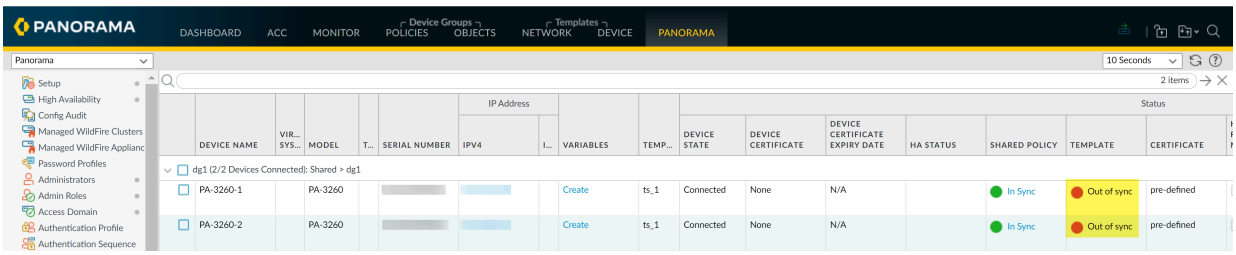
If your managed firewall automatically reverts its configuration due to a configuration change that caused a connection to break between the Panorama™ management server and the firewall, you can troubleshoot the out-of-sync firewalls to determine what changes were made and to determine what aspects of that last configuration push caused the firewall revert its configuration.

STEP 1 | Verify that the managed firewall automatically reverted to the last running configuration.

- On the firewall
  1. Launch the Firewall Web Interface.
  2. Click **Tasks** (bottom-right hand corner of the web interface).
  3. Verify that the last commit operation (either pushed from Panorama or committed locally) shows a `Reverted` status.



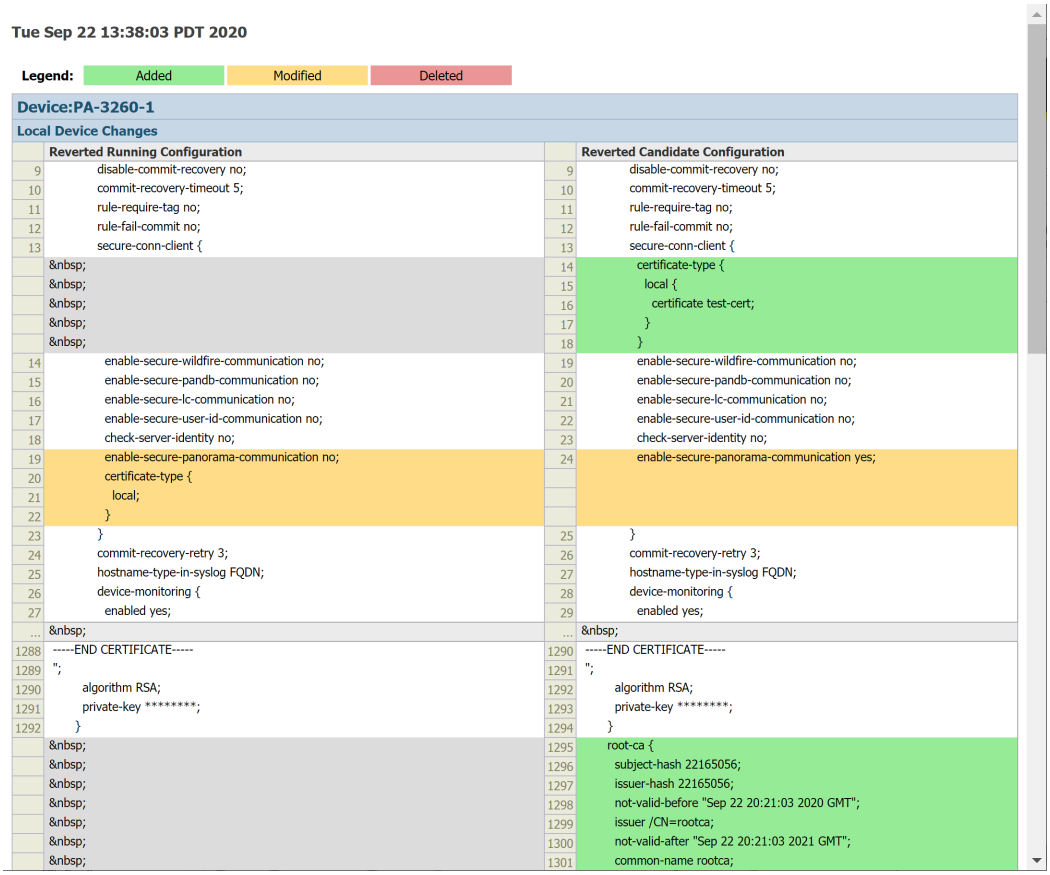
- On Panorama
  1. Log in to the Panorama Web Interface.
  2. Select **Panorama** > **Managed Devices** > **Summary**.
  3. View the Shared Policy and Template sync status. If you have recently pushed a configuration from Panorama to your managed firewalls and it reverted, the Shared Policy or Template display as `Out of Sync` (depending on what configuration changes were made).



STEP 2 | In the Last Merged Diff column for a managed firewall, **Show Last Merged Config Diff** ( ) to compare the current running configuration and the reverted configuration. In this

 ©2025 Palo Alto Networks, Inc.

example, a policy rule pushed from Panorama denied all traffic between the managed firewall and Panorama, which caused the firewall configuration to automatically revert.



Tue Sep 22 13:38:03 PDT 2020

Legend: Added | Modified | Deleted

Device:PA-3260-1

Local Device Changes

	Reverted Running Configuration		Reverted Candidate Configuration
9	disable-commit-recovery no;	9	disable-commit-recovery no;
10	commit-recovery-timeout 5;	10	commit-recovery-timeout 5;
11	rule-require-tag no;	11	rule-require-tag no;
12	rule-fail-commit no;	12	rule-fail-commit no;
13	secure-conn-client {	13	secure-conn-client {
		14	certificate-type {
		15	local {
		16	certificate test-cert;
		17	}
		18	}
14	enable-secure-wildfire-communication no;	19	enable-secure-wildfire-communication no;
15	enable-secure-pandb-communication no;	20	enable-secure-pandb-communication no;
16	enable-secure-lc-communication no;	21	enable-secure-lc-communication no;
17	enable-secure-user-id-communication no;	22	enable-secure-user-id-communication no;
18	check-server-identity no;	23	check-server-identity no;
19	enable-secure-panorama-communication no;	24	enable-secure-panorama-communication yes;
20	certificate-type {		
21	local;		
22	}		
23	}	25	}
24	commit-recovery-retry 3;	26	commit-recovery-retry 3;
25	hostname-type-in-syslog FQDN;	27	hostname-type-in-syslog FQDN;
26	device-monitoring {	28	device-monitoring {
27	enabled yes;	29	enabled yes;
...		...	
1288	-----END CERTIFICATE-----	1290	-----END CERTIFICATE-----
1289	";	1291	";
1290	algorithm RSA;	1292	algorithm RSA;
1291	private-key ********;	1293	private-key ********;
1292	}	1294	}
		1295	root-ca {
		1296	subject-hash 22165056;
		1297	issuer-hash 22165056;
		1298	not-valid-before "Sep 22 20:21:03 2020 GMT";
		1299	issuer /CN=rootca;
		1300	not-valid-after "Sep 22 20:21:03 2021 GMT";
		1301	common-name rootca;

STEP 3 | Modify configuration objects as needed as to not break the connection between the managed firewalls and Panorama before you re-push the configuration.

 ©2025 Palo Alto Networks, Inc.

View Task Success or Failure Status

Click the Task Manager icon  at the bottom right of the Panorama web interface to view the success or failure of a task. The Task Manager also displays a detailed message to help debug an issue. For details, see Use the Panorama Task Manager.

 ©2025 Palo Alto Networks, Inc.

Test Policy Match and Connectivity for Managed Devices

After you successfully push the device group and template stack configurations to your firewalls, Log Collectors, and WF-500 appliances, test that the correct traffic matches the policy rules pushed to your managed devices and that your firewalls can successfully connect to all appropriate network resources.

- Troubleshoot Policy Rule Traffic Match
- Troubleshoot Connectivity to Network Resources

Troubleshoot Policy Rule Traffic Match

To perform policy match tests for managed firewalls, test the policy rule configuration for your managed devices to ensure that the running configuration appropriately secures your network by allowing and denying the correct traffic. After the results are generated for traffic that was matched to configured rules, you can **Export to PDF** for auditing purposes.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Select **Panorama** > **Managed Devices** > **Troubleshooting** to perform a policy match.



*You may also run a policy match test from the **Policies** tab.*

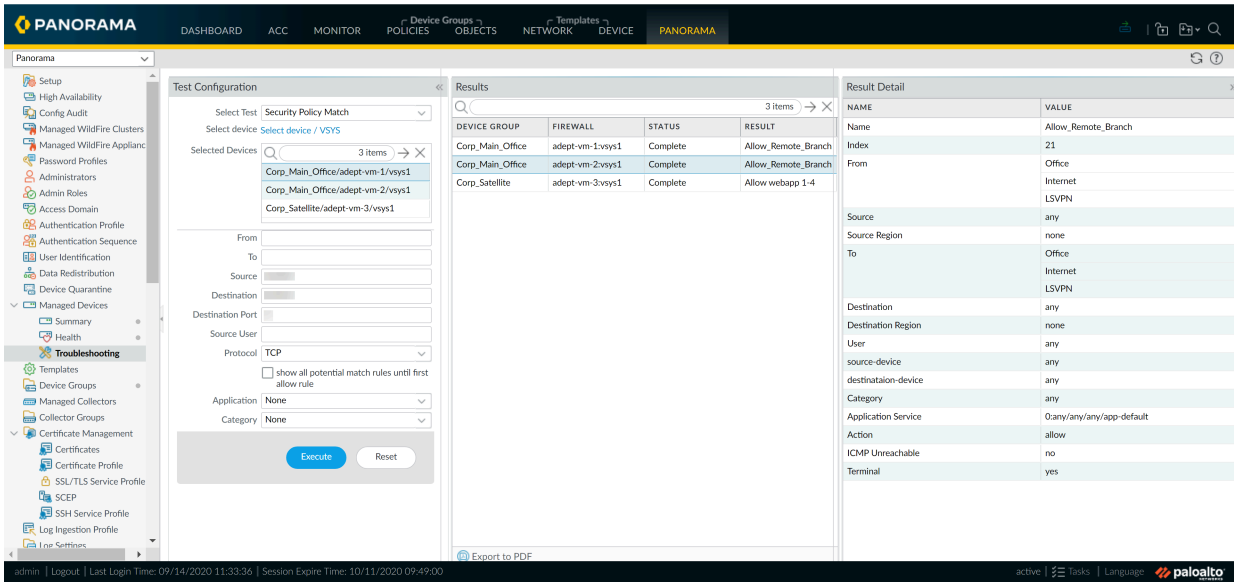
STEP 3 | Enter the required information to perform the policy match test. In this example, a Security policy match test is run.

1. Select **Security Policy Match** from the **Select Test** drop-down.
2. **Select device/VSYS** and select the managed firewalls to test.
3. Enter the Source IP address from which traffic originated.
4. Enter the Destination IP address of the target device for the traffic.
5. Enter the Protocol IP used for the traffic.
6. If necessary, enter any additional information relevant for your Security policy rule testing.

STEP 4 | **Execute** the Security policy match test.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Select the Security policy match Results to review the policy rules that match the test criteria.



Troubleshoot Connectivity to Network Resources

Perform connectivity tests for managed firewalls to ensure that your managed devices can connect to all appropriate network resources. Test the device configuration for your managed devices to ensure the running configuration appropriately secures your network by allowing you to verify that the configurations pushed to your managed devices still allow those devices to connect to resources such as your Log Collectors, configured External Dynamic Lists, and the Palo Alto Networks Update Server. Additionally, you can execute routing, WildFire®, Threat Vault, ping, and traceroute connectivity tests to verify that Panorama™ and managed devices can access any external network resources critical to the operation and security of your network. After the results are generated, you can **Export to PDF** for auditing purposes.

The Ping connectivity test is only supported for firewalls running PAN-OS 9.0 or later releases.

STEP 1 | Log in to the Panorama Web Interface.

STEP 2 | Select **Panorama** > **Managed Devices** > **Troubleshooting** to perform a connectivity test.

*You may also run a policy match test from the **Policies** tab.*

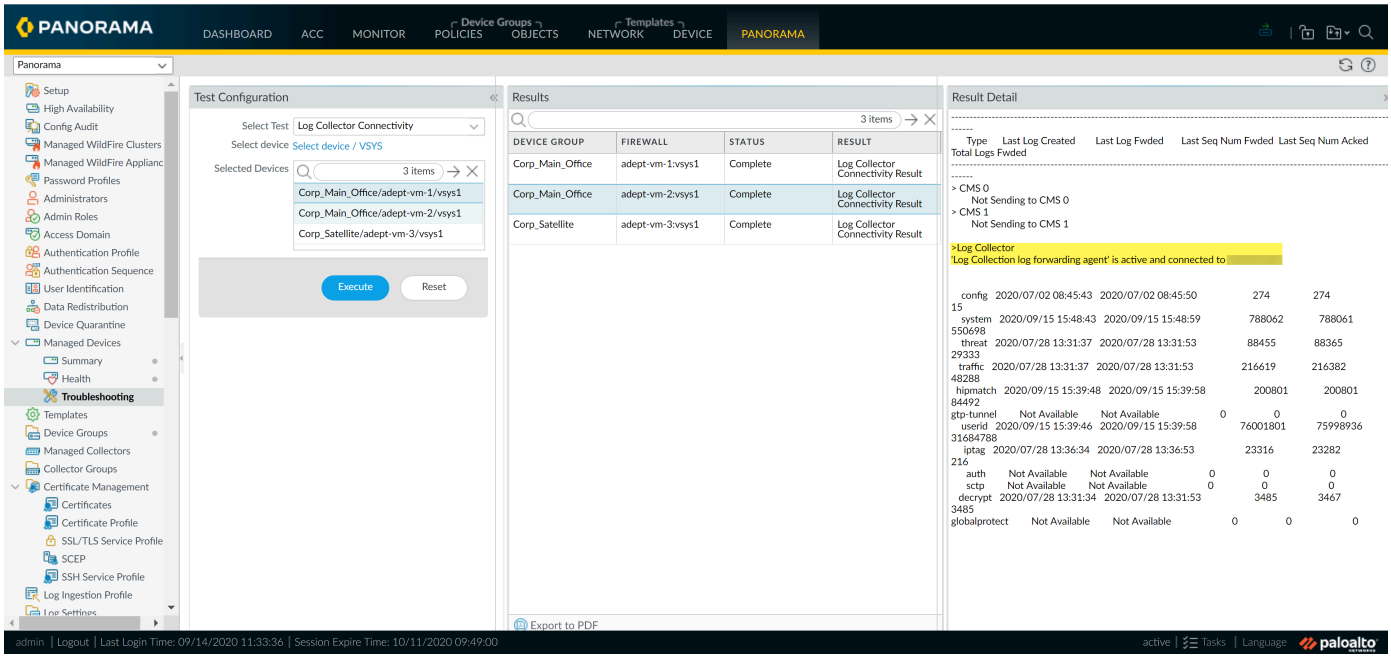
STEP 3 | Enter the required information to perform the connectivity test. In this example, a Log Collector Connectivity test is run.

1. Select **Log Collector Connectivity** from the **Select Test** drop-down.
2. **Select device/VSYS** and select the managed firewalls to test.
3. If necessary, enter any additional information relevant for your connectivity testing.

STEP 4 | **Execute** the Log Collector connectivity test.

 ©2025 Palo Alto Networks, Inc.

STEP 5 | Select the log collector connectivity Results to review the Log Collector connectivity status for the selected devices.



 ©2025 Palo Alto Networks, Inc.

Generate a Stats Dump File for a Managed Firewall

Generate a set of XML reports that summarize the network traffic over the last seven days for a single firewall managed by the Panorama™ management server or for all firewalls managed by Panorama. After you select a managed firewall and generate the stats dump file, you can download the stats dump file locally to your device.

The Palo Alto Networks or Authorized Partner systems engineer use the stat dump file to create a Security Lifecycle Review (SLR) and to perform security checkups after you successfully deploy your managed firewalls to help strength your security posture. The SLR highlights activity found on the network and the associated business or security risks that may be present. For more information on the SLR, contact your Palo Alto Networks or Authorized Partner systems engineer.



Stats dump file generation for multiple managed firewalls can take multiple hours to complete. During this time, you are unable to navigate from the stats dump file generation user interface so it is recommended to generate the stats dump file from the CLI so you can continue using the Panorama web interface.

Palo Alto Networks recommends generating a stats dump file for all managed firewalls from the Panorama CLI using the following command. Panorama must be able to reach your SCP or TFTP server to successfully export the stats dump file.

- ***SCP Server***

```
admin> scp export stats-dump to
 <username@hostname:SCP_export_path>
```

- ***TFTP Server***

```
admin> tftp export stats-dump to <tftp_host_address>
```

STEP 1 | Log in to the Panorama Web Interface.

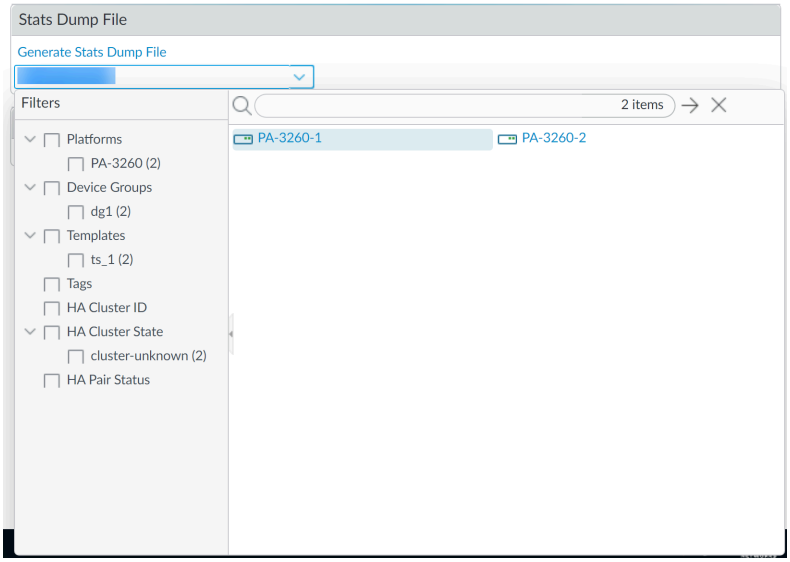
STEP 2 | Select **Panorama** > **Support** and navigate to the **Stats Dump File**.

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Select a managed firewall for which to generate a stats dump file.

It is recommended that you generate a stats dump file for a single managed firewall from the Panorama web interface.

A stats dump file is generated for **All devices** by default if you do not select a managed firewall.



STEP 4 | **Generate Stats Dump File**.

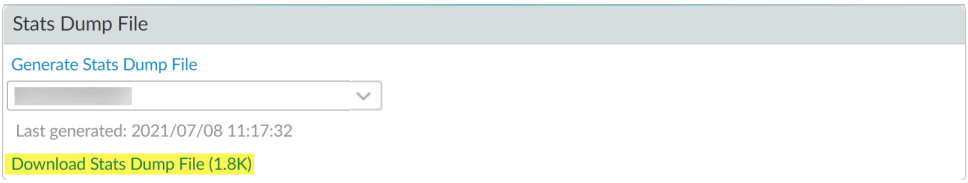
Click **Yes** when prompted to proceed generating the stats dump file.

A progress bar of the stats dump file generation status is displayed.

Generation may take up to an hour for a single managed firewall depending on the volume of log data. You are unable to navigate from the stats dump file generation status window during this time.

STEP 5 | Click **Download Stats Dump File** to download the stats dump file to your local device.

The downloaded stat dumps file is in a `tar.gz` file format.



 ©2025 Palo Alto Networks, Inc.

Recover Managed Device Connectivity to Panorama

PAN-OS 10.1 introduced the device registration authentication key to securely onboard managed firewalls, Dedicated Log Collectors, and WildFire appliance to the Panorama™ management server. The steps below describe how to recover the managed device connectivity to Panorama in the following scenarios:

- If a managed device disconnects from Panorama without reason and is not able to reconnect.
- You want to transition firewall management from a Panorama running PAN-OS 10.1 or later release to a different Panorama running PAN-OS 10.1 or a later release.
- If you reset Panorama or the managed firewall to factory default settings but the managed firewall is unable to connect to Panorama.

Recovering the managed device connectivity to Panorama applies only to managed devices that are running PAN-OS 10.1 or later release when onboarded to Panorama. The behavior described does not apply to managed devices running PAN-OS 10.0 and earlier releases or managed devices that were upgraded to PAN-OS 10.1 or later release while already managed by Panorama.

The following firewall platforms are not impacted by the described connectivity issues to Panorama.

- *Managed firewalls onboarded to Panorama using Zero Touch Provisioning (ZTP).*
- *CN-Series firewalls.*
- *Managed firewalls deployed on VMware NSX.*
- *VM-Series firewalls purchases from a public hypervisor marketplace. See PAYG firewalls for more information.*

STEP 1 | Reset the secure connection state of the managed device.

1. Log in to the managed device CLI.
   - Log in to the firewall CLI.
   - Log in to the Dedicated Log Collector CLI.
   - Log in to the WildFire appliance CLI.
2. Reset the secure connection state.

   This command resets the managed device connection and is irreversible.

   ```
   admin> request sc3 reset
   ```
3. Restart the management server on the managed device.

   ```
   admin> debug software restart process management-server
   ```

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Clear the secure connection state a managed device on Panorama and generate a new device registration authentication key.

Clearing the secure connection state for a managed device on Panorama is irreversible. This means that the managed device is disconnect and must be added back to Panorama.

1. Log in to the Panorama CLI.
2. Reset the secure connection state of a managed device on Panorama.

   This command resets the managed device connection to Panorama and is irreversible.

   ```
   admin> clear device-status deviceid <device_SN>
   ```

   Where **<device_SN>** is the serial number of the managed device you want to clear the connection state for.
3. Create a new device registration authentication key on Panorama.

   ```
   admin> request authkey add devtype <fw_or_lc) count
    <device_count> lifetime <key_lifetime> name <key_name> serial
    <device_SN>
   ```

   *The **devtype** and **serial** arguments are optional. Omit these two arguments to make a general use device registration authentication key that is not specific to a device type or device serial number.*
4. Verify the device registration authentication key you created are successfully created.

   ```
   admin> request authkey list *
   ```

   Make note of the key Name. The key Name is required to obtain the device registration authentication key required for onboarding.

   ```
   yoav@M-200(primary-active)> request authkey list *

   Name                          Cnt  Type     Expiry  Serial #
   ------------------------------------------------------------------------
   auth-key-fws                  100  fw       170470  1234567890,2345678901,3456789012,4567890123
   auth-key-lc                   100  lc       172852  0987654321,9876543210,8765432109,7654321098
   ```
5. Copy the device registration authentication Key value.

   ```
   admin> request authkey list <key_name>
   ```

   ```
   yoav@M-200(primary-active)> request authkey list auth-key-fws

   Name       : auth-key-fws
     Count    : 100
     Type     : fw
     Lifetime : 169813s
     Key      : [redacted]
     Serial # : 1234567890,2345678901,3456789012,4567890123
   ```

 ©2025 Palo Alto Networks, Inc.

STEP 3 | Add the device registration authentication key you created to the managed device.

1. Log in to the managed device CLI.
   - Log in to the firewall CLI.
   - Log in to the Dedicated Log Collector CLI.
   - Log in to the WildFire appliance CLI.
2. Add the device registration authentication key you created in the previous step.

   ```
   admin> request authkey set <auth_key>
   ```

   For **<auth_key>**, enter the Key value you copied in the previous step.

STEP 4 | Verify the managed device connectivity to Panorama.

```
admin> show panorama-status
```

Verify that the Panorama server `Connected` status displays `yes`.

If this procedure does not resolve the connectivity issue for your managed device, you must contact Palo Alto Networks Customer Support for further assistance as a full reset of all managed device connections on Panorama may be required.

 ©2025 Palo Alto Networks, Inc.

Restore an Expired Device Certificate

The device certificate installed on your Panorama™ management server, Dedicated Log Collector, or managed firewalls have a 90 day lifetime. Panorama, Dedicated Log Collectors, and managed firewalls with the device certificate installed automatically attempt to reinstall the device certificate 15 days before the certificate expires. However, you have the ability to manually reinstall the device certificate if it fails to reinstall automatically.

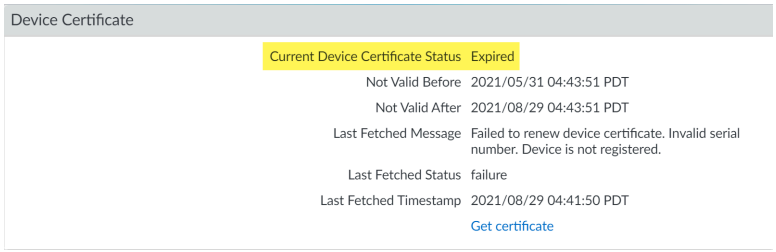
STEP 1 | Log in to the Panorama Web Interface.

 ©2025 Palo Alto Networks, Inc.

STEP 2 | Review the device certificate status for Panorama, Dedicated Log Collectors, and managed firewalls.

1. To review the Panorama device certificate status, select **Panorama** > **Setup** > **Management** and review the `Current Device Certificate Status` in the Device Certificate Section.

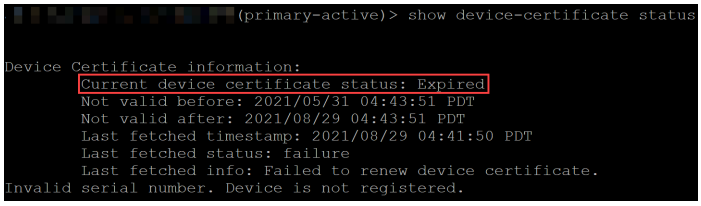
   The `Current Device Certificate Status` displays `Expired`.



2. To review the Dedicated Log Collector device certificate status, log in to the Dedicated Log Collector CLI and enter the following command:

```
admin>show device-certificate status
```

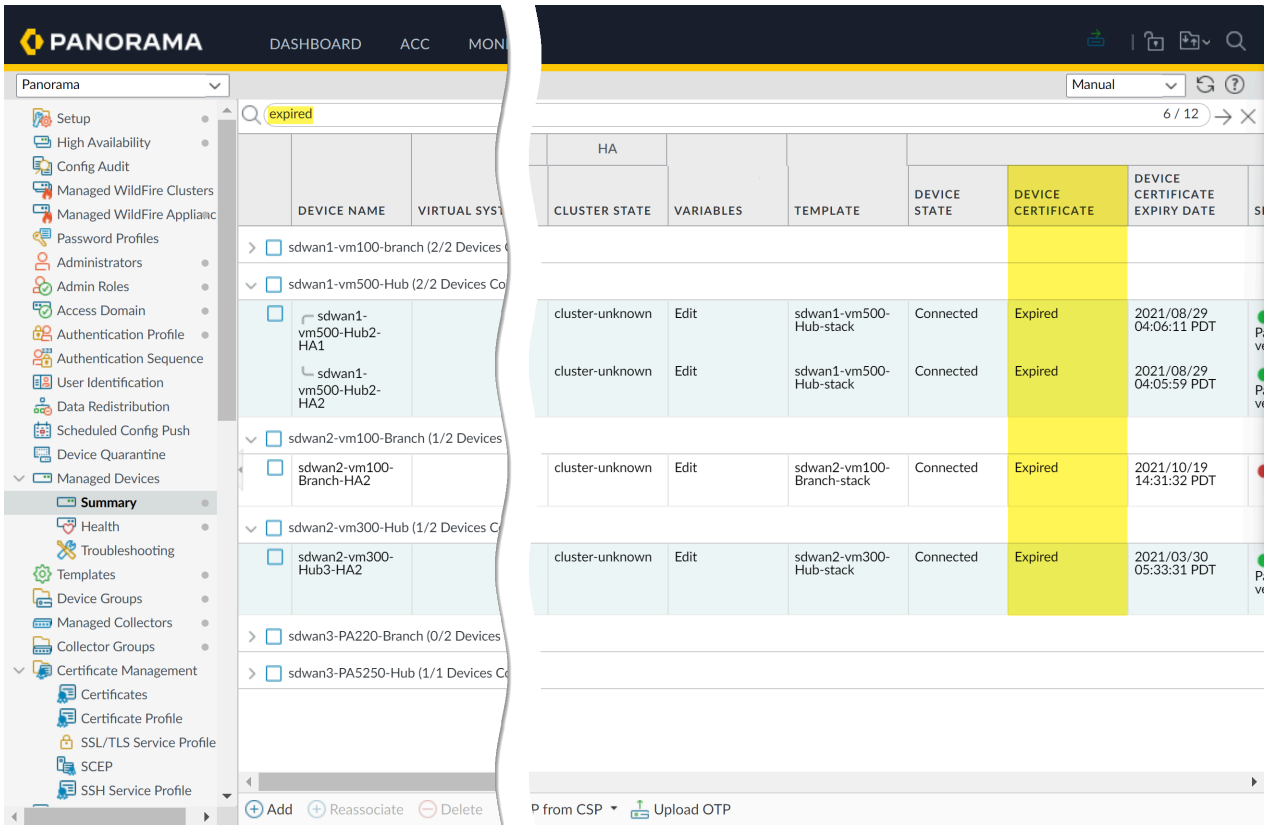
   The `Current Device Certificate Status` displays `Expired`.



3. To review the managed firewall device certificate status, select **Panorama** > **Managed Firewalls** > **Summary** and filter for **`expired`**.

   The Device Certificate column displays the current `Expired` device certificate status.

 ©2025 Palo Alto Networks, Inc.



STEP 3 | Reinstall the expired device certificate on Panorama, Dedicated Log Collectors, or managed firewalls.

If the ***request certificate fetch otp <otp_value>*** *command is not available, it means the Panorama, Log Collector, or managed firewall is a Trusted Platform Module (TPM) device.*

To restore the device certificate for a TPM device, run the following command:

```
request certificate fetch
```

- Install the Panorama Device Certificate
- Install the Device Certificate for a Dedicated Log Collector
- Install the Device Certificate for Managed Firewalls

 ©2025 Palo Alto Networks, Inc.

 ©2025 Palo Alto Networks, Inc.