



the network security company™

**Palo Alto Networks®**  
**Panorama Administrator's Guide**

**Panorama 5.1**

## Contact Information

### Corporate Headquarters:

Palo Alto Networks  
3300 Olcott Street  
Santa Clara, CA 95054

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

## About this Guide

This guide describes how to set up and use Panorama for centralized management; it is intended for administrators who want the basic framework to quickly set up the Panorama virtual appliance or the M-100 appliance for centralized administration of Palo Alto Networks firewalls.

If you have a M-100 appliance, this guide takes over after you complete [rack mounting your M-100 appliance](#).

For more information, refer to the following sources:

- ▲ [Palo Alto Networks Administrator's Guide](#)— for instructions on configuring the features on the firewall. The Palo Alto Networks Administrator's Guide will also help you with Panorama configuration items that are similar to the firewall and are not covered in this guide.
- ▲ <https://live.paloaltonetworks.com>— for access to the knowledge base, complete documentation set, discussion forums, and videos.
- ▲ <https://support.paloaltonetworks.com>— for contacting support, for information on the support programs, or to manage your account or devices.

To provide feedback on the documentation, please write to us at: [documentation@paloaltonetworks.com](mailto:documentation@paloaltonetworks.com).

Palo Alto Networks, Inc.

[www.paloaltonetworks.com](http://www.paloaltonetworks.com)

© 2013 Palo Alto Networks. All rights reserved.

Palo Alto Networks, PAN-OS, and Panorama are trademarks of Palo Alto Networks, Inc. All other trademarks are the property of their respective owners.

P/N 810-000150-00B

# Table of Contents

---

<b>Panorama Overview</b>	<b>1</b>
About Panorama	2
Panorama Platforms	3
About Centralized Configuration and Deployment Management	4
Context Switch—Device or Panorama	4
Templates	4
Device Groups	5
About Centralized Logging and Reporting	8
Logging Options	8
Managed Collectors and Collector Groups	8
Using Multiple Log Collectors in a Collector Group	9
Centralized Reporting	11
About Role-Based Access Control	12
Administrative Roles	12
Authentication Profiles and Sequences	13
Access Domains	13
Administrative Authentication	13
Panorama Recommended Deployments	15
Panorama for Centralized Management and Reporting	15
Panorama in a Distributed Log Collection Architecture	16
Plan Your Deployment	17
Deploy Panorama: Task Overview Checklist	19
<b>Set Up Panorama</b>	<b>21</b>
Set Up the Panorama Virtual Appliance	22
Prerequisites	22
Install Panorama on the ESX(i) Server	23
Perform Initial Configuration	24
Expand Log Storage Capacity on the Panorama Virtual Appliance	27
Set Up the M-100 Appliance	31
Perform Initial Configuration	32
Set Up the M-100 Appliance in Log Collector Mode	34
Increase Storage Capacity on the M-100 Appliance	35
Migrate from a Panorama Virtual Appliance to an M-100 Appliance	38
Prerequisites	38
Planning Considerations	39
Perform the Migration	40
Resume Managing the Devices	42

Install Licenses . . . . .	43
Register Panorama . . . . .	43
Activate/Retrieve the Licenses . . . . .	44
Install Content and Panorama Software Updates . . . . .	46
Navigate the Panorama User Interface . . . . .	48
Navigate the Web Interface . . . . .	48
Log in to the Web Interface . . . . .	49
Log in to the CLI . . . . .	50
Set Up Administrative Access . . . . .	51
Create an Administrative Account . . . . .	51
Define Access Domains . . . . .	53
Create an Authentication Profile . . . . .	54
Define an Authentication Sequence . . . . .	55
Configure Administrative Authentication . . . . .	55
<b>Manage Firewalls and Log Collection . . . . .</b>	<b>61</b>
Manage Your Firewalls . . . . .	62
Add Managed Devices . . . . .	62
Create Device Groups . . . . .	64
Create Templates . . . . .	70
Configure the Firewalls to Forward Logs to Panorama . . . . .	74
Commit Changes on Panorama . . . . .	79
Modify the Log Forwarding and Buffering Defaults . . . . .	80
Use Panorama to Configure Managed Devices: An Example . . . . .	81
Enable Logging . . . . .	90
Deploy Software Updates and Manage Licenses . . . . .	98
Replace a Managed Device with a New Device . . . . .	102
Before you Begin . . . . .	102
Restore the Configuration on the New Device . . . . .	104
Transition a Device to Central Management . . . . .	107
<b>Monitor Network Activity . . . . .</b>	<b>109</b>
Use Panorama for Visibility . . . . .	110
Monitor the Network with the ACC and AppScope . . . . .	110
Analyze Log Data . . . . .	113
Generate Reports . . . . .	113
Use Case: Monitor Applications Using Panorama . . . . .	116
Use Case: Use Panorama to Respond to an Incident . . . . .	120
<b>Panorama High Availability . . . . .</b>	<b>125</b>
High Availability Overview . . . . .	126
Failover Triggers . . . . .	127
Logging Considerations in HA . . . . .	128
Priority and Failover . . . . .	129
What Settings are Not Synchronized Between the HA Peers? . . . . .	130

Configure a Panorama High Availability Pair .....	131
Set Up High Availability on Panorama .....	131
Verify Failover. ....	134
Switch Priority to Resume NFS Logging. ....	134
Upgrade Panorama in High Availability .....	136
<b>Administer Panorama .....</b>	<b>139</b>
Manage Configuration Backups .....	140
Schedule Export of Configuration Files .....	141
Manage Panorama Configuration Backups .....	142
Configure the Number of Backups Stored on Panorama. ....	142
Load a Configuration Backup on a Managed Device .....	143
Compare Changes in Configuration .....	144
Restrict Access to Configuration Changes .....	145
Types of Locks .....	145
Locations for Taking a Lock .....	145
Take a Lock. ....	146
View Current Lock Holders .....	146
Enable Automatic Acquisition of the Commit Lock .....	146
Remove a Lock .....	147
Add Custom Logos .....	148
View Task Completion History .....	149
Reallocate Log Storage Quota .....	150
Monitor Panorama .....	152
Set up Email Alerts .....	153
Set up SNMP Access .....	154
Reboot or Shutdown Panorama. ....	158
Generate Diagnostic Files .....	159
Configure Password Profiles and Password Complexity .....	160
Replace the Virtual Disk on a Panorama Virtual Appliance. ....	162
<b>Troubleshooting .....</b>	<b>165</b>
Why does the Template commit fail? .....	165
Why is Panorama running a File System Integrity check? .....	166
Is there a separate connection for forwarding logs to Panorama? .....	166
Why does the log storage capacity for the Collector Group read 0 MB? .....	166
Why is Panorama in a suspended state? .....	167
Where do I view task completion status? .....	167





# 1 Panorama Overview

---


Panorama provides centralized management and visibility of multiple Palo Alto Networks next-generation firewalls. It allows you to oversee all applications, users, and content traversing the network from one location, and then use this knowledge to create application enablement policies that protect and control the entire network. Using Panorama for centralized policy and device management increases operational efficiency in managing and maintaining a distributed network of firewalls.

The following sections describe Panorama and provide guidelines for planning your Panorama deployment:

- ▲ [About Panorama](#)
- ▲ [Panorama Platforms](#)
- ▲ [About Centralized Configuration and Deployment Management](#)
- ▲ [About Centralized Logging and Reporting](#)
- ▲ [About Role-Based Access Control](#)
- ▲ [Panorama Recommended Deployments](#)
- ▲ [Plan Your Deployment](#)
- ▲ [Deploy Panorama: Task Overview Checklist](#)

## About Panorama

Panorama provides centralized management of the Palo Alto Networks next-generation firewalls, as shown in the following illustration:



Panorama allows you to effectively configure, manage, and monitor your Palo Alto Networks firewalls using central oversight with local control, as required. The three focal areas in which Panorama adds value are:

- ▲ Centralized configuration and deployment—To simplify central management and rapid deployment of the firewalls on your network, use Panorama to pre-stage the firewalls for deployment. You can then assemble the devices into groups, and create templates to apply a base network and device configuration and use device groups to administer globally shared and local policies. See [About Centralized Configuration and Deployment 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 [About Centralized Logging and Reporting](#).
- ▲ Distributed administration—Allows you to delegate or restrict access to global and local device configurations and policies. See [About Role-Based Access Control](#) for delegating appropriate levels of access for distributed administration.

Panorama is available in two platforms: as a virtual appliance and as a dedicated hardware appliance. For more information, see [Panorama Platforms](#).



## Panorama Platforms

Panorama is available in two platforms, each of which supports device management licenses for managing up to 25 devices, up to 100 devices, or up to 1000 devices:

- **Panorama Virtual Appliance**—The Panorama virtual appliance is installed on a VMware server. It allows for a simple installation and facilitates server consolidation for sites that need a virtual management appliance. It also supports integration with a Network File System (NFS) for increased storage and (> 2TB) log retention capabilities.

The Panorama virtual appliance is best suited for environments with fewer than 10 firewalls and log rates less than 10,000 logs/second.

- **M-100 Appliance**—A dedicated hardware appliance intended for large scale deployments. In environments with high logging rates and log retention requirements, this platform enables scaling of your log collection infrastructure. The appliance supports RAID 1 mirroring to protect against disk failures, and the default configuration ships with two 1TB drives; with additional RAID pairs, the M-100 appliance can support up to 4TB of log storage.

The M-100 appliance allows for separation of the central management function from the log collection function by supporting the following deployment modes:

- **Panorama mode:** The appliance performs both the central management and the log collection functions. This is the default mode.
- **Log Collector mode:** The appliance functions as a dedicated log collector, which can be managed by either an M-100 appliance in Panorama mode or a Panorama virtual appliance.

When deployed in Log Collector mode, the appliance does not have a web interface; administrative access is CLI only.

The platform choice is hinged on your need for a virtual appliance, the number of Palo Alto Networks firewalls you plan to manage, and your log collection requirements as detailed in the following table:

Considerations	VMware	M-100	
	Panorama Virtual Appliance	Panorama Mode	Distributed Log Architecture with Dedicated Log Collectors
Number of managed devices	10 or fewer firewalls	Up to 100 firewalls	up to 1000 firewalls
Log collection rate	<10,000 logs/second	<10,000 logs/second	> 10,000 logs/sec (Max 50,000 logs/sec per collector)

## About Centralized Configuration and Deployment Management

Panorama uses *Device Groups* and *Templates* to group devices into smaller and more logical sets that require similar configuration. All configuration elements, policies, and objects on the managed firewalls can be centrally managed on Panorama using Device Groups and Templates. In addition to managing configuration and policies, Panorama enables you to centrally manage licenses, software and associated content updates: SSL-VPN clients, GlobalProtect agents, dynamic content updates (Applications, Threats, WildFire and Antivirus).

### Context Switch—Device or Panorama

The Panorama web interface allows you to toggle between a Panorama-centric view to a device-centric view using the *context switch*. You can choose to manage the device centrally using Panorama and then switch context to a specific managed device to configure the device using the device's user interface. The similarity of the user interface on the managed firewalls and Panorama allows you to seamlessly move between the interfaces to administer and monitor devices as required.

If you have configured [access domains](#) to restrict administrative access to specific managed devices, the Panorama user interface only displays the devices/features for which the logged-in administrator has permissions.

### Templates

Templates are used to configure the settings that are required for the managed firewalls to operate on the network. They allow you to define a common base configuration using the **Network** and **Device** tabs on Panorama. For example, interface and zone configuration, server profiles for logging and SNMP access, network profiles for controlling access to zones and IKE gateways, can all be managed using Templates. When you group devices to define Template settings, consider grouping devices that are alike in hardware model, and require access to similar network resources, such as gateways and syslog servers.

Using templates, you can either push a limited common base configuration to a group of devices and then configure the rest of the settings manually on the device. Or, push a larger common base configuration and then override the template settings on the device to adapt for device-specific changes. When you override a setting on the device, the setting is saved to the local configuration of the device and is no longer managed by the Panorama template. You can, however, use Panorama to force the template configuration onto the device or restore the template settings on the device. For example, you can define a common NTP server in the template, but override the NTP server configuration on the device to accommodate for the local time zone on the device. If you then decide to restore the template settings, you can easily undo or revert the local changes that you implemented on the device.

Templates cannot be used to define an operational state change such as FIPS mode or to enable multi-vsyt mode on the firewalls. For more information, see [What can Templates not be Used for?](#)

## Device Groups

To use Panorama effectively, you must group the firewalls on your network into logical units called *device groups*. A device group allows grouping based on network segmentation, geographic location, or by the need to implement similar policy configurations. A device group can include physical firewalls, virtual firewalls and/or a virtual system. By default, all managed devices belong to the *Shared* device group on Panorama.

Device Groups enable central management of policies and objects using the **Policies** and **Objects** tabs on Panorama. Objects are configuration elements that are referenced in policies. Some of the objects that firewall policies make use of are: IP addresses, URL categories, security profiles, users, services, and applications.

Using Device Groups you can create shared objects or device group-specific objects and then use these objects to create a hierarchy of rules (and rulebases) to enforce how inbound and outbound traffic is handled by the managed firewalls. For example, a corporate acceptable use policy could be defined as a set of shared policies. Then, to allow only the regional offices to access peer-to-peer traffic such as bittorrent, you can create a security rule as a shared policy and target it to the regional offices or make it a device group rule that is pushed to the regional offices. See [Use Panorama to Configure Managed Devices: An Example](#).

## About Policies

Device Groups provide a way to implement a layered approach for managing policies across a network of managed firewalls. The layered approach allows for deployment of corporate policies centrally, as *shared policies*, in conjunction with *device group-specific policies* and policies that are *locally* administered on the device.

Both shared policies and device group-specific policies allow you to craft pre-rules and post-rules to manage all the rulebases from a central location: Security, NAT, QoS, Policy Based Forwarding, Decryption, Application Override, Captive Portal, and DoS Protection.

- **Pre-rules**—Rules that are added to the top of the rule order and are evaluated first. You can use pre-rules to enforce the Acceptable Use Policy for an organization; for example, to block access to specific URL categories, or to allow DNS traffic for all users. Pre-rules can be of two types: Shared pre-rules that are shared across all managed devices and Device Groups, and Device Group pre-rules that are specific to a Device Group.
- **Post-rules**—Rules that are added at the bottom of the rule order and are evaluated after the pre-rules and the rules locally defined on the device. Post-rules typically include rules to deny access to traffic based on the App-ID, User-ID, or Service. Like pre-rules, post rules are also of two types: Shared post-rules that are shared across all managed devices and Device Groups, and Device Group post-rules that are specific to a Device Group.

The evaluation order of the rules is:



When the traffic matches a policy rule, the defined action is triggered and all subsequent policies are disregarded.

This ability to layer policies, creates a hierarchy of rules where local policies are placed between the pre- and post-rules, and can be edited by switching to the local firewall context, or by accessing the device locally. This cascade of rules is visually demarcated for each device group (and managed device), and provides the ability to scan through a large numbers of rules.

	Dashboard	ACC	Monitor	Polices	Objects	Network	Device				
Pre-rules	<div> <div></div> <div></div> <div></div> <div></div> <div></div> <div></div> <div></div> <div></div> <div></div> <div></div> <div></div> </div>										
	Name	Tag	Zone	Address	User	HIP Profile	Zone	Address	Application	Service	Action
	mjs-dg-rule	none	trust	capt-a	any	any	untrust	corp-nets	any	any	
	rdp	mjs-dg	trust	any	domain/user...	any	untrust	any	ms-rdp	application-...	
	rdp-2	mjs-dg	trust	A	any	any	untrust	any	ms-rdp	application-...	
		mikes-test-rule		B							
				C							
				D							
	allow-all-to-untr...	none	trust	any	any	any	untrust	any	any	any	
	allow-m-100	alert	untrust	any	any	any	trust	any	ssh	m-100-https	
Local rules on the device								ssl	m-100-ssh		
	laptop-rama	none	trust	mike...	any	any	trust	any	any	any	
	test	none	trust	C	any	any	untrust	mjs-mac...	any	any	
	web	none	trust	any	any	any	untrust	any	web-browsing	application-...	
	allow-and-log	none	trust	any	any	any	untrust	any	any	any	
Post-rules	deny-and-log	none	untrust	any	any	any	trust	any	any		

Pre-rules and post-rules pushed from Panorama can be viewed on the managed firewalls, but they can only be edited in Panorama. Local device rules can be edited by either the local administrator or a Panorama administrator who has switched to a local firewall context.

## About Objects

Objects are configuration elements that are referenced in policies. Some of the objects that firewall policies make use of are: IP addresses, URL categories, security profiles, users, services, and applications. Because objects can be reused across policies, creating *shared objects* or *device group objects* reduces duplication of these configuration elements. For example, creating shared address objects and address groups or shared service objects and service groups allows you to create one instance of the object and reference it in any rulebase to manage the firewalls across multiple device groups. Because shared objects are defined once but used many times, they reduce administrative overhead, and maintain consistency and accuracy everywhere the shared object is used.

Both shared objects and device group objects can be used in pre-rules, post-rules and in rules locally defined on a device. When you create an object on Panorama, you can configure the behavior on whether:

- The device group object takes precedence over a shared object, when both objects have the same name. By default, the shared object takes precedence. This behavior ensures that an shared object always supersedes a device group object with the same name.

However, if a device has a locally created object with the same name as a shared or a device group object that is pushed from Panorama, a commit failure will occur.

- All shared and device group objects that are defined on Panorama are pushed to the managed devices. By default, all objects—those that are and are not referenced in policies—are pushed to the managed devices.

## About Centralized Logging and Reporting

Panorama allows aggregates data from all managed firewalls, providing 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 devices.

The Application Command Center (ACC) on Panorama provides a single pane for unified reporting across all the firewalls; it allows you to centrally analyze, investigate, and report on network traffic and security incidents. On Panorama, you can view logs and generate reports from logs forwarded to Panorama or to the managed Log Collectors, if configured, or you can query the managed devices 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 Log Collectors) or by accessing the logs stored locally on the managed devices.

If you choose not to configure the managed firewalls to forward logs to Panorama, you can schedule reports to be run on each managed firewall and forward the results to Panorama for a combined view of user activity and network traffic. Although this view does not provide granular drill-down on specific data and activities, it still provides a unified reporting approach.

### Logging Options

Both the Panorama virtual appliance and the M-100 appliance can perform log collection for logs that are forwarded from the managed devices. The options for logging vary on each platform.

- **On a Panorama virtual appliance**, there are three logging options: use the 10GB of internal storage space allocated for logging as soon as you install the virtual appliance, or add a virtual disk that can support up to 2TB of storage, or mount a Network File System (NFS) datastore, where you can determine the storage capacity that is allocated for logging.
- **On the M-100 appliance**, the default shipping configuration includes 1TB disks in a RAID pair, which can be increased to 4TB RAID storage. When the M-100 appliance is in Panorama mode, you can enable the RAID disks and use these disks as the default Log Collector. With the M-100 appliance in Log Collector mode, you must use Panorama to assign the devices to the Log Collector appliance(s). In a deployment with multiple Log Collector appliances, Panorama queries all managed Log Collectors to generate an aggregated view of traffic and cohesive reports.

For easy scaling, begin with a single Panorama and incrementally add dedicated Log Collectors, as your needs expand.

### Managed Collectors and Collector Groups

A Log Collector is an M-100 appliance that is configured to function in *Log Collector mode*. It is a dedicated Log Collector appliance that is configured and managed using Panorama, and hence also called a Managed Collector. It can be managed by either an M-100 appliance in Panorama mode or by a Panorama virtual appliance. When added as a Managed Collector and connected to Panorama, the Log Collector can be administered using the Panorama web interface. Otherwise, administrative access to the Log Collector is only available through the CLI using the default administrative user (*admin*) account. Additional administrative user accounts are not supported.

A Collector Group is one or more M-100 appliances that operate as a single logical log collection unit, and the logs are uniformly distributed amongst all the disks in a Log Collector and across all members in the Collector Group. Spreading the logs uniformly across the disks and Log Collectors, maximizes the use of the available storage space. Each Panorama can manage up to 16 Collector Groups. In order to manage a Log Collector, you must add it to a Collector Group. Although a Collector Group can contain multiple Log Collectors, Palo Alto Networks recommends placing only one Log Collector in a Collector Group unless [more than 4TB of storage space is required](#) in a Collector Group.

The Collector Group configuration specifies which managed firewalls can send logs to the Log Collectors in the group. After the Log Collectors are configured and the firewalls are enabled for forwarding logs, each device forwards its logs to the assigned Log Collector.



If you are using Panorama to manage firewalls running both PAN-OS version 5.0 and a PAN-OS version earlier than 5.0, make note of the following compatibility requirement:

- Only devices running PAN-OS v5.0 can send logs to a dedicated Log Collector (an M-100 appliance configured in the Log Collector mode).
- Devices running PAN-OS versions earlier than 5.0a can only send logs to a Panorama virtual appliance or to an M-100 appliance in Panorama mode.

Managed Collectors and Collector Groups are integral to the Distributed Log Collection architecture on Panorama. The Distributed Log Collection architecture allows for easy scalability and incremental addition of dedicated Log Collectors as your logging needs grow. The M-100 appliance in Panorama mode can log to its default Collector Group and then be expanded to a Distributed Log Collection architecture with one or more Collector Groups that include M-100 appliances in the Log Collector mode.

## Using Multiple Log Collectors in a Collector Group

Although Palo Alto Networks recommends placing only one Log Collector in a Collector Group, if you have a scenario where you require more than 4TB of log storage capacity in a Collector Group, you may need to add multiple Log Collectors to the Collector Group. In the following scenarios, a Collector Group may require multiple Log Collectors:

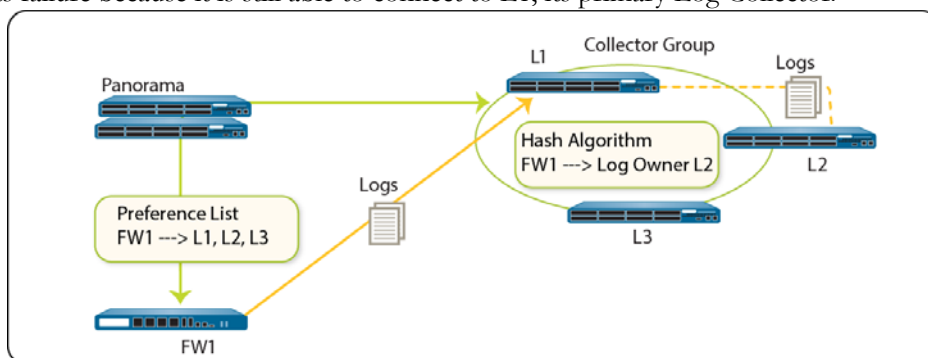
- A single firewall generates more than 4TB of logs. For example, if a managed firewall generates 12 TB of logs, you will require at least three Log Collectors in the Collector Group.
- A groups of firewalls that are forwarding logs to Collector Group and the capacity requirement exceeds 4TB of storage space.

If a Collector Group contains multiple Log Collectors, the available storage space is used as one logical unit and the logs are uniformly distributed across all the Log Collectors in the Collector Group. The log distribution is based on the disk capacity of the Log Collectors (that ranges from 1TB to 4TB, depending on the number of disk pairs) 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, the logs may not necessarily be written to the first Log Collector specified in the preference list.

For example, consider the following preference list:

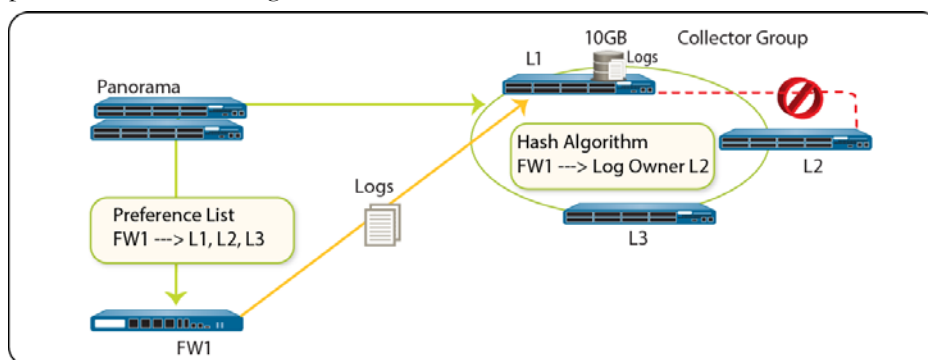
Managed Firewall	Log Forwarding Preference List Defined on a Collector Group
FW1	L1,L2,L3
FW2	L4,L5,L6

Using this list, FW1 will forward logs to L1, its primary Log Collector, but the hash algorithm could determine that the logs will be written on L2. If L2 becomes inaccessible or has a chassis failure, FW1 will not know about its failure because it is still able to connect to L1, its primary Log Collector.



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

With multiple Log Collectors in a Collector Group, the firewall does not buffer logs to its local storage when it is able to connect to its Primary Log Collector. Therefore, FW1 will continue sending logs to L1. Because L2 is unavailable, the Primary Log Collector L1 buffers the logs to its HDD, which has 10GB of log space. If L2 remains unavailable and the logs pending for L2 exceed 10GB, L1 will overwrite the older log entries in order to continue logging. In such an event, loss of logs is a risk. Therefore, if using multiple Log Collectors in a Collector Group, make sure to obtain an On-Site-Spare (OSS) or a cold standby unit to enable prompt replacement should a Log Collector failure occur.





## Centralized Reporting

Panorama aggregates logs from all managed devices and enables reporting on the aggregated data for a global view of application use, user activity, and traffic patterns across the entire network infrastructure. 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 devices 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 architecture, 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 devices 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 device and run reports on data that is stored locally on the managed devices.

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.

## About Role-Based Access Control

Role-based access control allows you to specify the privileges and responsibilities accorded to every administrative user. On Panorama, you can define administrative accounts with specific roles, profiles, or [access domains](#) to regulate access to specific features on Panorama and the managed devices; these options allow you to limit administrative access to only the devices and areas of the management interface that each administrator requires to perform the job. By default, every Panorama server comes pre-configured with a default administrative account (admin) that provides full read-write access (also known as superuser access). As a best practice, 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 (or modification) and enables logging of the actions of each individual administrator.

For every administrative user, you can also define an authentication profile that determines how the user's access credentials are verified. To enforce more granular administrative access, use access domains in order to restrict administrative access to a particular device, device group or template.

### Administrative Roles

The way you configure administrator accounts depends on the security requirements within your organization, whether you have existing authentication services you want to integrate with, and how many different administrative roles are required. A *role* defines the type of access the associated administrator has to the system. There are two types of roles:

- **Dynamic Roles**—Built-in roles that provide access to Panorama and the managed devices: Superuser (full-access), Superuser (read-only), and Panorama administrator.

The Panorama administrator cannot perform the following actions:

- Create, modify, or delete Administrators
- Create, modify, or delete Admin Roles or Access Domains
- Export, validate, revert, save, load, or import the configuration from the **Device > Setup** tab
- Configure **Scheduled Config Export** functionality on the **Panorama** tab.

With dynamic roles, there is no need to update the role definitions as new features are added because the roles automatically update.

- **Admin Role Profiles**—Create your own role definitions in order to provide more granular access control to the various functional areas of the web interface, CLI and/or XML API. The two Admin Role Profiles available are: Panorama, and Device Group and Template. You could create an Admin Role Profile for your operations staff that provides access to specific Device Groups and/or Templates so that they have access to the device and network configuration areas of the web interface and a separate profile for your security administrators that provides access to security policy definition, logs, and reports on Panorama. Keep in mind that with an Admin Role Profile you must update the profiles to explicitly assign privileges for new features/components that are added to the product. By default, access to all new components and features is disabled.

See [Set Up Administrative Access](#) for creating administrative roles.

## Authentication Profiles and Sequences

Among its other uses, an authentication profile defines how an administrative user is authenticated on Panorama upon login. If you create a local user account on Panorama, you can authenticate the user to the local database, or use an external RADIUS, LDAP, or Kerberos server for authentication. If you do not want to create a local user account, and want to manage both account administration and authentication using an external authentication mechanism, you must use RADIUS. For a high-level overview of the process, see [Using RADIUS Vendor Specific Attributes \(VSAs\)](#).

To authenticate to multiple authentication sources—local, RADIUS, LDAP, and/or Kerberos—define an authentication sequence. An authentication sequence is a ranked order of authentication profiles that an administrative user is matched against. Panorama checks against the local database first, and then each profile in sequence until the user is successfully authenticated. The user is denied access to Panorama only if authentication fails for all the profiles defined in the authentication sequence.

To create authentication profiles and sequences, see [Create an Authentication Profile](#) and [Define an Authentication Sequence](#).

## Access Domains

An access domain defines the features and permissions accorded to an administrative user, enabling granular control over the administrative user's ability to switch context and access the features on the user interface of the managed firewalls. Access Domains can also limit access to a subset of the Device Groups and/or Templates created on Panorama and therefore restrict the user's ability to configure and manage devices.

The access domain is linked to RADIUS vendor-specific attributes (VSAs) and is supported only if a RADIUS server is used for administrator authentication. If RADIUS is not used, the access domain settings are ignored. For information on defining an access domain, see [Define Access Domains](#).

## Administrative Authentication

There are four ways to authenticate administrative users:

- **Local administrator account with local authentication**—Both the administrator account credentials and the authentication mechanisms are local to the firewall. To further secure the local administrator account, create a password profile that defines a validity period for passwords and/or set device-wide password complexity settings. For more information, see [Create an Administrative Account](#).
- **Local administrator account with certificate- or key-based authentication**—With this option, the administrator accounts are local to the firewall, but authentication is based on SSH keys (for CLI access) or client certificates/common access cards (for the web interface). For details on how to configure this type of administrative access, see [Enable Certificate-Based Authentication for the Web Interface](#) and [Enable SSH Key-Based Authentication for the Command Line Interface](#).

- **Local administrator account with external authentication**—The administrator accounts are managed on the local firewall, but the authentication functions are offloaded to an existing LDAP, Kerberos, or RADIUS service. To configure this type of account, you must first create an authentication profile that defines how to access the external authentication service and then create an account for each administrator that references the profile. For more information, refer to “Setting Up Authentication Profiles” in Chapter 3 of the [Palo Alto Networks Administrator's Guide](#).
- ▲ **External administrator account and authentication**—Account administration and authentication are handled by an external RADIUS server. To use this option, you must define Vendor Specific Attributes (VSAs) on your RADIUS server that map to the admin role. For a high-level overview of the process, see [Using RADIUS Vendor Specific Attributes \(VSAs\)](#). For details on how to configure this type of administrative access, refer to the [Radius Vendor Specific Attributes \(VSA\)](#) article.

## Panorama Recommended Deployments

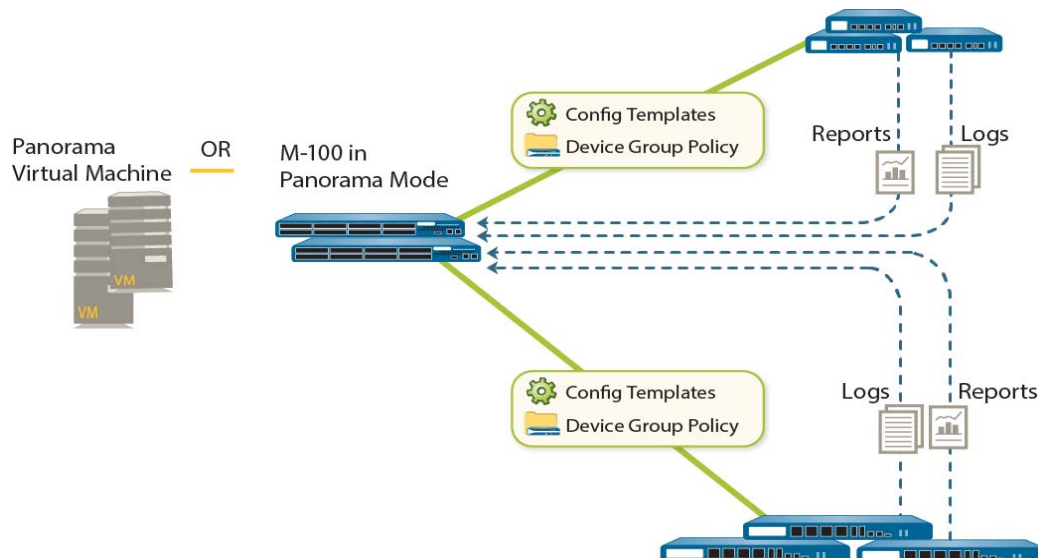
A Panorama deployment is comprised of the Panorama Management server that has a browser-based interface, (optional) Log Collectors, and the Palo Alto Networks firewalls to be managed. The recommended Panorama deployments are:

- ▲ Panorama for Centralized Management and Reporting
- ▲ Panorama in a Distributed Log Collection Architecture

### Panorama for Centralized Management and Reporting

The following diagram shows how the Panorama virtual appliance or the M-100 appliance can be deployed in a redundant configuration to provide the following benefits:

- Centralized management—Centralized policy and device management that allows for rapid deployment and management of up to one thousand firewalls.
- Visibility—Centralized logging and reporting to analyze and report on user-generated traffic and potential threats.
- Role-based access control—Appropriate levels of administrative control at the device level or global level for administration and management.



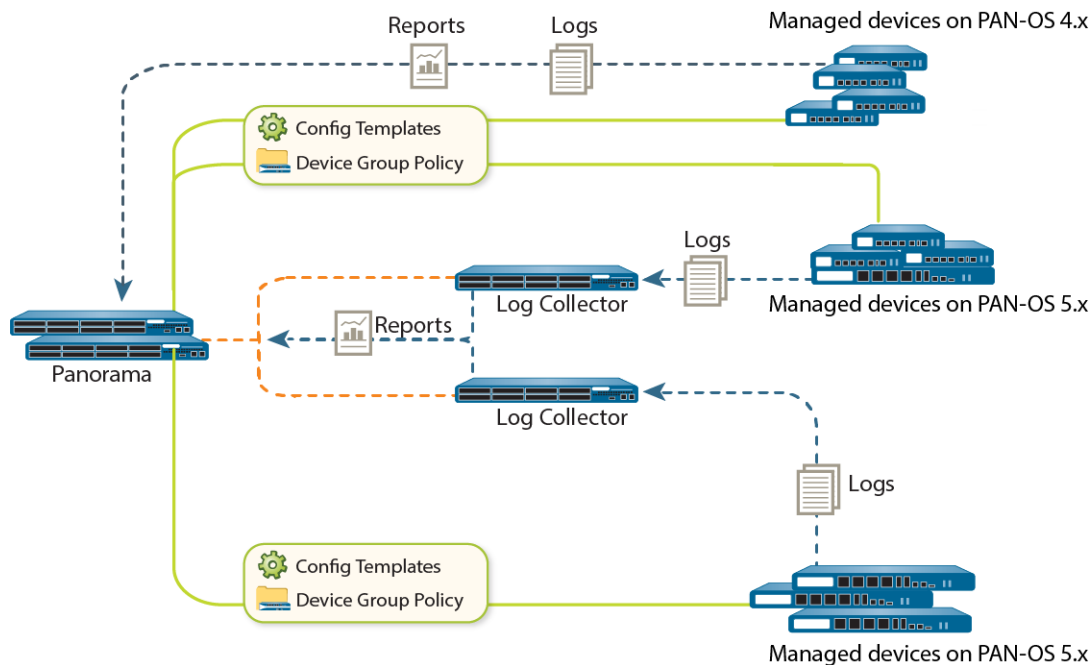
## Panorama in a Distributed Log Collection Architecture

The hardware-based Panorama—the M-100 appliance—can be deployed either as a Panorama Management server that performs management and log collection functions or as a dedicated Log Collector that provides a comprehensive log collection solution for the firewalls on your network. Using the M-100 appliance as a Log Collector allows for a more robust environment where the log collection process is offloaded to a dedicated appliance. Using a dedicated appliance in a Distributed Log Collection (DLC) architecture provides redundancy, improved scalability, and capacity for longer term log storage.

In a DLC architecture, the Panorama Management server (Panorama virtual appliance or an M-100 in Panorama mode) manages the firewalls and the Log Collectors. Using Panorama, the firewalls are configured to send logs to one or more Log Collectors; Panorama can then be used to query the Log Collectors and provide an aggregated view of network traffic. In a DLC configuration, the logs stored on the Log Collectors are accessible from both the primary and secondary Panorama peers in a high availability (HA) pair.

In the following topology, the Panorama peers in high availability mode manage the deployment and configuration of firewalls running PAN-OS 4.x and 5.x. This solution provides the following benefits:

- Allows for improved performance in the management functions on Panorama
- Provides high-volume log storage on a dedicated hardware appliance
- Provides horizontal scalability and redundancy with RAID 1 storage



## Plan Your Deployment

- Verify the PAN-OS versions of the firewalls to be managed. To manage the firewalls, Panorama must be running the same major release or a later version than the firewalls that it will manage. For example, a Panorama 4.0 appliance cannot manage devices running PAN-OS 5.0.
- Plan to use the same URL filtering database (BrightCloud or PAN-DB) across all managed firewalls. If some firewalls are using the BrightCloud database and others are using PAN-DB, Panorama can only manage security policies for one or the other URL filtering database. URL filtering rules for the other database must be managed locally on the firewalls that use that database.
- Plan to use Panorama in a high availability configuration; set it up as an active/passive high availability pair. See [Panorama High Availability](#).
- Estimate the log storage capacity for your network. To size the log storage capacity for your security and compliance needs, you must consider a number of factors such as the network topology, number of firewalls sending logs, type of log traffic for example, URL and threat logs versus traffic logs, the [rate](#) at which logs are generated and the number of days that you want to store logs on Panorama. For details, see the article: [Panorama Logging Suggestions](#).
- For meaningful reports on network activity, plan a logging solution:
  - Do you need to forward logs to a syslog server, in addition to Panorama?
  - 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 need to forward logs?
  - Do you need redundancy in logging? With Panorama virtual appliances in HA, each peer can log to its virtual disk. The managed devices can send logs to both peers in the HA pair. This option provides redundancy in logging and is best suited to support up to 2TB of log storage capacity.
  - Do you log to an NFS? Support for NFS is only provided on the Panorama virtual appliance. Consider using NFS if more than 2TB of log storage capacity is required. If using NFS, note that the managed devices can only send logs to the primary peer in the HA pair, and only the active-primary Panorama is mounted to the NFS and can write to it.
- 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 devices in the network.

If you have already deployed and configured the Palo Alto Networks firewalls on your network, determine whether to transition the devices to centralized management. This process requires a migration of all configuration and policies from your firewalls to Panorama, see [Transition a Device to Central Management](#).
- Determine what administrative access privileges, roles, and permissions are required for permitting access to the managed firewalls and Panorama. See [Set Up Administrative Access](#).
- Plan the required Device Groups. To do this, determine whether to group devices based on device function, security policy, geographic location, or network segmentation. For example, group devices by function— such as those that support the organizational needs of partners or the R&D functional groups; or group devices that perform the same function, such as gateway devices, branch office devices or datacenter devices. See [Device Groups](#).

- Plan a layering strategy for administering policies. Think through how policies must be inherited and evaluated and how to best implement shared rules, device-group rules, and device-specific rules to meet your network needs.
  - For visibility and centralized policy management, consider using Panorama for administering policies, even if you would like to create device-specific exceptions to shared/device-group policies. To apply a rule to a subset of devices in a device group, you can *target* the rule(s) to the specific device(s), see [Target Policies to a Subset of Devices](#).
  - Consider whether to create smaller device groups based on commonality or to create larger device groups to scale more easily. See [Use Panorama to Configure Managed Devices: An Example](#).
- Plan your device organization for how configuration settings (using templates) are inherited and enforced. For example, think through how to assign devices to templates based on hardware platforms, geographic proximity and similar network set up needs for time zones, DNS server, and interface settings. See [Use Panorama to Configure Managed Devices: An Example](#).



## Deploy Panorama: Task Overview Checklist

The following task list summarizes the steps to get started with Panorama:

- Step 1.** (M-100 appliance only) Rack mount the appliance. Refer to the [M-100 Hardware Reference Guide](#).
- Step 2** Perform initial configuration to enable network access to Panorama. See [Set Up the Panorama Virtual Appliance](#) or [Set Up the M-100 Appliance](#).
- Step 3** [Install Licenses](#).
- Step 4** [Install Content and Panorama Software Updates](#).
- Step 5** [Add Managed Devices](#)
- Step 6** [Create Device Groups](#) and [Create Templates](#).
- Step 7** (Optional) Enable log collection to a dedicated Log Collector. See [Enable Logging](#).
- Step 8** Monitor network activity using the visibility and reporting tools on Panorama. See [Monitor the Network with the ACC and AppScope](#) and [Generate Reports](#).
- Step 9** Set up Panorama in a high availability configuration. See [Panorama High Availability](#).

For a use case example, for how to begin using Panorama for central management, see the workflow in [Use Panorama to Configure Managed Devices: An Example](#).





## 2 Set Up Panorama

---

For centralized reporting and cohesive policy management across all the firewalls on your network, Panorama can be deployed as a virtual appliance or as a hardware appliance (the M-100 appliance).

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

- ▲ [Set Up the Panorama Virtual Appliance](#)
- ▲ [Set Up the M-100 Appliance](#)
- ▲ [Migrate from a Panorama Virtual Appliance to an M-100 Appliance](#)
- ▲ [Navigate the Panorama User Interface](#)
- ▲ [Set Up Administrative Access](#)

## Set Up the Panorama Virtual Appliance

The Panorama virtual appliance consolidates the Panorama management and logging functions into a single virtual appliance. This solution enables use of an existing VMware virtual infrastructure to easily deploy and centrally administer and monitor the Palo Alto Networks firewalls in your network as described in the following sections:

- ▲ Prerequisites
- ▲ Install Panorama on the ESX(i) Server
- ▲ Perform Initial Configuration
- ▲ Expand Log Storage Capacity on the Panorama Virtual Appliance



The Panorama virtual appliance cannot be used as a dedicated Log Collector. Only an M-100 appliance in Log Collector mode provides dedicated log collection capabilities. You can, however, manage a Log Collector using the Panorama virtual appliance. See [Set Up the M-100 Appliance](#) for details.

### Prerequisites

To set up a Panorama virtual appliance efficiently, check that your server meets the following requirements before you begin:

- **Minimum System Requirements**

For Panorama version 5.1	For Panorama version 5.0 or earlier
Panorama version 5.1 is a 64-bit kernel-based VM	Panorama version 5.0 or earlier use a 32-bit kernel-based VM
<ul style="list-style-type: none"> <li>• VMware ESX(i) 4.1 or later</li> <li>• Quad Core CPU (2GHz); use 3GHz if you have 10 or more firewalls</li> <li>• 4GB RAM; 16GB recommended if have 10 or more firewalls</li> <li>• 40GB disk space</li> <li>• A client computer with one of the following: VMware vSphere Client or VMware Infrastructure Client that is compatible with your ESX(i) server</li> </ul>	<ul style="list-style-type: none"> <li>• VMware ESX(i) 3.5 or later</li> <li>• 2GHz CPU; use Quad Core CPU for optimal performance with high logging rates</li> <li>• 2GB RAM; 4GB recommended if have 10 or more firewalls</li> <li>• 40GB disk space</li> <li>• A client computer with one of the following: VMware vSphere Client or VMware Infrastructure Client that is compatible with your ESX(i) server</li> </ul>



VMware concepts and terminology are not covered in this document. This guide assumes familiarity with the VMware suite of products that are required to create the virtual appliance.

- Register the Panorama serial number on the support site at <https://support.paloaltonetworks.com>; the serial number was sent to you by email. After registering the serial number on the support site, you gain access to the Panorama software downloads page.

## Install Panorama on the ESX(i) Server

Use these instructions to install a new Panorama virtual appliance. If you are upgrading your existing Panorama virtual appliance, skip to [Install Content and Panorama Software Updates](#).

CREATE THE VIRTUAL PANORAMA	
<p><b>Step 1.</b> Download and extract the Panorama base image zip file to the server on which you will be installing Panorama.</p> <p>The virtual appliance installation uses the Open Virtual Machine Format (OVF) template file, which is included in the base image.</p>	<ol style="list-style-type: none"> <li>1. Go to <a href="https://support.paloaltonetworks.com/">https://support.paloaltonetworks.com/</a> and download the <b>Panorama Base Image</b> zip file.</li> <li>2. Unzip the Panorama base image zip file, and extract the <b>panorama-esx.ovf</b> file. This .ovf template file is required for installing Panorama.</li> </ol>
<p><b>Step 2</b> Access the ESX(i) server.</p>	<p>Launch the VMware vSphere Client and connect to the VMware server.</p>
<p><b>Step 3</b> Install Panorama.</p> <p>Starting with Panorama 5.1, the Panorama virtual appliance is installed as a 64-bit virtual machine.</p>	<ol style="list-style-type: none"> <li>1. Choose <b>File &gt; Deploy OVF Template</b>.</li> <li>2. <b>Browse</b> to select the panorama-esx.ovf file from the recently unzipped Panorama base image, and click <b>Next</b>.</li> <li>3. Confirm that the product name and description match the downloaded version, and click <b>Next</b>.</li> <li>4. Enter a descriptive name for the Panorama virtual appliance, and click <b>Next</b>.</li> <li>5. Select a <b>Datastore Location</b> on which to install the Panorama image, and click <b>Next</b>.</li> <li>6. Select <b>Thick Provision Lazy Zeroed</b> as the disk format, and click <b>Next</b>.</li> <li>7. Specify which networks in the inventory must be used for the Panorama virtual appliance.</li> <li>8. Confirm the selected options and then click <b>Finish</b> to begin the installation process.</li> </ol>

CREATE THE VIRTUAL PANORAMA	
	<p>9. When the installation completes, select the Panorama virtual appliance, and click <b>Edit Settings...</b> to define the following settings:</p> <ol style="list-style-type: none"> <li>Verify that you have allocated the appropriate amount of memory. <ul style="list-style-type: none"> <li>Panorama 5.1: at least 4GB</li> <li>Panorama 5.0 or earlier: 2 to 4GB</li> </ul> </li> <li>Select the appropriate guest operating system. <ul style="list-style-type: none"> <li>Panorama 5.1: <b>Linux</b> as your <b>Guest Operating System</b> and <b>Version</b> as <b>Other Linux (64-bit)</b>.</li> <li>Panorama 5.0 or earlier: <b>Linux</b> as your <b>Guest Operating System</b> and <b>Version</b> as <b>Other Linux (32-bit)</b></li> </ul> </li> <li>Choose the SCSI controller. <ul style="list-style-type: none"> <li>Panorama 5.1: <b>SCSI controller</b> is <b>LSI Logic Parallel</b></li> <li>Panorama 5.0 or earlier: <b>SCSI controller</b> is <b>Bus Logic Parallel</b></li> </ul> </li> </ol>
<p><b>Step 4</b> Power on the Panorama virtual appliance.</p>	<p>Click the <b>Power On</b> button.</p> <p>When the Panorama virtual appliance boots, the installation process is complete.</p>

Continue with [Perform Initial Configuration](#).

## Perform Initial Configuration

Use the Panorama virtual appliance console on the ESX(i) server to set up network access to the Panorama virtual appliance. To complete initial configuration, you must first configure the management interface, then access the Panorama web interface to add the serial number for the virtual appliance, and define the time zone for the Panorama virtual appliance. For unified reporting, consider using GMT or UTC as the uniform time zone across all the managed devices and Panorama.

CONFIGURE THE MANAGEMENT INTERFACE	
<p><b>Step 1.</b> Gather the required information from your network administrator.</p>	<ul style="list-style-type: none"> <li>IP address for MGT port</li> <li>Netmask</li> <li>Default gateway</li> <li>DNS server IP address</li> </ul>
<p><b>Step 2</b> Access the console of the Panorama virtual appliance.</p>	<ol style="list-style-type: none"> <li>Select the <b>Console</b> tab on the ESX(i) server for the virtual Panorama. Press enter to access the login screen.</li> <li>Enter the default username/password (admin/admin) to log in.</li> <li>Enter <b>configure</b> to switch to configuration mode.</li> </ol>

CONFIGURE THE MANAGEMENT INTERFACE	
<p><b>Step 3</b> Configure the network access settings for the management interface.</p> <p>The management interface is used for management traffic, HA connectivity synchronization, log collection, and for communication with Log Collector appliances.</p>	<p>Enter the following command:</p> <pre>set deviceconfig system ip-address &lt;Panorama-IP&gt; netmask &lt;netmask&gt; default-gateway &lt;gateway-IP&gt; dns-setting servers primary &lt;DNS-IP&gt;</pre> <p>where &lt;Panorama-IP&gt; is the IP address you want to assign to the Panorama management interface, &lt;netmask&gt; is the subnet mask, &lt;gateway-IP&gt; is the IP address of the network gateway, and &lt;DNS-IP&gt; is the IP address of the DNS server.</p>
<p><b>Step 4</b> Commit your changes and exit the configuration mode.</p>	<p>Enter <b>commit</b>.</p> <p>Enter <b>exit</b>.</p>
<p><b>Step 5</b> Verify network access to external services required for firewall management, such as the Palo Alto Networks Update Server.</p>	<p>To verify that Panorama has external network access, use the ping utility. Verify connectivity to the default gateway, DNS server, and the Palo Alto Networks Update Server as shown in the following example:</p> <pre>admin@Panorama-Corp&gt; ping host updates.paloaltonetworks.com PING updates.paloaltonetworks.com (67.192.236.252) 56(84) bytes of data. 64 bytes from 67.192.236.252: icmp_seq=1 ttl=243 time=40.5 ms 64 bytes from 67.192.236.252: icmp_seq=1 ttl=243 time=53.6 ms 64 bytes from 67.192.236.252: icmp_seq=1 ttl=243 time=79.5 ms</pre> <p><b>Note</b> After verifying connectivity, press Ctrl+C to stop the pings.</p>
ADD SERIAL NUMBER AND TIME ZONE	
<p><b>Step 1.</b> Log in to the Panorama web interface.</p>	<p>Using a secure connection (https) from a web browser, log in using the IP address and password you assigned to the management interface (https://&lt;IP address&gt;).</p>
<p><b>Step 2</b> (Optional) Modify the management interface settings.</p>	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Setup &gt; Management</b> and then click the Edit icon in the Management Interface Settings section of the screen.</li> <li>2. Select which management services to allow on the interface. For example, to enable SSH access, select <b>SSH</b>. As a best practice, make sure Telnet and HTTP are not selected because these services use plaintext and are not as secure as the other services.</li> <li>3. Click <b>OK</b>. Click <b>Commit</b> and select Panorama as the <b>Type</b> and click <b>OK</b>.</li> </ol>

ADD SERIAL NUMBER AND TIME ZONE (CONTINUED)	
<p><b>Step 3</b> Add the Panorama serial number.</p> <p>The serial number was sent to you with the order fulfillment email.</p>	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Setup &gt; Management</b> and click the Edit icon in the General Settings section of the screen.</li> <li>2. Enter the <b>Serial Number</b>.</li> </ol>
<p><b>Step 4</b> Configure time zone, and general firewall settings.</p>	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Setup &gt; Management</b> and click the Edit icon in the General Settings section of the screen.</li> <li>2. Align the clock on Panorama and the managed firewalls to use the same time zone, for example GMT or UTC.</li> </ol> <p>Timestamps are recorded when the logs are received on Panorama and when they were generated on the firewalls. Aligning the time zones on both Panorama and the managed devices ensures that the timestamps are in sync, and the process of querying logs and generating reports on Panorama is harmonious.</p> <ol style="list-style-type: none"> <li>3. Enter a <b>Hostname</b> for the server and enter the network <b>Domain</b> name. The domain name is just a label; it will not be used to join the domain.</li> <li>4. Enter the <b>Latitude</b> and <b>Longitude</b> to enable accurate placement of the server on the world map.</li> <li>5. Click <b>OK</b>.</li> </ol>
<p><b>Step 5</b> Change the default admin password.</p> <p><b>Note</b> To ensure that the management interface remains secure, consider enforcing <b>Minimum Password Complexity</b> and defining an interval at which administrators must change their passwords.</p>	<ol style="list-style-type: none"> <li>1. Click on the <b>admin</b> link in the lower left part of the management console. A dialog to change the administrator's password displays.</li> <li>2. Enter the old password, and the new password in the appropriate fields and store the new password in a safe location.</li> <li>3. Click <b>OK</b>.</li> </ol>
<p><b>Step 6</b> Save your configuration changes.</p>	<p>Click <b>Commit</b> and select Panorama as the <b>Type</b> and click <b>OK</b>.</p>



## Expand Log Storage Capacity on the Panorama Virtual Appliance

By default, the Panorama virtual appliance is set up with a single disk partition for all data, and ~10GB of this space is allocated for log storage. See the article [Panorama Logging Suggestions](#) to estimate the log storage capacity for your requirements, and then use one of the following options to expand the log storage capacity on the Panorama virtual appliance:

- [Add a Virtual Disk](#) on your ESX(i) server to expand storage to a maximum of 2TB.
- [Set up Access to an NFS Datastore](#) (NFS) datastore. Use this option if more than 2TB of log storage capacity is required.

### Add a Virtual Disk

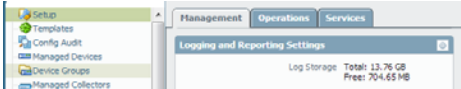
The Panorama virtual appliance by default installs with a virtual disk of size 34GB of which 10.89GB is used for logging. To enable more than the ~10GB of storage, use the following procedure to create a new virtual disk that can support up to 2TB of storage capacity.



The Panorama virtual appliance can only use one virtual disk. When configured to use a virtual disk, the virtual appliance does not use the default 10GB internal storage for logging. Therefore, if it loses connectivity to the virtual disk, logs could be lost 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.

ADD A VIRTUAL DISK	
<b>Step 1.</b> Power off the Panorama virtual appliance.	
<b>Step 2.</b> On the ESX(i) server, add the virtual disk to the Panorama virtual appliance.	<ol style="list-style-type: none"> <li>1. Select the Panorama virtual appliance on the ESX(i) server.</li> <li>2. Click <b>Edit Settings</b>.</li> <li>3. Click <b>Add</b> to launch the Add Hardware wizard, and select the following options when prompted: <ol style="list-style-type: none"> <li>a. Select <b>Hard Disk</b> for the hardware type.</li> <li>b. Select <b>Create a new virtual disk</b>.</li> <li>c. Select <b>SCSI</b> as the virtual disk type.</li> <li>d. Select the <b>Thick provisioning</b> disk format.</li> <li>e. In the location field, select <b>Store with the virtual machine option</b>.</li> </ol> </li> </ol> <p><b>Note</b> The datastore does not have to reside on the ESX(i) server.</p> <ol style="list-style-type: none"> <li>f. Verify that the settings look correct and click <b>Finish</b> to exit the wizard. The new disk is added to the list of devices for the virtual appliance.</li> </ol>

ADD A VIRTUAL DISK	
Step 3 Power on the Panorama virtual appliance.	<p>When powered on, the virtual disk is initialized for first-time use. The time that the initialization process takes to complete varies by the size of the new virtual disk.</p> <p>When the virtual disk is initialized and ready, all existing logs on the internal storage are moved over to the new virtual disk. All new entries will now be written to the virtual disk.</p>
Step 4 Verify the size of the virtual disk.	<div><div><div>1. Select <b>Panorama &gt; Setup &gt; Management</b>.</div><div>2. In the Logging and Reporting Settings section, verify that the <b>Log Storage</b> capacity accurately displays the new disk capacity.</div></div><div>A screenshot of the Panorama Management console. The left sidebar shows a tree view with 'Panorama' selected, and sub-items like 'Templates', 'Config Audit', 'Managed Devices', 'Device Groups', and 'Managed Collectors'. The main panel has tabs for 'Management', 'Operations', and 'Services'. The 'Management' tab is active, showing 'Logging and Reporting Settings'. It displays 'Log Storage' with 'Total: 13.76 GB' and 'Free: 704.65 MB'.</div></div>

## Set up Access to an NFS Datastore

Mounting the Panorama virtual appliance to an NFS datastore provides the ability to write logs to a centralized location and offers the flexibility to expand the log storage capacity beyond 2TB. Before setting up an NFS datastore in a Panorama high availability configuration, see [Logging Considerations in HA](#).

MOUNT AN NFS DATASTORE	
<p><b>Step 1.</b> Set up access to the datastore.</p>	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Setup &gt; Operations</b>.</li> <li>2. Click <b>Storage Partition Setup</b> link in the Miscellaneous section.</li> <li>3. Select <b>NFS V3</b>.</li> <li>4. Enter the IP address of the NFS <b>Server</b>.</li> <li>5. Enter the location/path for storing the log files in the <b>Log Directory</b> field. For example, export/panorama.</li> <li>6. Select the protocol—<b>TCP</b> or <b>UDP</b>—and enter the <b>Port</b> for accessing the NFS server.</li> </ol> <p><b>Note</b> 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.</p> <ol style="list-style-type: none"> <li>7. For optimal NFS performance, in the <b>Read Size</b> and <b>Write Size</b> 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.</li> <li>8. Select <b>Test Logging Partition</b> to verify that Panorama is able to access the NFS server IP address and the directory location specified above.</li> <li>9. (Optional) Select the <b>Copy on Setup</b> option. This setting copies the existing logs stored on Panorama to the NFS volume. If you have a lot of existing logs, enabling the copy on setup option, might initiate the transfer of a large volume of data.</li> <li>10. Click <b>Commit</b> and select <b>Panorama</b> as the <b>Commit Type</b> to save the changes.</li> </ol>
<p><b>Step 2</b> Reboot the Panorama virtual appliance.</p> <p>Until a reboot is initiated, logs will be written to the local storage disk on the Panorama virtual appliance.</p>	<p>To begin writing logs to the NFS datastore, reboot the virtual Panorama.</p> <ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Setup &gt; Operations</b>.</li> <li>2. In the Device Operations section, select <b>Reboot Panorama</b>.</li> </ol>

## Determine the Log Rate on the Palo Alto Networks Firewall

Use these instructions at different times during the day to approximate the normal and peak log generation rate on each firewall. In order to accurately estimate the amount of storage required for logs on your network, in addition to the log rate on the firewall, you must consider several other factors. For details, see the article: [Panorama Logging Suggestions](#).

VIEW THE LOG GENERATION RATE	
Step 1. Access the CLI on each Palo Alto Networks firewall.	See <a href="#">Log in to the CLI</a> ; the process of accessing the CLI on the firewall is the same as that on Panorama.
Step 2 View the current log generation rate.	<div>Enter the following CLI command to gauge the log rate on the firewall:</div> <div><b>debug log-receiver statistics</b></div> <div>Logging statistics</div> <div>-----</div> <div>Log incoming rate: 246/sec</div> <div>Log written rate: 246/sec</div>

## Where to go Next?

Now that initial configuration is complete, continue with the following sections for additional configuration instructions:

- ▲ [Install Licenses](#)
- ▲ [Install Content and Panorama Software Updates](#)
- ▲ [Navigate the Panorama User Interface](#)
- ▲ [Set Up Administrative Access](#)
- ▲ [Manage Your Firewalls](#)

## Set Up the M-100 Appliance

The M-100 management appliance is a high performance hardware platform that can be deployed in two modes:

- **Panorama mode:** The appliance performs both the central management and the log collection functions. This is the default mode.
- **Log Collector mode:** The appliance functions as a dedicated log collector, which can be managed by either an M-100 appliance in Panorama mode or a Panorama virtual appliance. If you have large volumes of log data being forwarded from multiple firewalls, the M-100 appliance in Log Collector mode provides increased scale and performance. When deployed in Log Collector mode, the appliance does not have a web interface; administrative access is CLI only.

Use the following workflow for setting up the M-100 appliance:

M-100 Appliance in Panorama Mode	M-100 Appliance in Log Collector Mode
<p><b>Step 1.</b> Rack mount the M-100 appliance. Refer to the <a href="#">M-100 Hardware Reference Guide</a> for instructions.</p> <p><b>Step 2</b> <a href="#">Perform Initial Configuration</a></p> <p><b>Step 3</b> <a href="#">Activate/Retrieve the Licenses</a></p> <p><b>Step 4</b> <a href="#">Install Content and Panorama Software Updates</a></p> <p><b>Step 5</b> (Optional) <a href="#">Increase Storage Capacity on the M-100 Appliance</a></p> <p><b>Step 6</b> <a href="#">Set Up Administrative Access</a></p> <p><b>Step 7</b> <a href="#">Manage Your Firewalls</a></p> <p><b>Step 8</b> <a href="#">Enable Logging</a></p>	<p><b>Step 1.</b> Rack mount the M-100 appliance. Refer to the <a href="#">M-100 Hardware Reference Guide</a> for instructions.</p> <p><b>Step 2</b> <a href="#">Perform Initial Configuration</a></p> <p><b>Step 3</b> <a href="#">Activate/Retrieve the Licenses</a></p> <p><b>Step 4</b> <a href="#">Install Content and Panorama Software Updates</a></p> <p><b>Step 5</b> (Optional) <a href="#">Increase Storage Capacity on the M-100 Appliance</a></p> <p><b>Step 6</b> <a href="#">Set Up the M-100 Appliance in Log Collector Mode</a></p> <p><b>Step 7</b> <a href="#">Enable Logging</a></p>

## Perform Initial Configuration

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 MGT interface or using a direct serial port connection to the console port on the M-100 appliance.

CONFIGURE THE MANAGEMENT INTERFACE	
<b>Step 1.</b> Gather the required information from your network administrator.	<ul style="list-style-type: none"> <li>• IP address for MGT port</li> <li>• Netmask</li> <li>• Default gateway</li> <li>• DNS server address</li> </ul>
<b>Step 2</b> Connect your computer to the M-100 appliance.	<p>You can connect to the M-100 appliance in one of the following ways:</p> <ul style="list-style-type: none"> <li>• Attach a serial cable from a computer to the Console port on the M-100 appliance and connect using a terminal emulation software (9600-8-N-1).</li> <li>• Attach an RJ-45 Ethernet cable from a computer to the MGT port on the M-100 appliance. From a browser, go to <b>https://192.168.1.1</b>. Note that this may require changing the IP address on the computer to an address in the 192.168.1.0 network, such as 192.168.1.2, in order to access this URL.</li> </ul>
<b>Step 3</b> When prompted, log in to the appliance.	<p>Log in using the default username and password (admin/admin). The appliance will begin to initialize.</p>
<b>Step 4</b> Configure the network access settings for the MGT interface.  The management interface is used for management traffic, HA connectivity synchronization, log collection, and for communicating with the Log Collector appliance(s).	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Setup</b> and then click the Edit icon in the Management Interface Settings section of the screen.</li> <li>2. Enter the <b>IP Address</b>, <b>Netmask</b>, and <b>Default Gateway</b>.</li> <li>3. (Optional) Select which management services to allow on the interface. For example, enable SSH. As a best practice, make sure Telnet and HTTP are not selected because these services use plaintext and are not as secure as the other services.</li> <li>4. Click <b>OK</b>. Click <b>Commit</b> and select Panorama as the <b>Commit Type</b> and click <b>OK</b>.</li> </ol>

CONFIGURE THE MANAGEMENT INTERFACE (CONTINUED)	
<p><b>Step 5</b> Configure the hostname, time zone, and general settings.</p>	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Setup &gt; Management</b> and click the Edit icon in the General Settings section of the screen.</li> <li>2. Align the clock on Panorama and the managed firewalls to use the same <b>Time Zone</b>, for example GMT or UTC.  Setting the same time zone on both Panorama and the managed devices ensures that the timestamps on the logs are in sync. Timestamps are recorded when the logs are received on Panorama and when they were generated on the firewalls. Aligning the time zones on both Panorama and the managed devices ensures that the timestamps are in sync, and the process of querying logs and generating reports on Panorama is harmonious.</li> <li>3. Enter a <b>Hostname</b> for the server. This hostname will be used as the display name/label for the appliance. For example, this is the name that will be displayed at the CLI prompt; and the name displayed in the Collector Name field when you add the appliance as a Managed Collector on the <b>Panorama &gt; Managed Collectors</b> tab.</li> <li>4. Enter your network <b>Domain</b> name. The domain name is just a label; it will not be used to join the domain.</li> <li>5. (Optional) Enter the <b>Latitude</b> and <b>Longitude</b> to enable accurate placement of the server on the world map. The latitude and longitude values that will be used in the <b>App Scope &gt; Traffic Maps</b> and <b>App Scope &gt; Threat Maps</b>.</li> <li>6. Click <b>OK</b>.</li> </ol>
<p><b>Step 6</b> Change the default admin password.</p> <p><b>Note</b> To ensure that the management interface remains secure, enforce <b>Minimum Password Complexity</b> and specify the interval at which administrators must change their passwords.</p>	<ol style="list-style-type: none"> <li>1. Click on the <b>admin</b> link in the lower left part of the management console. A dialog to change the administrator's password displays.</li> <li>2. Enter the old password, and the new password in the appropriate fields and store the new password in a safe location. Click <b>OK</b>.</li> <li>3. Click <b>Commit</b> and select Panorama as the <b>Commit Type</b>.</li> </ol>

**CONFIGURE THE MANAGEMENT INTERFACE (CONTINUED)**

<b>Step 7</b> Verify network access to external services required for firewall management, such as the Palo Alto Networks Update Server	<p>To verify that Panorama has external network access, use the ping utility. Verify connectivity to the default gateway, DNS server, and the Palo Alto Networks Update Server as shown in the following example:</p> <pre>admin@Panorama-Corp&gt; ping host updates.paloaltonetworks.com  PING updates.paloaltonetworks.com (67.192.236.252) 56(84) bytes of data.  64 bytes from 67.192.236.252: icmp_seq=1 ttl=243 time=40.5 ms  64 bytes from 67.192.236.252: icmp_seq=1 ttl=243 time=53.6 ms  64 bytes from 67.192.236.252: icmp_seq=1 ttl=243 time=79.5 ms</pre> <p><b>Note</b> After verifying connectivity, press Ctrl+C to stop the pings.</p>
---	---

Continue with [Install Licenses](#) and [Install Content and Panorama Software Updates](#), regardless of whether you plan on using the M-100 appliance in Panorama mode or in Log Collector mode.

## Set Up the M-100 Appliance in Log Collector Mode

Using the M-100 appliance as a Log Collector offloads the task of processing logs to a dedicated appliance. Use the instructions in this section to convert the M-100 appliance from Panorama mode to the Log Collector mode for deploying it as a dedicated Log Collector. Ensure that the Panorama appliance that will manage the firewalls and the Log Collector is already set up.



In the Log Collector mode, the M-100 appliance does not support the web interface for configuration tasks; only SSH access is supported. Therefore, before changing the mode on the M-100 appliance [Perform Initial Configuration](#) and use the web interface in Panorama mode to [Activate/Retrieve the Licenses](#).

To send logs to an M-100 appliance in Log Collector mode, the Palo Alto Networks firewalls must be running PAN-OS v5.0 or later versions. Palo Alto Networks firewalls running PAN-OS versions earlier than 5.0 can only send logs to an M-100 appliance in Panorama mode or to a Panorama virtual appliance.

**SWITCH FROM PANORAMA MODE TO LOG COLLECTOR MODE**

<b>Step 1.</b> Access the Command Line Interface (CLI) on the M-100 appliance.	<p>Connect to the M-100 appliance in one of the following ways:</p> <ul style="list-style-type: none"> <li>• Attach a serial cable from a computer to the Console port on the M-100 appliance. Then, connect using a terminal emulation software (9600-8-N-1).</li> <li>• Use a terminal emulation software such as PuTTY, to open an SSH session to the IP address assigned to the M-100 appliance during initial configuration.</li> </ul>
--	--



SWITCH FROM PANORAMA MODE TO LOG COLLECTOR MODE (CONTINUED)	
<b>Step 2</b> When prompted, log in to the appliance.	Use the default <i>admin</i> account and the password assigned during initial configuration.
<b>Step 3</b> Switch from the Panorama mode to the Log Collector mode.	<ol style="list-style-type: none"> <li>To switch to log collector mode, enter the following command: <b>request system logger-mode logger</b></li> <li>Enter <b>Yes</b> to confirm the change to log collector mode. The appliance will reboot.</li> </ol>
<b>Step 4</b> Verify that the appliance is in Log Collector mode.	<ol style="list-style-type: none"> <li>Log back in to the CLI on the M-100 appliance.</li> <li>Enter the following command: <b>show system info   match logger_mode</b> The response printed on screen reads as <i>logger_mode: True</i> If the value displays as <b>False</b>, the M-100 appliance is still in Panorama mode.</li> </ol>
<b>Step 5</b> Specify the IP address of the Panorama appliance that is managing the Log Collector.	Enter the following command in the CLI: <b>configure</b> <b>set panorama-server &lt;ip_address&gt;</b> <b>commit</b>

Now that you have successfully set up your M-100 appliance, for further instructions on assigning a Log Collector to a firewall, defining Collector Groups, and managing the Log Collector using Panorama, see [Enable Logging](#).

## Increase Storage Capacity on the M-100 Appliance

The M-100 appliance ships with two disks in a RAID1 configuration. Each M-100 appliance allows for the addition of up to three additional disk pairs in RAID1, each with a storage capacity of 1TB, to reach a maximum capacity of 4 TB RAID storage.



If adding disk pairs to an already deployed M-100 appliance, there is no need to take the system offline to expand the storage capacity. When the additional disk pair(s) become available, the M-100 appliance will redistribute the logs amongst the available disk pairs. This log redistribution process happens in the background and does not impact uptime or availability of the M-100 appliance.

SET UP THE DISKS IN A RAID PAIR	
<b>Step 1.</b> Install the new disks in the appropriate drive bays.	Make sure to add the drives sequentially in the next open disk bay slot for the disk pair. For example, add B1/B2 before C1/C2.  For information on adding the physical drives, refer to the <a href="#">M-100 Hardware Reference Guide</a> .

SET UP THE DISKS IN A RAID PAIR (CONTINUED)	
<p><b>Step 2</b> Access the Command Line Interface (CLI) on the M-100 appliance.</p>	<p>You can connect to the M-100 appliance in one of the following ways:</p> <ul style="list-style-type: none"> <li>• Connect a serial cable from your computer to the Console port and connect to the M-100 appliance using terminal emulation software (9600-8-N-1).</li> <li>• Use a terminal emulation software such as PuTTY to open an SSH session to the IP address of the M-100 appliance.</li> </ul>
<p><b>Step 3</b> When prompted, log in to the appliance.</p>	<p>Use the default <i>admin</i> account and the password assigned.</p>
<p><b>Step 4</b> Set up each additional disk pair in a RAID configuration.</p> <p><b>Note</b> The time required to mirror the data on the drive may vary from several minutes to a couple hours, depending on the amount of data on the drive.</p>	<p>This example uses the drives in the disk bays B1 and B2.</p> <ol style="list-style-type: none"> <li>1. Enter the following commands and confirm the request when prompted:           <pre>request system raid add B1 request system raid add B2</pre> </li> <li>2. To monitor the progress of the RAID configuration, enter the following command:           <pre>show system raid detail</pre> <p>When the RAID set up is complete, the following response displays:</p> <pre> 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     status      : active sync    Disk id B2     Present     model       : ST91000640NS     size        : 953869 MB     status      : active sync           </pre> </li> </ol>

**SET UP THE DISKS IN A RAID PAIR (CONTINUED)**

<p><b>Step 5</b> Make the disk pair available for logging.</p> <p>To enable the disk pairs for logging, this appliance must have been added as a Managed Collector on Panorama. If you have not already added it, see <a href="#">Add a Log Collector to Panorama</a>.</p>	<ol style="list-style-type: none"> <li>1. Access the Panorama management server that is managing this log collector (if it is a different appliance).</li> <li>2. On the <b>Panorama &gt; Managed Collectors</b> tab, select the log collector and follow the instructions in <a href="#">Step 6</a> in <a href="#">Add a Log Collector to Panorama</a>.</li> </ol>
<p><b>Step 6</b> Save your configuration changes.</p>	<p>Click <b>Commit</b>. Select <b>Panorama</b> as the <b>Commit Type</b> and click <b>OK</b>.</p>

For further instructions on adding a Log Collector as a Managed Collector on Panorama, defining Collector Groups, assigning a Log Collector to a firewall, see [Enable Logging](#).

# Migrate from a Panorama Virtual Appliance to an M-100 Appliance

On a Panorama virtual appliance that manages 10 or more devices and has a logging rate of over 10,000 logs per second, migrating to the M-100 appliance will provide improved response time on the web interface and speedier execution of reports. The M-100 appliance also provides up to 4TB of RAID storage. Use the instructions in this section to migrate the configuration from the Panorama virtual appliance over to an M-100 appliance.

- ▲ [Prerequisites](#)
- ▲ [Planning Considerations](#)
- ▲ [Perform the Migration](#)
- ▲ [Resume Managing the Devices](#)

## Prerequisites

To proceed with the migration of your current subscription, you must have:

- Purchased an M-100 appliance
- Obtained a migration upgrade, and purchased a new subscription that includes software and hardware support.

To process the migration upgrade you must have contacted your sales representative with the following:

- Serial number of the virtual Panorama that you plan to phase out
- Support terms for the M-100 appliance and the auth-code that you received when you purchased the appliance
- Effective date for the migration

Palo Alto Networks will automatically apply the associated authorization codes to the serial number of your management appliance, phase out support on your existing virtual Panorama, and trigger the support for your M-100 appliance on the chosen effective date.

Starting at the effective date, you will have a limited time to complete the migration process. At the end of the period, the support entitlement on the Panorama virtual appliance will be terminated and you will no longer be able to receive software or threat updates. Refer to this [article](#) for the details on the license migration process.

## Planning Considerations

- Plan on completing this migration during a maintenance window. Although the firewalls can buffer the logs and forward them to Panorama when the connection is reestablished, completing the migration during a maintenance window minimizes loss of log data during the transition time when the Panorama virtual appliance goes offline and the M-100 appliance comes online.
- Consider whether to maintain access to the Panorama virtual appliance after completing the migration. Because the log format on the Panorama virtual appliance is incompatible with that on the M-100 appliance, existing log data cannot migrate over to the M-100 appliance. Therefore, to access the old logs the Panorama virtual appliance must remain accessible.
- Decide whether to use the same IP address on the M-100 appliance or to assign a new one. Palo Alto Networks recommends reusing the same management IP address to prevent the need to reconfigure each managed device to point to a new IP address.



If you have log compliance requirements, plan to reconfigure a new IP address on the Panorama virtual appliance to maintain access to the old logs for generating reports.

- Keep a *new* IP address, at hand for use in setting up connectivity to the M-100 appliance during initial configuration. If you have decided to transfer the IP address that was assigned to the Panorama virtual appliance, this new IP address will be used temporarily. When you restore the configuration file from the Panorama virtual appliance on the M-100 appliance, this *new* IP address will be overwritten.

## Perform the Migration

To migrate the configuration from the Panorama virtual appliance to the M-100 appliance, complete the following tasks:

### MIGRATE FROM THE PANORAMA VIRTUAL APPLIANCE

**Step 1.** Complete these tasks on the Panorama virtual appliance.

1. Upgrade to the latest Panorama version.	See <a href="#">Install Content and Panorama Software Updates</a> .
2. Export the running configuration on the virtual Panorama.	<ol style="list-style-type: none"> <li>1. In the <b>Panorama &gt; Setup &gt; Operations</b> tab, Configuration Management section, select <b>Export named Panorama configuration snapshot</b>.</li> <li>2. Select the active configuration (running-config.xml) and click <b>OK</b>. The file is downloaded and saved to the local machine.</li> <li>3. Rename the file.</li> </ol>
3. Power off the VM or change the IP address.	<p>If you plan on reusing the MGT interface IP address that was configured on the Panorama virtual appliance on the M-100 appliance, you can either power off the virtual appliance or assign a new IP address to the MGT port on the virtual appliance.</p> <p>To change the IP address, on the <b>Panorama &gt; Setup</b> tab, edit the <b>Management Interface Settings</b> section and enter the new IP address.</p>

**Step 2** Complete these tasks on the M-100 appliance.

1. Set up network access.	<p>See <a href="#">Perform Initial Configuration</a> for instructions.</p> <p>Consider assigning a new <i>temporary</i> IP address during initial configuration on the M-100 appliance and reusing the IP address that was assigned to the Panorama virtual appliance. The temporary IP address will be overwritten when you import the configuration later in this process.</p>
2. Install the same Panorama version as that running on the Panorama virtual appliance.	<p>Install the same Panorama version that you selected in Step 1 above. For instructions on performing the upgrade, see <a href="#">Install Content and Panorama Software Updates</a>.</p>
3. Register Panorama and retrieve the license.	See <a href="#">Install Licenses</a> .

MIGRATE FROM THE PANORAMA VIRTUAL APPLIANCE (CONTINUED)	
4. Import and load the configuration file.	<ol style="list-style-type: none"> <li>1. In the <b>Panorama &gt; Setup &gt; Operations</b> tab, Configuration Management section, select <b>Import named Panorama configuration snapshot</b>.</li> <li>2. <b>Browse</b> to select the running-config.xml (or the renamed file) and click <b>OK</b>.</li> <li>3. Select the <b>Load named Panorama configuration snapshot</b> link to load the configuration file you just imported. Any errors that occur when loading the configuration file are displayed onscreen.</li> <li>4. If there were errors save them to a local file. Review and resolve each error to make certain that all components of the configuration have been migrated over.</li> </ol>
5. Review and modify the configuration on Panorama.	<ol style="list-style-type: none"> <li>1. If you do not plan to reuse the same network access settings for the MGT interface, modify the values:               <ol style="list-style-type: none"> <li>a. Select <b>Panorama &gt; Setup</b> and then click the Edit icon in the Management Interface Settings section of the screen.</li> <li>b. Enter the <b>IP Address</b>, <b>Netmask</b>, and <b>Default Gateway</b>.</li> <li>c. Confirm that the list of IP addresses defined in the Permitted IP Addresses list is accurate.</li> </ol> </li> <li>2. To change the hostname, edit the General Settings section of the <b>Panorama &gt; Setup</b> tab.</li> <li>3. Confirm that the administrative access settings (administrators, roles and access domains) configured on the appliance is accurate on the <b>Panorama &gt; Administrators</b> tab, <b>Panorama &gt; Admin Roles</b> tab, and the <b>Panorama &gt; Access Domains</b> tab.</li> </ol>
6. Add the default Log Collector back to the M-100 appliance.	When importing the configuration from the Panorama virtual appliance, the default log collector is removed from the M-100 appliance. To add the log collector back on the M-100 appliance, use the instructions in <a href="#">Add a Log Collector to Panorama</a> .
7. Save all your changes to Panorama.	After reviewing the configuration changes, click <b>Commit</b> . Select Panorama as the <b>Commit Type</b> and click <b>OK</b> .

## Resume Managing the Devices



To resume central management, you must restore connectivity to the managed devices. Complete this task during a maintenance window to minimize network disruption.

### VERIFY THE STATUS OF THE MANAGED DEVICES

**Step 1.** Log in to Panorama.

Using a secure connection (https) from a web browser, log in using the IP address and password assigned during initial configuration (https://<IP address>).

**Step 2** Synchronize the configuration on Panorama with that of the managed device.

1. Select **Panorama > Managed Devices**, and verify that the **Connected** status of each device displays as . The status for the Templates and Device Groups will display as  Out of sync.
2. To synchronize the device groups:
  - a. Click **Commit** and select **Device Groups** as the **Commit Type**.
  - b. Select each device group and click **OK**.
3. To synchronize the templates:
  - a. Click **Commit** and select **Panorama** as the **Commit Type**.
  - b. Click **Commit** and select **Template** as the **Commit Type**.

**Step 3** Verify that the status of the devices, templates and shared policy is **Connected** and **In sync**.

								Status	
<input type="checkbox"/>	Device Name	Virtual System	Tags	Serial Number	IP Address	Template	Connected	Shared Policy	Template
 Desk_FWs (3/3 Devices Connected)									
<input type="checkbox"/>	Corp_gateway		Corp, NAmerica	001606000100	192.168.1.1	MJS_Template		 In sync	 In sync



## Install Licenses

Before you can begin using Panorama for centralized management, logging, and reporting, you must register the Panorama and retrieve the licenses.

Every instance of Panorama requires valid licenses that entitle you to manage the devices and to obtain support. The device management license enforces the maximum number of devices that can be managed by Panorama. The support license enables Panorama software updates and dynamic content updates for the latest application and threat signatures, among other updates, that are published by Palo Alto Networks.

To purchase licenses, contact your Palo Alto Networks Systems Engineer or reseller. After obtaining a license, navigate to **Panorama > Licenses** to perform the following tasks depending on how you receive your licenses:

- **Retrieve license keys from license server**—Use this option if the license has been activated on the support portal.
- **Activate feature using authorization code**—Use the authorization code to activate a license that has not been previously activated on the support portal.
- **Manually upload license key**—Use this option if Panorama does not have connectivity to the Palo Alto Networks update server. In this case, first download the license key file from the support site to an Internet-connected computer and then upload it to Panorama.


## Register Panorama

To manage all the assets you purchase from Palo Alto Networks, you must create an account and register the serial numbers with the account.

REGISTER WITH PALO ALTO NETWORKS	
<b>Step 1.</b> Log in to the Panorama web interface.	Using a secure connection (https) from a web browser, log in using the IP address and password you assigned during initial configuration (https://<IP address>).
<b>Step 2</b> Locate your serial number and copy it to the clipboard.	<ul style="list-style-type: none"> <li>• The serial number for the M-100 appliance displays on the <b>Dashboard</b>; locate the <b>Serial Number</b> in the General Information section of the screen.</li> <li>• For the Panorama virtual appliance, the serial number was included in the order fulfillment email.</li> </ul>
<b>Step 3</b> Go to the Palo Alto Networks Support site.	In a new browser tab or window, go to <a href="https://support.paloaltonetworks.com">https://support.paloaltonetworks.com</a> .

**REGISTER WITH PALO ALTO NETWORKS (CONTINUED)**

**Step 4** Register Panorama. The way you register depends on whether you already have a login to the support site.



- If this is the first Palo Alto Networks appliance you are registering and you do not yet have a login, click **Register** on the right side of the page. To register, provide your email address and the serial number for Panorama (which you can paste from your clipboard). When prompted, set up a username and password for access to the Palo Alto Networks support community.
- If you already have a support account, log in and then click **My Devices**. Scroll down to Register Device section at the bottom of the screen and enter the serial number for Panorama (which you can paste from your clipboard), your city and postal code and then click **Register Device**.

## Activate/Retrieve the Licenses

Every Panorama appliance —the virtual form factor and the hardware-based appliance—requires a valid license.



If you are running an evaluation license on your Panorama virtual appliance and want to apply a Panorama license that you have purchased:

1. Register the Panorama serial number on the Palo Alto Networks Support site. See [Step 4 in Register Panorama](#).
2. Select **Panorama > Setup > Management** and click the Edit icon in the General Settings section.
3. Enter the **Serial Number** for the Panorama virtual appliance and click **Commit**, to commit your changes to Panorama. The license is automatically applied to Panorama.

**ACTIVATE THE LICENSE**

**Step 1.** Locate the authorization codes for the product/subscription you purchased.

When you placed your order, you must have received an email from Palo Alto Networks customer service listing the auth-code associated with the purchase. If you cannot locate this email, contact customer support to obtain your codes before you proceed.

ACTIVATE THE LICENSE (CONTINUED)	
<p><b>Step 2</b>    Activate the license.</p> <p>The M-100 appliance requires both the support subscription and the device management license.</p> <p>The Panorama virtual appliance only requires a support subscription. The device management capability was enabled when you added the serial number.</p> <p><b>Note</b>    If the management port on Panorama does not have Internet access, manually download the license files from the support site and upload it to Panorama using the <b>Manually upload license key</b> option.</p>	<ol style="list-style-type: none"> <li>1. To activate your support subscription, select <b>Panorama &gt; Support</b>.</li> <li>2. Select <b>Activate feature using authorization code</b>. Enter the <b>Authorization Code</b> and then click <b>OK</b>.</li> <li>3. Verify that the subscription was successfully activated. <div data-bbox="917 426 1339 573" data-label="Image"> <p>A screenshot of the 'Support' portal. It displays the following information: Phone: 866-898-9087, Email: support@paloaltonetworks.com, ExpiryDate: October 01, 2015, Level: Premium, and Description: 24 x 7 phone support; advanced replacement hardware service. At the bottom, there is a link that says 'Activate support using authorization code'.</p> </div> </li> <li>4. (Only required for the M-100 appliance) In the <b>Panorama &gt; Licenses</b> tab, select <b>Activate feature using authorization code</b>.</li> <li>5. When prompted, enter the <b>Authorization Code</b> for using Panorama and click <b>OK</b>.</li> <li>6. Verify that the license was successfully activated and that it displays support for the appropriate number of devices. For example: <div data-bbox="917 831 1372 972" data-label="Image"> <p>A screenshot of the 'Device Management License' portal. It displays the following information: Date Issued: September 04, 2012, Date Expires: Never, and Description: Device management license to manage up to 25 devices.</p> </div> </li> </ol> <p><b>Note</b>    For the Panorama virtual appliance, you can view the device management license only on the Support portal.</p>
<p><b>Step 3</b>    (Not required if you completed Step 2) Retrieve license keys from the license server.</p>	<p>Use the <b>Retrieve license keys from the license server</b> option if you have activated the license keys on the Support portal.</p> <p>Select <b>Panorama &gt; Support</b>, and select <b>Retrieve license keys from the license server</b>.</p>

## Install Content and Panorama Software Updates

A valid support subscription enables to the software image and the Release Notes for Panorama. To take advantage of the latest fixes and security enhancements, it is a good idea to upgrade to the latest software update or to the update version recommended by your reseller or Palo Alto Networks Systems Engineer.

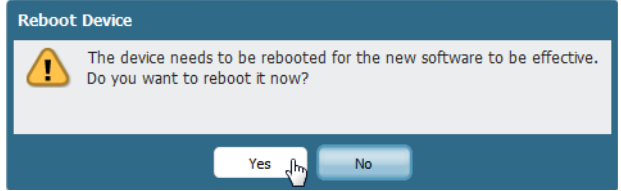


**Important:** Panorama version 5.1 is only available as a 64-bit OS. Before upgrading a Panorama virtual appliance to v5.1, make sure that the ESX(i) host supports a 64-bit OS and that it meets the minimum system requirements for the 64-bit OS. See [Prerequisites](#) for more information.

If managing devices with additional subscriptions, such as Threat Prevention or WildFire, Panorama also requires the content updates for the applications and threats databases. Your support subscription allows you to obtain these updates. The applications and threat databases are referenced in the policy configurations and are used when generating reports; these databases are used to match the identifiers recorded in the logs with the corresponding threat, URL, or application names. Therefore, to prevent a mismatch, Palo Alto Networks recommends that you install the same applications and threats database version on Panorama and on the managed devices.

PERFORM THE CONTENT AND PANORAMA VERSION UPDATES																																					
<p><b>Step 1.</b> Launch the Panorama web interface and go to the dynamic updates page.</p> <p>Before updating the software, install the latest content updates supported in the release.</p>	<ol style="list-style-type: none"><li>Using a secure connection (https) from a web browser, log in using the IP address and password you assigned during initial configuration (https://&lt;IP address&gt;).</li><li>Select <b>Panorama &gt; Dynamic Updates</b>.</li></ol>																																				
<p><b>Step 2</b> Check for, download, and install the latest content database update.</p> <p>Install the Application and Threat updates before installing the Antivirus update.</p>	<ol style="list-style-type: none"><li>Click <b>Check Now</b> to check for the latest updates. If the value in the Action column is <b>Download</b> it indicates that an update is available.</li><li>Click <b>Download</b> to obtain the desired version.</li><li>Click the <b>Install</b> link in the Action column. When the installation completes, a check mark displays in the Currently Installed column.</li></ol>																																				
<p><b>Step 3</b> Check for software updates.</p>	<ol style="list-style-type: none"><li>Select <b>Panorama &gt; Software</b>.</li><li>Click <b>Check Now</b> to check for the latest updates. If the value in the Action column is <b>Download</b> it indicates that an update is available.</li></ol>																																				
<p><b>Step 4</b> Download the update.</p> <p><b>Note</b> If Panorama does not have Internet access from the management port, you can download the software update from the <a href="#">Palo Alto Networks Support Site</a>. You can then manually <b>Upload</b> it to Panorama.</p>	<p>Locate the version you want to upgrade to, and click <b>Download</b>. When the download completes, the value in the Action column changes to <b>Install</b>.</p> <table><tr><th>Version</th><th>Size</th><th>Release Date</th><th>Downloaded</th><th>Currently Installed</th><th>Action</th></tr><tr><td>5.0.0</td><td>259 MB</td><td>2012/11/01 19:58:24</td><td>✓</td><td></td><td>Install</td></tr><tr><td>4.1.9</td><td>169 MB</td><td>2012/11/05 23:40:31</td><td></td><td></td><td>Download</td></tr><tr><td>4.1.8</td><td>168 MB</td><td>2012/09/22 21:01:08</td><td>✓</td><td>✓</td><td>Download</td></tr><tr><td>4.1.8-h3</td><td>168 MB</td><td>2012/10/18 23:49:21</td><td></td><td></td><td>Download</td></tr><tr><td>4.1.7</td><td>152 MB</td><td>2012/07/29 00:56:06</td><td></td><td></td><td>Download</td></tr></table>	Version	Size	Release Date	Downloaded	Currently Installed	Action	5.0.0	259 MB	2012/11/01 19:58:24	✓		Install	4.1.9	169 MB	2012/11/05 23:40:31			Download	4.1.8	168 MB	2012/09/22 21:01:08	✓	✓	Download	4.1.8-h3	168 MB	2012/10/18 23:49:21			Download	4.1.7	152 MB	2012/07/29 00:56:06			Download
Version	Size	Release Date	Downloaded	Currently Installed	Action																																
5.0.0	259 MB	2012/11/01 19:58:24	✓		Install																																
4.1.9	169 MB	2012/11/05 23:40:31			Download																																
4.1.8	168 MB	2012/09/22 21:01:08	✓	✓	Download																																
4.1.8-h3	168 MB	2012/10/18 23:49:21			Download																																
4.1.7	152 MB	2012/07/29 00:56:06			Download																																

**PERFORM THE CONTENT AND PANORAMA VERSION UPDATES (CONTINUED)**

<p><b>Step 5</b> Install the update.</p>	<ol style="list-style-type: none"> <li>1. Click <b>Install</b>.</li> <li>2. Reboot Panorama: <ul style="list-style-type: none"> <li>• If prompted to reboot, click <b>Yes</b>.</li> </ul> </li> </ol>  <p>The dialog box titled "Reboot Device" contains a yellow warning icon and the text: "The device needs to be rebooted for the new software to be effective. Do you want to reboot it now?". At the bottom, there are two buttons: "Yes" and "No". A mouse cursor is pointing at the "Yes" button.</p> <ul style="list-style-type: none"> <li>• If you are not prompted to reboot, select <b>Panorama &gt; Setup &gt; Operations</b> and click <b>Reboot Panorama</b> in the Device Operations section of the screen.</li> </ul>
<p><b>Step 6</b> (Only required for a Panorama virtual appliances upgrading to Panorama version 5.1) Modify the settings on the Panorama virtual appliance.</p> <p><b>Important:</b> Before powering on a Panorama virtual appliance on 5.1, make sure that the ESX(i) host supports a 64-bit OS and that it meets the minimum system requirements for the 64-bit OS. See <a href="#">Prerequisites</a> for more information.</p>	<p>After Panorama reboots, complete the following tasks:</p> <ol style="list-style-type: none"> <li>1. Power off the virtual appliance.</li> <li>2. Right click and select <b>Edit Settings...</b> to modify these parameters: <ol style="list-style-type: none"> <li>a. On the Options tab, change the Guest Operating System from <b>Other Linux (32-bit)</b> to <b>Other Linux (64-bit)</b>.</li> <li>b. On the Hardware tab, change the SCSI Controller from <b>BusLogic Parallel</b> to <b>LSI Logic Parallel</b>.</li> <li>c. On the Hardware tab, change the memory allocation to 4GB minimum; 16GB for managing 10 or more firewalls.</li> </ol> </li> <li>3. Power on the virtual appliance.</li> </ol>

To continue with managing the firewalls and enabling log collection, see [Chapter 3, Manage Firewalls and Log Collection](#).

## Navigate the Panorama User Interface

Panorama provides three user interfaces: a web interface, a command line interface (CLI), and a REST management API.

- ▲ **Web Interface**—The Panorama web interface is purposefully designed with a similar look and feel to the firewall. If you are already familiar with the firewall, you will be able to navigate and complete administrative tasks and generate reports tasks from the Panorama web interface with relative ease. This graphical interface allows you to access Panorama using HTTPS and it is the best way to perform administrative tasks. You can enable HTTP access to Panorama, if required on the Management Interface Settings section on the **Panorama > Setup > Management** tab. See [Navigate the Web Interface](#) and [Log in to the Web Interface](#).
- ▲ **Command Line Interface**—The Command Line Interface is a no-frills interface that allows you to type through the commands in rapid succession to complete a series of tasks. The CLI supports two command modes—operational and configuration—and each mode has its own hierarchy of commands and statements. When you get familiar with the nesting structure and the syntax for the commands, the CLI allows quick response times and offers administrative efficiency. See [Log in to the CLI](#).
- ▲ **REST Management API**—The XML-based REST API is provided as a web service that is implemented using HTTP/HTTPS requests and responses. It allows you to streamline your operations and integrate with existing, internally developed applications and repositories. For information on how to use the Panorama API interface, refer to the document [PAN-OS and Panorama XML-Based REST API](#). To access the online community for developing scripts, visit: <https://live.paloaltonetworks.com/community/devcenter>

## Navigate the Web Interface

Use the Panorama web interface to configure Panorama, manage and monitor the managed devices and Log Collectors, and to access the web interface of each managed device using the Device Context. Refer to the online help on Panorama for details on the options in each tab in the web interface.

The Panorama web interface includes the following tabs:

Tab	Sub-Tab	Description
Dashboard		Displays general information about the Panorama network access settings and model. It includes widgets that display information on applications, logs, and system resources and settings.
ACC		Displays the overall risk and threat level on the network, based on information gathered from the managed devices.
Monitor		Provides access to logs and reports.
Panorama		Configure Panorama, manage licenses, set up high availability, access software updates and security alerts, manage administrative access, and manage the deployed firewalls and log collectors.

Tab	Sub-Tab	Description (Continued)
Device Groups (You must <a href="#">Create Device Groups</a> for this tab to display.)	Policies	Create centralized policies and apply the configuration to multiple devices/device groups.
	Objects	Define policy objects that can be referenced in policy and shared across all managed devices/device groups.
Templates (You must <a href="#">Add a New Template</a> for this tab to display.)	Network	Configure network setting, such as network profiles, that can be applied to the managed devices.
	Device	Configure device configuration, such as server profiles and admin roles, that can be applied to the managed devices.

## Log in to the Web Interface

LOG IN TO THE WEB INTERFACE	
<b>Step 1.</b> Log in to the Panorama web interface.	Using a secure connection (https) from a web browser, log in using the IP address and password you assigned during initial configuration (https://<IP address>).
<b>Step 2</b> (Optional) Enable HTTP and SSH access.	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Setup &gt; Management</b> and then click the Edit icon in the Management Interface Settings section of the screen.</li> <li>2. Select which management services to allow on the interface. For example, select <b>HTTP</b> and <b>SSH</b>.</li> <li>3. Click <b>OK</b>.</li> </ol>

## Log in to the CLI

You can log in to the Panorama CLI using a serial port connection or access remotely using an SSH client.

LOG IN TO THE CLI	
<ul style="list-style-type: none"> <li>Use SSH to log in to the Panorama CLI.</li> </ul> <p><b>Note</b> The same instructions apply for an M-100 appliance in Log Collector mode.</p>	<ol style="list-style-type: none"> <li>Make sure that you have the following: <ul style="list-style-type: none"> <li>A computer with network access to Panorama</li> <li>Panorama IP address</li> <li>SSH is enabled on the Management interface. To enable SSH access, see <a href="#">(Optional) Enable HTTP and SSH access</a>.</li> </ul> </li> <li>To access the CLI using SSH: <ol style="list-style-type: none"> <li>Enter the Panorama IP address in the SSH client.</li> <li>Use port 22.</li> <li>Enter your administrative access credentials when prompted. After successfully logging in, the CLI prompt displays in operational mode. For example: admin@ABC_Sydney&gt;</li> </ol> <p>To enable key-based authentication, see <a href="#">Enable SSH Key-Based Authentication for the Command Line Interface</a>.</p> </li> </ol>
<ul style="list-style-type: none"> <li>Change to configuration mode.</li> </ul>	<p>To go into configuration mode, enter the following command at the prompt:</p> <pre>admin@ABC_Sydney&gt; <b>configure</b></pre> <p>The prompt changes to admin@ABC_Sydney#</p>
<ul style="list-style-type: none"> <li>Use a serial port connection to log in to the Panorama CLI.</li> </ul>	<ol style="list-style-type: none"> <li>Make sure that you have the following: <ul style="list-style-type: none"> <li>A null-modem serial cable that connects Panorama to a computer with a DB-9 serial port</li> <li>A terminal emulation program running on the computer</li> </ul> </li> <li>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.</li> <li>Enter your administrative access credentials when prompted.</li> </ol>



## Set Up Administrative Access

By default, Panorama includes a default administrative account (`admin`), with full read-write access to all the functionality on Panorama. As a best practice, create a separate administrative account for each person who needs access to the administrative or reporting functions of Panorama. This prevents unauthorized configuration (or modification) and enables logging of the actions of each individual administrator.

Panorama allows you to define and restrict access as broadly or granularly as required, depending on the security requirements within your organization. For example, you may decide that a datacenter administrator can have access to all the device and networking configuration, while a security administrator can have control over security policy definition, the log viewer and reporting, and other key individuals can have limited CLI or XML API access.



You cannot add an administrative account to an M-100 appliance in Log Collector mode. Only the default administrative user account with the default username `admin` is available.

The following sections describe the supported methods for setting up administrative accounts and provides procedures for setting up basic administrative access:





- ▲ [Create an Administrative Account](#)
- ▲ [Define Access Domains](#)
- ▲ [Define an Authentication Sequence](#)
- ▲ [Define Access Domains](#)
- ▲ [Configure Administrative Authentication](#); for information on the different options available to authenticate administrative users, see [Administrative Authentication](#).

### Create an Administrative Account

An administrative user must have an account and be assigned to a *role*. The role defines the type of access the associated administrator has to Panorama; you can assign the administrative user to a built-in Dynamic Role or to a custom role (Admin Role Profile) that you define. If you plan to use Admin Role Profiles rather than Dynamic Roles, create the profiles that define what type of access, if any, to give to the different sections of the web interface, the CLI, and XML API for each administrator assigned to the role. For more information on roles, see [Administrative Roles](#).

For each administrative user you can also define the minimum password complexity, a password profile, and use an authentication profile to use an external authentication service to validate the administrator's credentials.

The following example shows how to create a local administrator account with local authentication:

CREATE A LOCAL ADMINISTRATOR ACCOUNT	
<p><b>Step 1.</b> Create an Admin Role profile.</p> <p>This step is only required if using custom roles instead of using the built-in Dynamic Roles available on Panorama.</p>	<p>Complete the following steps for each role you want to create:</p> <ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Admin Roles</b> and then click <b>Add</b>.</li> <li>2. Select <b>Panorama</b> or <b>Device Group and Template</b> to define the scope of administrative privileges to assign. The access privileges defined for Panorama are enforced when the administrator logs in to Panorama; the Device Group and Template role enforces read-only access to the Managed Devices, Templates, and Device Groups nodes on the <b>Panorama</b> tab. Access to all other tabs can be modified as required.</li> <li>3. For the <b>Web UI</b> and /or <b>XML API</b> tabs, set the access levels—Enable , Read Only , Disable —for each functional area of the interface by clicking the icon to toggle it to the desired setting: <ul style="list-style-type: none"> <li>• For Panorama access, define access to the <b>Web UI</b>, <b>XML API</b>, and <b>Command Line</b>. The <b>Command Line</b> tab, does not allow granular access. You must select from the predefined options: <b>superuser</b>, <b>superreader</b>, <b>Panorama-admin</b> or <b>None</b>.</li> <li>• For access to firewalls (Device Group and Template), only one tab is available: <b>Web UI</b>. From Panorama, you cannot enable access to the CLI or XML API on a device because there are no predefined roles that restrict access. Therefore, to prevent privilege-level escalation, the ability to manage access to the CLI and XML API is not available from Panorama.</li> </ul> </li> <li>4. Enter a <b>Name</b> for the profile and then click <b>OK</b> to save it.</li> </ol>
<p><b>Step 2</b> (Optional) Set requirements for local user-defined passwords.</p>	<ul style="list-style-type: none"> <li>• <b>Create Password Profiles</b>—Define how often administrators must change their passwords. Create multiple password profiles and apply them to administrator accounts as required to enforce security. To create a password profile, select <b>Panorama &gt; Password Profiles</b> and then click <b>Add</b>.</li> <li>• <b>Configure minimum password complexity settings</b>—Define rules that govern password complexity, which forces administrators to create passwords that are harder to guess, crack, or compromise. Unlike password profiles, which can be applied to individual accounts, these rules are device-wide and apply to all passwords. To configure the settings, select <b>Panorama &gt; Setup</b> and then click the Edit  icon in the Minimum Password Complexity section.</li> </ul>

CREATE A LOCAL ADMINISTRATOR ACCOUNT (CONTINUED)	
<b>Step 3</b> Create an account for each administrator.	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Administrators</b> and then click <b>Add</b>.</li> <li>2. Enter a user <b>Name</b> and <b>Password</b> for the administrator.</li> <li>3. Select the <b>Role</b> to assign to this administrator. Select a predefined Dynamic role or a custom role-based profile as defined in Step 1.</li> <li>4. (Optional) Select the <b>Authentication Profile</b> to use for validating an administrative user's credentials to an external authentication server. See <a href="#">Create an Authentication Profile</a>.</li> <li>5. (Optional) Select a <b>Password Profile</b>. See Step 2.</li> <li>6. Click <b>OK</b> to save the account.</li> </ol>
<b>Step 4</b> Save the configuration changes.	<ol style="list-style-type: none"> <li>7. Click <b>Commit</b>, and select <b>Panorama</b> in the <b>Commit Type</b> option.</li> </ol>

## Define Access Domains

An *access domain* provides a way to limit administrative access to specified device groups (to manage policies and objects) and templates (to manage network and device settings), and the ability to switch context to the web interface on the managed devices. Access domain settings are only relevant if:

- A Custom Admin Role Profile with a **Device Group and Template** role is defined.
- A RADIUS server is used for administrator authentication. The access domain is linked to RADIUS vendor-specific attributes (VSAs). On the RADIUS server, a VSA attribute number and value is defined for each administrative user. The value defined must match the access domain configured on Panorama. When an administrator attempts to log in Panorama, Panorama queries the RADIUS server for the administrator's access domain and attribute number. Based on the response from the RADIUS server, the administrator is authorized for access and is restricted to the devices/virtual systems, device groups and templates specified in the access domain. For details on the supported RADIUS VSAs, see [Using RADIUS Vendor Specific Attributes \(VSAs\)](#).

DEFINE AN ACCESS DOMAIN	
<b>Step 1.</b> Create an access domain.	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Access Domain</b> and then click <b>Add</b>.</li> <li>2. Enter a user <b>Name</b> to identify the domain.</li> </ol>
<b>Step 2</b> Specify the device groups, templates and device contexts that the user can administer.	In the <b>Device Groups</b> , <b>Templates</b> , and <b>Device Context</b> tabs, click <b>Add</b> and pick from the filtered list or drop-down that displays.
<b>Step 3</b> Save the configuration changes.	Click <b>Commit</b> , and select <b>Panorama</b> in the <b>Commit Type</b> option.

Create an Authentication Profile

An authentication profile specifies the authentication service that validates the administrator’s credentials and defines how to access that authentication service. Panorama can be configured to access the local database, a RADIUS server, Kerberos server, or an LDAP server.

If you are using an external authentication server, create a server profile (**Panorama > Server Profiles**) before creating an authentication profile. Panorama requires the server profile to access the authentication service.

CREATE AN AUTHENTICATION PROFILE

Step 1

Create an authentication profile.

1.

Select **Panorama > Authentication Profile** and then click **Add**.

2.

Enter a user **Name** to identify the authentication profile.

Step 2

Define the conditions for locking out the administrative user.

1.

Enter the **Lockout Time**. This is the number of minutes that a user is locked out upon reaching the maximum number of failed attempts (0-60 minutes; default 0). 0 means that the lockout is in effect until it is manually unlocked.

2.

Enter the **Failed Attempts** count. This is the number of failed login attempts that are allowed before the account is locked out (1-10; default 0). By default, the failed attempt count is 0 and the user is not locked out despite repeated failure to authenticate.

Step 3

Specify the users and groups that are explicitly allowed to authenticate.

By adding an allow list to an authentication profile, you can limit access to specific users in a user group/directory.

For the **Allow List**, pick one of the following:

- Select the **All** check box to allow all users.
- Click **Add** and enter the first few characters of a name in the field to list all the users and user groups that start with those characters. Repeat to add as many users/user groups as required.

Step 4


Select the authentication service and attach the server profile.

1.

In the **Authentication** drop-down, select the type of authentication you will use.

2.

Select the appropriate server profile in the **Server Profile** drop-down.

Lockout						
Failed Attempts (#)	Lockout Time (min)	Allow List	Authenticat...	Server Profile	Others	Locked Users
default	default	 all	RADIUS	MJS-RADIUS		none

Step 5

Commit your changes.

Click **Commit**, and select **Panorama** in the **Commit Type** option.

## Define an Authentication Sequence

An authentication sequence is an ordered list of authentication profiles that allows the use of more than one authentication service. Authentication sequences provide flexibility in environments where multiple databases exist for different users and user groups. When defining an authentication sequence, Panorama attempts to authenticate the administrator using each of the configured server profiles in sequence. For example, an authentication sequence can instruct Panorama to check LDAP first, RADIUS next, and the local database last, until a successful authentication occurs; if it fails, the administrator is denied access.

DEFINE AN AUTHENTICATION SEQUENCE	
Step 1. Create an authentication sequence.	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Authentication Sequence</b> and then click <b>Add</b>.</li> <li>2. Enter a user <b>Name</b> to identify the authentication sequence.</li> <li>3. Click <b>Add</b> to select the chronological sequence of authentication profiles against which the administrator's credentials must be checked.</li> </ol>
Step 2 (Optional) Define the conditions for locking out the administrative user.	<ol style="list-style-type: none"> <li>1. Enter the <b>Lockout Time</b>. This is the number of minutes that a user is locked out upon reaching the maximum number of failed attempts (0-60 minutes; default 0). 0 means that the lockout is in effect until it is manually unlocked.</li> <li>2. Enter the <b>Failed Attempts</b> count. This is the number of failed login attempts that are allowed before the account is locked out (1-10; default 0). By default, the failed attempt count is 0 and the user is not locked out despite repeated failure to authenticate.</li> </ol>
Step 3 Save the configuration changes.	Click <b>Commit</b> , and select <b>Panorama</b> in the <b>Commit Type</b> option.

## Configure Administrative Authentication

Administrators can authenticate locally to Panorama using passwords or certificates, or they can authenticate to an external authentication server.

There are three options for setting up administrative authentication on Panorama:

- ▲ Create a local user account and authenticate locally using password authentication, certificate-based, or key-based authentication. See [Create an Administrative Account](#); [Enable Certificate-Based Authentication for the Web Interface](#) and [Enable SSH Key-Based Authentication for the Command Line Interface](#).
- ▲ Create a local user account but authenticate to an external RADIUS/LDAP/Kerberos server using authentication profiles:
  - Create a server profile on the **Panorama > Server Profile** tab. A server profile is required for each external service with which Panorama must interact. The server details required for to establish the connection with Panorama varies by the authentication service you plan to use.
  - Create an authentication profile. See [Create an Authentication Profile](#).

- (Role-based access only) Define an Admin Role Profile that specifies whether the user has access to Panorama or Device Groups and Templates; see [Create an Admin Role profile](#). For dynamic roles, an Admin Role Profile is not required.
- ▲ Use RADIUS Vendor Specific Attributes (VSAs) for managing administrative access to Panorama. Use this option if you do not want to create a local account on Panorama for an administrative user, and would like to use your current infrastructure to manage authentication and password management on a RADIUS server. For a high-level overview of the process, see [Using RADIUS Vendor Specific Attributes \(VSAs\)](#).

Enable Certificate-Based Authentication for the Web Interface

As a more secure alternative to using a password to authenticate a user, enable certificate-based authentication for securing access to Panorama. With certificate-based authentication a digital signature is exchanged and verified, in lieu of a password.



To enable certificate-based authentication, you must configure Panorama to use a client certificate profile (see [Step 4](#) and [Step 5](#)). When you enable a client certificate profile, each administrator must use a client certificate for access to Panorama.

Use the following instructions to enable certificate-based authentication. This example uses a CA certificate generated on Panorama.

ENABLE CERTIFICATE-BASED AUTHENTICATION	
<p><b>Step 1.</b> Generate a CA certificate on Panorama.</p> <p><b>Note</b> To use a certificate from a trusted third-party or enterprise CA, you must import that CA certificate in to Panorama.</p>	<p>To generate a CA certificate on Panorama:</p> <ol style="list-style-type: none"><li>1. Log in to the Panorama web interface.</li><li>2. Select <b>Panorama &gt; Certificate Management &gt; Certificates</b> and click <b>Generate</b>.</li><li>3. Enter a <b>Certificate Name</b>. Add the IP address or FQDN of Panorama for listing in the <b>Common Name</b> field of the certificate. Optionally, you can change the cryptographic settings, and define certificate options such as country, organization, or state.</li><li>4. Make sure to leave the <b>Signed By</b> option blank and select the <b>Certificate Authority</b> option.</li><li>5. Click <b>Generate</b> to create the certificate using the details you specified above.</li></ol>

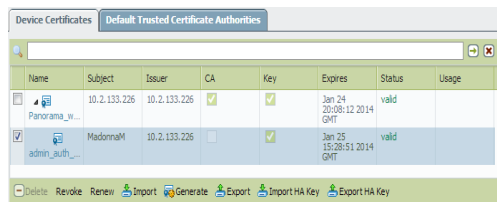
**ENABLE CERTIFICATE-BASED AUTHENTICATION (CONTINUED)**

**Step 2** Create and export the client certificate that will be used to authenticate an administrator.



The 'Generate Certificate' dialog box contains the following fields and options:

- Certificate Name:** Admin\_auth\_cert
- Common Name:** MadonnaS
- IP or FQDN to appear on the certificate:** (empty)
- Signed By:** Panorama\_web\_access (selected from a dropdown menu)
- Certificate Authority:** ☐



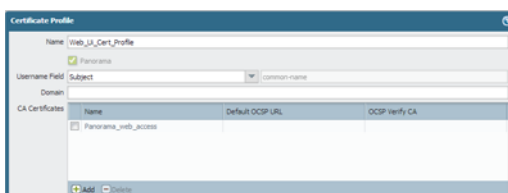
Name	Subject	Issuer	CA	Key	Expires	Status	Usage
Panorama_w...	10.2.133.226	10.2.133.226	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Jan 24 20:08:12 2014 GMT	valid	
admin_auth...	MadonnaM	10.2.133.226	<input type="checkbox"/>	<input checked="" type="checkbox"/>	Jan 25 13:20:51 2014 GMT	valid	

1. Use the CA certificate to generate a client certificate for the specified administrative user.
  - a. Select **Panorama > Certificate Management > Certificates** and click **Generate**.
  - b. In the **Common Name** field, enter the name of the administrator for whom you are generating the certificate. The name syntax must match the format used by the local or external authentication mechanism.
  - c. In the **Signed by** field, select the same CA certificate that you created in Step 1.
  - d. Click **Generate** to create the certificate.
2. Export the client certificate you just generated.
  - a. Select the certificate that you just generated and click **Export**.
  - b. To encrypt the private key, select **PKCS12** as the **File Format**.
  - c. Enter a passphrase to encrypt the private key and confirm the entry.
  - d. Click **OK** to export the certificate.

**Step 3** Create or modify an administrator account to enable client certificate authentication on the account.

1. Select **Panorama > Administrators** and then click **Add**.
2. Enter a login name for the administrator; the name is case-sensitive.
3. Select **Use only client certificate authentication (Web)** to enable the use of the certificate for authentication.
4. Select the **Role** to assign to this administrator. You can either select one of the predefined dynamic roles or select a custom role and attach an authentication profile that specifies the access privileges for this administrator.
5. (Optional) For custom roles, select the device groups, templates and the device context that the administrative user can modify.
6. Click **OK** to save the account settings.

**Step 4** Create the Client Certificate Profile that will be used for securing access to the web interface.



The 'Certificate Profile' dialog box contains the following fields and options:

- Name:** Web\_UA\_Cert\_Profile
- Username Field:** Subject (selected from a dropdown menu)
- Domain:** (empty)
- CA Certificates:**
  - ☒ Panorama
  - ☐ Panorama\_web\_access

1. Select **Panorama > Certificate Management > Certificate Profile** and click **Add**.
2. Enter a name for the certificate profile and in the **Username Field** select **Subject**.
3. Select **Add** in the CA Certificates section and from the **CA Certificate** drop-down, select the CA certificate you created in Step 1.

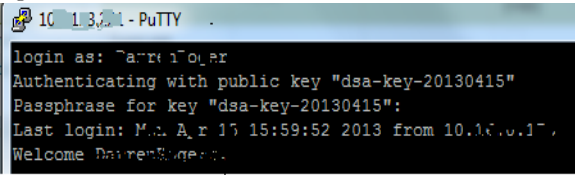
ENABLE CERTIFICATE-BASED AUTHENTICATION (CONTINUED)	
<b>Step 5</b> Configure Panorama to use the client certificate profile for authentication.	<ol style="list-style-type: none"> <li>1. On the <b>Panorama &gt; Setup</b> tab, click the Edit icon in the Authentication Settings section of the screen.</li> <li>2. In the <b>Certificate Profile</b> field, select the client certificate profile created in <a href="#">Step 4</a>.</li> <li>3. Click <b>OK</b> to save your changes.</li> </ol>
<b>Step 6</b> Save the configuration changes.	Click <b>Commit</b> and select <b>Panorama</b> as the <b>Commit Type</b> . You will be logged out of the device.
<b>Step 7</b> Import the administrator's client certificate into the web browser on the client system that the administrator will use to access the Panorama web interface.	For example, in Firefox: <ol style="list-style-type: none"> <li>1. Select <b>Tools &gt; Options &gt; Advanced</b>.</li> <li>2. Click <b>View Certificates</b>.</li> <li>3. Select the <b>Your Certificates</b> tab and click <b>Import</b>. Browse to the location where you saved the client certificate.</li> <li>4. When prompted, enter the passphrase to decrypt the private key.</li> </ol>
<b>Step 8</b> Verify that certificate-based authentication is configured.	<ol style="list-style-type: none"> <li>1. From a client system that has the client certificate loaded, access the Panorama IP address or hostname.</li> <li>2. When prompted, select the client certificate you imported in Step 7. A certificate warning will display.</li> <li>3. Add the certificate to the exception list and log in to the Panorama web interface.</li> </ol>

## Enable SSH Key-Based Authentication for the Command Line Interface

To enable SSH key-based authentication, complete the following workflow for every administrative user:

ENABLE SSH (PUBLIC KEY BASED) AUTHENTICATION	
<b>Step 1.</b> Use an SSH key generation tool to create an asymmetric keypair on the client machine.  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 required to generate the keypair, refer to the product documentation for your SSH client.  The public key and private key are two separate files; save both to a location that can be accessed by Panorama. For added security, enter a passphrase to encrypt the private key. The administrator will be prompted for this passphrase when logging in to Panorama.



ENABLE SSH (PUBLIC KEY BASED) AUTHENTICATION (CONTINUED)	
<p><b>Step 2</b> Create an account for the administrator and enable certificate-based authentication.</p>	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Administrators</b> and then click <b>Add</b>.</li> <li>2. Enter a user <b>Name</b> and <b>Password</b> for the administrator. Make sure to enter a strong/complex password and record it in safe location; Panorama will only prompt for this password in the event that the certificates are corrupted or a system failure occurs.</li> <li>3. (Optional) Select an <b>Authentication Profile</b>.</li> <li>4. Enable <b>Use Public Key Authentication (SSH)</b>.</li> <li>5. Click <b>Import Key</b> and browse to import the public key created in <a href="#">Step 1</a>.</li> <li>6. Select the <b>Role</b> to assign to this administrator. You can either select one of the predefined Dynamic roles or a custom Role-Based profile. For details, see <a href="#">Create an Admin Role profile</a>.</li> <li>7. Click <b>OK</b> to save the account.</li> </ol>
<p><b>Step 3</b> Save the configuration changes.</p>	<ol style="list-style-type: none"> <li>8. Click <b>Commit</b>, and select <b>Panorama</b> as the <b>Commit Type</b> option.</li> </ol>
<p><b>Step 4</b> Verify that the SSH client uses the private key to authenticate the public key presented by Panorama.</p>	<ol style="list-style-type: none"> <li>1. Configure the SSH client to use the private key to authenticate to Panorama.</li> <li>2. Log in to the CLI on Panorama.            </li> <li>3. If prompted, enter the passphrase you defined when creating the keys in <a href="#">Step 1</a>.</li> </ol>


## Using RADIUS Vendor Specific Attributes (VSAs)

To use RADIUS VSAs you must complete the following tasks:

- On Panorama:
  - Configure a RADIUS server profile (**Panorama > Server Profiles > RADIUS**).
  - Create an authentication profile that specifies RADIUS as the protocol for authentication and attach the RADIUS server profile (**Panorama > Authentication Profiles**).
  - Create an custom administrative role profile with a Device Group and Template role (**Panorama > Admin Roles**).
  - Configure Panorama to use the authentication profile for authentication (**Setup > Management > Authentication Settings > Authentication Profile**).
  - (Required only if using the VSA: PaloAlto-Panorama-Admin-Access-Domain) If you want to restrict administrative access to specific managed devices, Templates, and/or Device Groups, define an access domain (**Panorama > Access Domains**).
- On the RADIUS server:
  - Add the Panorama IP address or hostname as the RADIUS client.
  - Define the VSAs supported by Panorama. To define an attribute, use the vendor code (25461), attribute name (make sure that it matches the name of the admin role profile/access domain defined on Panorama; it is case sensitive), number, format (string):
    - PaloAlto-Panorama-Admin-Role, attribute #3
    - PaloAlto-Panorama-Admin-Access-Domain, attribute #4

For detailed instructions on setting up authentication using RADIUS VSAs, refer to the following documents:

- On Windows 2003 Server and Cisco ACS 4.0: <https://live.paloaltonetworks.com/docs/DOC-1765>
- On Cisco ACS 5.2: <https://live.paloaltonetworks.com/docs/DOC-1979>



## 3 Manage Firewalls and Log Collection

---

Panorama provides two main functions: it centralizes the process of administering Palo Alto Networks firewalls and provides visibility into network traffic through aggregated reporting. In order to administer the devices and generate reports on network traffic, you must add the firewalls to Panorama as managed devices and configure the firewalls to forward logs to Panorama or a Log Collector. This section includes the following topics:

- ▲ [Manage Your Firewalls](#)
- ▲ [Enable Logging](#)
- ▲ [Deploy Software Updates and Manage Licenses](#)
- ▲ [Replace a Managed Device with a New Device](#), to replace a Return Merchandise Authorization (RMA) device
- ▲ [Transition a Device to Central Management](#)

# Manage Your Firewalls

Because Panorama is designed to be a focal point for firewall administration, the workflow in this document is best suited for first-time firewall deployment. Review [Plan Your Deployment](#) and then continue with the following sections:

- ▲ [Add Managed Devices](#)
- ▲ [Create Device Groups](#)
- ▲ [Create Templates](#)
- ▲ [Configure the Firewalls to Forward Logs to Panorama](#)
- ▲ [Commit Changes on Panorama](#)
- ▲ [Modify the Log Forwarding and Buffering Defaults](#)
- ▲ [Use Panorama to Configure Managed Devices: An Example](#)



To view the **Objects**, **Policies** tabs on the Panorama web interface, you must first create at least one Device Group and at least one Template for the **Network** and **Device** tabs to display. These tabs include the configuration options required to configure and manage the firewalls on your network.

If you have already configured and deployed firewalls on your network, the process of migrating the configuration, local policies and objects from the firewalls to a centralized management approach requires an understanding of scripting and the use of the REST API on the firewalls. To make this transition efficient, Palo Alto Networks recommends using trained and certified partners who are familiar with the planning, implementation, and verification stages of the migration process. Contact your authorized reseller or partner for more information on the support offerings that are available to you. For a brief overview of the process, see [Transition a Device to Central Management](#) and for more details refer to this article: [Panorama Device Migration Tech Note](#).




## Add Managed Devices

Adding firewalls to Panorama is the first step in centrally managing them using Panorama. Before you begin, collect the serial numbers for all the devices you want to manage using Panorama.

To manage a firewall using Panorama, prepare each firewall as follows:

- Perform initial configuration on the firewall so that the device is accessible and can communicate with Panorama over the network.
- Add the Panorama IP address(es) (one server or two, if Panorama is configured in a high availability pair) in the Panorama Settings section of the **Device > Setup > Management** tab and commit the changes.
- Set up the data interfaces. For each interface you plan to use, select the interface type and attach it to a security zone so that you can push configuration and policy from Panorama.

You can then add the firewalls as managed devices on Panorama as follows:

ADD DEVICES TO MANAGE	
<p>Step 1. Add device(s) to Panorama.</p>	<ol style="list-style-type: none"><li>1. Select <b>Panorama &gt; Managed Devices</b>.</li><li>2. Click <b>Add</b> and enter the serial number for each device that you want to manage centrally using Panorama. Add only one entry per line.</li><li>3. Click <b>OK</b>. The newly added devices will display in the Managed Devices pane.</li><li>4. (Optional) Add a <b>Tag</b>. Tags make it easier for you to find a device from a large list; they help you to dynamically filter and refine the list of firewalls that display. For example, if you add a tag called branch office, you can filter for all branch office devices across your network.<ol style="list-style-type: none"><li>a. Select the check box next to the managed device.</li><li>b. Click <b>Tag</b> . Click <b>Add</b> and enter a text string of up to 31 characters. Do not use an empty space.</li><li>c. Click <b>OK</b>.</li></ol></li><li>5. Click <b>Commit</b>, and select <b>Panorama</b> as the <b>Commit Type</b>.</li></ol>
<p>Step 2. Verify that the device is connected to Panorama.</p>	<p>If the firewall is accessible on the network and the Panorama IP address is configured on the device, Panorama must be able to connect  to the device.</p> 

## Create Device Groups

After you add the devices, you can group the devices in to *device groups*. A device group can include one or more firewalls or virtual systems that need similar policies and objects and can therefore be effectively managed as a logical unit.

When managing firewalls that are configured in an active-passive high availability (HA) configuration, make sure to place both devices in the same device group in Panorama. This is essential to make sure that the same policies and objects are pushed to both devices in the HA pair. Panorama pushed policies are not synchronized between firewall HA peers.

CREATE DEVICE GROUPS	
<p><b>Step 1.</b> Create Device Group(s).</p> <p><b>Note</b> A device can only belong to one Device Group; for devices with multiple virtual systems, each virtual system can belong to a different Device Group.</p>	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Device Groups</b>, and click <b>Add</b>.</li> <li>2. Enter a unique <b>Name</b> and a <b>Description</b> to identify the device group.</li> <li>3. Use the filters to select the devices that you would like to add to the group.</li> <li>4. (Optional) Select the <b>Group HA Peers</b> check box for firewalls that are set up as an HA pair. Adding both devices or virtual systems to the same device group allows you to push shared policies and objects simultaneously to both peers.</li> </ol> <p><b>Note</b> To group HA peers, the devices must be running PAN-OS 5.0 or later.</p> <ol style="list-style-type: none"> <li>5. (Optional) If you plan to use user or groups in policy, you must assign a <b>Master</b> device for the device group. The master device is the firewall from which Panorama gathers username and user group information for use in policies.</li> <li>6. Click <b>OK</b>.</li> <li>7. Click <b>Commit</b>, and select <b>Panorama</b> as the <b>Commit Type</b>. Save the changes to the running configuration on Panorama.</li> <li>8. Click <b>Commit</b>, and select <b>Device Group</b> as the <b>Commit Type</b>. Push the changes to the devices in the device group.</li> </ol>
<p><b>Step 2</b> Begin centrally administering policies on the devices in the device group(s).</p>	<ul style="list-style-type: none"> <li>• <a href="#">Create Objects for Use in Shared or Device Group Policy</a></li> <li>• <a href="#">Manage Shared Objects</a></li> <li>• <a href="#">Target Policies to a Subset of Devices</a></li> <li>• <a href="#">View Rule Hierarchy and Find Unused Rules</a></li> </ul> <p>For an example, see <a href="#">Use Panorama to Configure Managed Devices: An Example</a></p>

## Create Objects for Use in Shared or Device Group Policy

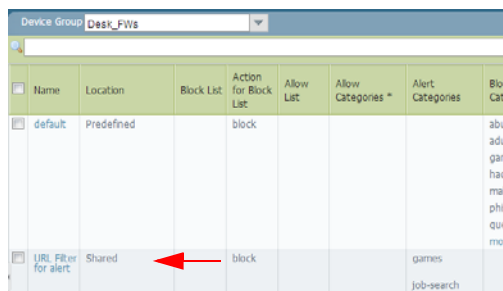
An *Object* is a container for grouping discrete identities such as IP addresses, URLs, applications, or users, for use in policy enforcement. You can use Panorama to create and clone all objects in the **Objects** tab such as **Address/Address Group, Region or User/User Group**. These policy objects can be shared across all managed devices or be specific to a device group.

- A *shared object* is a reusable component that is created on Panorama. It is shared across all device groups and can be referenced in shared policies or in device group policies. It reduces administrative overhead and ensures consistency in configuring multiple firewalls.
- A *device group* object is specific to the device group in which it is defined. It can be used only in the device group where it is created and is not visible when configuring other device groups or shared rules and objects. For example, a device group object for a set of web server IP addresses that is created in the datacenter device group is not available for use in any other device group or for use in shared policies.

### CREATE OBJECTS

Create a shared object.

In this example, we will add a shared object for URL Filtering categories for which we want to be trigger an alert.



Name	Location	Block List	Action for Block List	Allow List	Allow Categories *	Alert Categories	Block Categories
default	Predefined		block				abuse, adult, gaming, hacking, malware, phishing, quiet, monitoring
URL Filter for alert	Shared		block			games, job-search	

1. Select the **Objects > Security Profiles > URL Filtering** tab and click **Add**.

If the **Objects** tab does not display, see [Add Managed Devices](#) to add a device group. The Panorama web interface displays the **Objects** tab only if you have created a device group.

2. Enter a **Name** and a **Description**.
3. Select the **Shared** check box. If you do not select the checkbox, the object will be a part of the device group that currently displays in the **Device Group** drop-down.
4. Select the check box next to the **URL Categories** for which you want to be notified and select **Alert** in the **Action** column, then click **OK**.
5. Click **Commit**, and select **Panorama** as the **Commit Type**.

## CREATE OBJECTS

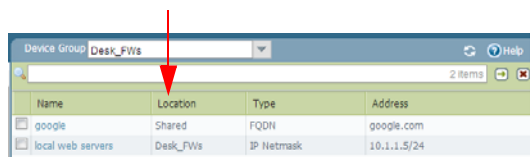
Create a device group object.

In this example, we will add a device group object for specific web servers on your network.

1. Select the Device Group for which you plan to use this object in the **Device Group** drop-down.
2. Select the **Objects > Addresses** tab.
3. Select **Address** and click **Add**.
4. Verify that the **Shared** check box is not selected.
5. Enter a **Name**, a **Description**, and select the **Type** of address object from the drop-down. For example, select **IP Range** and include the IP address range for the web servers for which you would like to create an address object.
6. Click **OK**.
7. Commit your changes.
  - a. Click **Commit**, and select **Panorama** as the **Commit Type**. This saves the changes to the running configuration on Panorama.
  - b. Click **Commit**, and select **Device Group** as the **Commit Type**. This pushes the changes to the devices included in the Device Group.

View shared objects and device group objects in Panorama.

To demonstrate the difference between a shared object and a device group object, the following screenshot includes a shared address object that was created on Panorama.



Name	Location	Type	Address
google	Shared	FQDN	google.com
local web servers	Desk_FW's	IP Netmask	10.1.1.5/24

The **Location** column in the **Objects** tab displays whether an object is shared or is specific to a device group.

1. Select the device group, for which you just created a device group object, in the **Device Group** drop-down.
2. Select the **Objects > Addresses** tab and verify that the device group object displays; note that the device group name in the Location column matches the selection in the **Device Group** drop-down.

**Note** If a different device group is selected in the **Device Group** drop-down, only the device group objects (and shared objects) created for the selected device group will display.

## Manage Shared Objects

You can configure how Panorama handles shared objects. Consider whether you:

- Would like to configure Panorama to only push shared objects that are referenced either in shared policies or device group policies to the managed device. Say for example, all objects in your deployment are defined as shared objects, but you would like to push only the relevant objects for each device group. The **Share Unused Address and Service Objects** checkbox allows you to limit the objects that are pushed to the managed devices.

By default, Panorama pushes all shared objects (used and unused) to the managed devices. On lower-end platforms, such as the PA-200, consider pushing only the relevant shared objects to the managed devices. This is because the number of objects that can be stored on the lower-end platforms is considerably lower




than that of the mid- to high-end platforms. Also, if you have many address and service objects that are unused, clearing the **Share Unused Address and Service Objects** check box reduces the commit times significantly on the devices because the configuration pushed to each device is smaller.

Disabling this option may, however, increase the commit time on Panorama. This is because Panorama has to dynamically check whether a particular object is referenced in policy.

### UNUSED SHARED OBJECTS

- Disable the sharing of unused address and service objects to devices.



1. Select **Panorama > Setup > Management**, and click the Edit button in the Panorama Settings section.
2. Clear the **Share Unused Address and Service Objects with Devices** check box.

- Would like to ensure that a shared object takes precedence over an object that has the same name as a device group object.

By default, shared objects do not override any device group object with the same name as a shared object.

If you would like to prevent overrides to objects that have been defined as shared objects on Panorama, you can enable the option for **Shared Objects Take Precedence**. When enabled, all device group objects with the same name will be discarded and the shared object settings will be pushed to the managed devices.

### PRECEDENCE OF SHARED OBJECTS

- Ensure that shared objects always take priority over device group objects.

1. Select **Panorama > Setup > Management**, and click the Edit button in the Panorama Settings section.
2. Select the **Shared Objects Take Precedence** check box.

## Target Policies to a Subset of Devices

A policy *target* allows you to specify the devices in a Device Group to which to push policy. It allows you to exclude one or more devices or virtual systems, or to only apply the rule to specific devices or virtual systems in a Device Group.

The ability to target a policy enables you to keep policies centralized on Panorama; it offers visibility and efficiency in managing the rules. Instead of creating local rules on a device or virtual system, targeted policy rules allow you to define the rules (as shared or device-group pre- or post-rules) on Panorama.

TARGET A POLICY	
<p><b>Step 1.</b> Create a policy.</p>	<ol style="list-style-type: none"> <li>1. Select the <b>Device Group</b> for which you want to define policy.</li> <li>2. Select the <b>Policies</b> tab, and select the rulebase for which you would like to create policy. For example, define a pre-rule in the Security policies rulebase that permits users on the internal network to access the servers in the DMZ: <ol style="list-style-type: none"> <li>a. Click <b>Add</b> in <b>Policies &gt; Security &gt; Pre-Rules</b>.</li> <li>b. Give the rule a descriptive name in the <b>General</b> tab.</li> <li>c. In the <b>Source</b> tab, set the <b>Source Zone</b> to Trust.</li> <li>d. In the <b>Destination</b> tab, set the <b>Destination Zone</b> to DMZ.</li> <li>e. In the <b>Service/ URL Category</b> tab, set the <b>Service</b> to <b>application-default</b>.</li> <li>f. In the <b>Actions</b> tab, set the <b>Action Setting</b> to <b>Allow</b>.</li> <li>g. Leave all the other options at the default values.</li> </ol> </li> </ol>
<p><b>Step 2</b> Target the policy to include or exclude a subset of devices.</p>	<p>To apply the policy to a selected set of devices.</p> <ol style="list-style-type: none"> <li>1. Select the <b>Target</b> tab in the Policy Rule window.</li> <li>2. Select the devices on which you would like the rule to apply.</li> </ol> <p><b>Note</b> By default, although the check box for the virtual systems in the Device Group is unchecked, all the virtual systems will inherit the rule on commit. Select the check box for one or more virtual systems to which you want the rule to apply.</p> <ol style="list-style-type: none"> <li>3. (Optional) To exclude a subset of devices from inheriting the policy rule, select the check box <b>Install on all but specified devices</b>.</li> <li>4. Click <b>OK</b>.</li> <li>5. Save the configuration changes. <ol style="list-style-type: none"> <li>a. Click <b>Commit</b>, and select <b>Panorama</b> as the <b>Commit Type</b> to save the changes to the running configuration on Panorama.</li> <li>b. Click <b>Commit</b>, and select <b>Device Group</b> as the <b>Commit Type</b> to push the changes to the devices selected in the Device Group.</li> </ol> </li> </ol>

## View Rule Hierarchy and Find Unused Rules

The ordering of rules is essential for securing your network. Policy rules are evaluated from top to bottom. A packet is matched against the first rule that meets the defined criteria; after a match is triggered the subsequent rules are not evaluated. Therefore, the more specific rules must precede more generic ones in order to enforce the best match criteria.

Use the following procedure to verify the ordering of rules and make changes as appropriate:

### PREVIEW RULES AND VIEW UNUSED RULES

**Step 1.** View the rule hierarchy for each rulebase.

1. Select the **Policies** tab, and click **Preview Rules**.

2. Use the following filters for previewing rules:

- **Rulebase:** Select a rulebase and view the rules defined for that rulebase—Security, NAT, QoS, Policy Based Forwarding, Decryption, Captive Portal, Application Override, or DoS Protection.
- **Device Group:** For the selected rulebase, you can view all **Shared** policies or select a specific **Device Group** for which you want to view the combined list of policies inherited from Panorama and those defined locally.
- **Device:** For the selected rulebase and Device Group, you can view the list of policies that will be evaluated on a specific device in the Device Group.

Combined Rules Preview							
Rulebase: Security		Device Group: PK_Branch Office		Device: PK-PA-200/ntsys1			
	Source			Destination			
Name	Zone	Address	User	Zone	Address	Application	Service
Pre-PK-Sec-Rule	any	any	any	any	any	any	any
General Security Policy	L3-Trusted	any	any	L3-Untrusted	any	PK-Safe-Apps	any
VPN	L3-Untrusted	172.16.0.0/16	any	L3-Trusted	any	any	any
Post-PK-Sec-Rule	any	any	any	any	any	any	any

All the Shared and Device Group rules that are inherited from Panorama are displayed in green, and the local rules on the device are displayed in blue; the local rules are encased between the pre-rules and post-rules.

3. Close the Combined Rules window to exit the preview mode.

**Step 2** Find unused rules, and optionally delete or disable the rules. Each device maintains a flag for the rules that have a match. Panorama monitors each device, fetches, and aggregates the list of rules that do not have a match. Because the flag is reset when a dataplane reset occurs on a reboot or a restart, as a best practice, monitor this list periodically to determine whether the rule has had a match since the last check before you delete or disable it.

1. Select the **Policies** tab, and click **Highlight Unused Rules**.

The rules are not currently used display with a dotted yellow background.

2. (Optional) To delete an unused rule, select the rule and click **Delete**.

3. (Optional) To disable a rule, select the rule and click **Disable**. The disabled rule displays in an italicized font.

**Step 3** Rearrange the rules within a selected pre-rule or post-rule rulebase, if required.

1. In a rulebase, select the rule you want to move.

2. Click the **Move Up**, **Move Down**, **Move Top** or **Move Bottom** options to reorder the placement of the rule.

**Note** To rearrange local rules on the device, switch to the local device context.

---

**PREVIEW RULES AND VIEW UNUSED RULES (CONTINUED)**

---

**Step 4** If you modified the rules, save the changes.

1. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.
  2. Click **Commit**, and select **Device Group** as the **Commit Type** to push the changes to the devices selected in the Device Group.
- 

## Create Templates

Panorama Templates allow you manage the configuration options on the **Device** and **Network** tabs on the managed firewalls. Using templates you can define a base configuration for centrally staging new firewalls and then make device-specific exceptions in configuration, if required. You can for example, use templates to define administrative access to the device, set up User-ID, manage certificates, set up the firewalls in a high availability pair, define log settings and server profiles on the managed firewalls.

When creating templates, make sure to assign similar devices to a template. For example, group devices with a single virtual system in a one template and devices enabled for multiple virtual systems in another template, or group devices that require very similar network interface and zone configuration in a template.

See the following sections for information on working with templates:

- [What can Templates not be Used for?](#)
- [Add a New Template](#)
- [Override Template Settings](#)
- [Disable Template Settings](#)

### What can Templates not be Used for?

For firewalls running PAN-OS 4.x, the use of Panorama templates is limited to the following:

- Creating response pages
- Defining authentication profiles and sequences
- Creating self-signed certificates on Panorama or importing certificates
- Creating client authentication certificates (known as Certificate Profiles in Panorama 5.0 and later)
- Creating server profiles: SNMP Trap, Syslog, Email, NetFlow, RADIUS, LDAP, and Kerberos

For firewalls running PAN-OS 5.x and later, templates allow you configure a wide array of settings, with the following exceptions:

- Cannot enable operational modes such as multi-vsyt mode, FIPS mode, or CC mode using templates; these operational settings must be configured locally on each device.

- Cannot configure the IP address details for the firewall HA pair.  
The HA1 peer IP address; HA1 backup peer IP address; HA2 peer IP address; and HA2 backup peer IP address must be configured locally on each managed device.
- Cannot configure a master key and diagnostics.


## Add a New Template

Until you add a template on Panorama, the **Device** and **Network** tabs required to define the network set up elements and device configuration elements on the firewall will not display. Use these instructions to add a new template.

ADD A TEMPLATE	
<p><b>Step 1.</b> Add a new Template.</p>	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Templates</b>.</li> <li>2. Click <b>Add</b> and enter a unique name and a description to identify the template.</li> <li>3. (Optional) Select the <b>Virtual Systems</b> check box if this template will be used for devices that are multi-vsyz capable and are enabled for multi-vsyz functionality.</li> </ol> <p><b>Note</b> A commit failure will occur if a template enabled for devices with multi-vsyz capability is pushed to devices that are not multi-vsyz capable or are not enabled for the multi-vsyz functionality.</p> <ol style="list-style-type: none"> <li>4. Specify the <b>Operational Mode</b> for the devices to which the template will be applied. The default is <b>normal</b>; change to <b>cc</b> or <b>fips</b>, as required. The template commit will fail if there is a mismatch in the operational mode specified on the template with what is enabled on the devices included in the template.</li> <li>5. (Optional) Select the <b>VPN Disable Mode</b> when creating templates for hardware models that have the -NV indicator in the model name; these models are hard coded to disallow VPN configuration for countries that do not allow VPN connectivity.</li> <li>6. Select the <b>Devices</b> for which you plan to use this template. To apply a template to a device, you have to select the devices individually.</li> </ol> <p><b>Note</b> If a new managed device is added to Panorama, you must add the new device to the appropriate template. When you commit your changes to the Template, the configuration is pushed to all the devices assigned to the template.</p> <ol style="list-style-type: none"> <li>7. (Optional) Select the <b>Group HA Peers</b> check box for firewalls that are set up as an HA pair. Adding both devices or virtual systems to the same settings to both peers.</li> <li>8. Click <b>OK</b>.</li> <li>9. Click <b>Commit</b>, and select <b>Panorama</b> as the <b>Commit Type</b> to save the changes to the running configuration on Panorama.</li> <li>10. Click <b>Commit</b>, and select <b>Template</b> as the <b>Commit Type</b> to push the changes to the devices included in the template.</li> </ol>

## ADD A TEMPLATE (CONTINUED)

**Step 2** Verify that the template is available.



After you add the first template, the **Device** and **Network** tabs will display on Panorama.

In the **Network** and **Device** tabs, a **Template** drop-down displays. Verify that the newly added template displays in the drop-down.


**Step 3** Apply a configuration change using the template.




Let's specify a base configuration that defines a Primary DNS server for the devices in the template.

1. In the **Template** drop-down, select the template that you want to configure.
2. Select **Device > Setup > Services**, and edit the Services section.
3. Enter an IP address for the **Primary DNS Server**.
4. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.
5. Click **Commit**, and select **Template** as the **Commit Type** to push the changes to the devices included in the selected template.

**Step 4** Verify that the device is configured with the template settings that you pushed from Panorama.



1. Switch to the device context for a firewall that you pushed the setting to using the template.
2. Go to **Device > Setup > Services**. The IP address that you pushed using the template displays. The template icon  also displays.

**Note** To delete/remove a template, you must disable the template on the managed device locally. You must have superuser privileges on the device to disable the template.

## Override Template Settings

While templates allow you to create a base configuration that can be applied to multiple devices, you might want to configure device-specific settings that are not applicable to all the devices in a template. Template overrides allow for exceptions or modifications to meet your deployment needs. If, for example, a template was used to create a base configuration but a few devices in a test lab environment need different settings for the DNS server IP address or the NTP server, you can override the settings defined in the template.



### OVERRIDE TEMPLATE SETTINGS

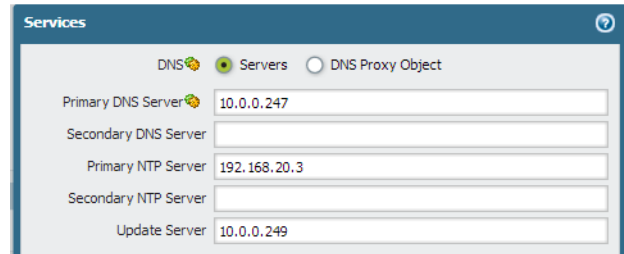
**Step 1.** Access the web interface of the managed device.

You can either directly launch the IP address of the firewall or you can switch to the device context on Panorama.

**OVERRIDE TEMPLATE SETTINGS (CONTINUED)**

**Step 2** Navigate to the setting that you need to modify on the device. In this example, we will override the DNS server IP address that you assigned using a template in [Add a New Template](#).

1. Go to **Device > Setup > Services** and edit the Services section.
2. To override the template, click the icon  to override the value defined for the Primary DNS server IP address.
3. Enter a new value for the Primary DNS Server. Note that the template override icon  now displays to indicate that the value that was pushed using a template has been modified on the device.



4. Click **OK**.
5. Click **Commit** to save your changes on the device.

## Disable Template Settings

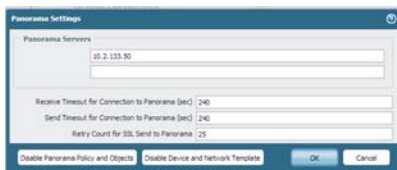
If you want to stop using templates for managing the configuration on a managed device, you can disable the template. When disabling a template, you can choose to copy the template settings to the local device configuration or to delete the values that were previously pushed using the template.

**OVERRIDE TEMPLATE SETTINGS**

1. Access the web interface of the managed device. You can either directly launch the IP address of the firewall or you can switch to the device context on Panorama.

**Note** Superuser privileges are required for disabling templates.

2. Select **Device > Setup > Management**, and click the Edit button on the Panorama Settings section.
3. Select **Disable Device and Network Template**.



4. (Optional) Select **Import Device and Network Template before disabling**, to save the configuration settings locally on the device. If this option is not selected, all Panorama pushed settings will be deleted from the device.
5. Click **OK**.
6. Click **Commit** to save the changes.

## Configure the Firewalls to Forward Logs to Panorama

By default, all log files are generated and stored locally on the firewall. In order to aggregate logs on Panorama, you must configure each firewall to forward logs to Panorama.

If you have compliance policies that require data archival, you can also forward logs to an external service for archive, notification, and/or analysis. Panorama cannot forward logs received from the firewalls to an external server.

To set up log forwarding, complete the following tasks:

- Create a *Server Profile* for each external service to which you want the firewalls to forward logs (Syslog, Email, SNMP Trap) forwarding destination. A server profile defines how to access the remote server, and authenticate to the service, if necessary. You do not need a server profile, if you only plan to forward logs to Panorama or to a Log Collector.

- Configure each log type for forwarding.

For each log type you can specify whether to forward to Syslog, email, and/or an SNMP trap receiver, in addition to Panorama. If you have a distributed log collection architecture, when forwarding to Panorama is enabled, the log forwarding preference list is used to forward logs to the configured Log Collectors. While you can perform these tasks manually on each firewall, you can use device groups and templates on Panorama for a more streamlined workflow.

Log type	Description	Workflow Using Panorama
Traffic Logs	To forward Traffic logs, you must set up a Log Forwarding Profile and add it to the security policies for which you want forwarding to occur. Only traffic that matches a specific rule is logged and forwarded.	Use Device Groups to create a Log Forwarding profile for forwarding Traffic and Threat logs ( <b>Objects &gt; Log Forwarding</b> ) to Panorama and to an external service/syslog server, if required.
Threat Logs (includes URL Filtering Logs, WildFire, and Data Filtering Logs)	To forward Threat logs, you must create a Log Forwarding Profile that specifies which severity levels you want to forward and then add it to the security policies for which you want forwarding to occur. You must also attach a Security Profile (Antivirus, Anti-spyware, Vulnerability, URL Filtering, File Blocking, Data Filtering, or DoS Protection) to the security policy. A Threat log entry will only be created (and therefore forwarded) if the associated traffic matches a security profile.  The result of files analyzed by WildFire are included with Threat logs. WildFire log entries with a benign verdict are logged as <b>Informational</b> , while those with a malware verdict are logged as <b>Medium</b> . So, to enable forwarding to logs to Panorama and to a syslog server, you must enable events of informational and medium severity levels.	If you are forwarding logs to an external service/syslog server, you must create a Syslog Server Profile ( <b>Device &gt; Server Profiles &gt; Syslog</b> ). The Log Forwarding Profile uses the Syslog Server Profile, which you will configure using Templates, to access the server.  All the Traffic and Threat logs forwarded to Panorama can be viewed on the corresponding tab in the <b>Monitor &gt; Logs</b> tab.



Log type	Description	Workflow Using Panorama
System Logs	System logs show system events such as HA failures, link status changes, and administrative access to the device. For each severity level that you want to forward logs for, you must select forwarding to Panorama, Email, SNMP Traps, and a syslog server, if required.	<p>For System, Config, and HIP Match logs, you must configure a template, and select the Panorama check box to enable forwarding to Panorama in the corresponding tab in the <b>Device &gt; Log Settings</b> tab.</p> <p>For forwarding syslogs to an external service for archival to traditional syslog servers or to SIEM servers (for example, Splunk, Arcsight, Qradar) you also must set up a server profile using a Template (<b>Device &gt; Server Profiles &gt; Syslog</b>).</p>
Config Logs	Configuration logs record changes to the configuration. To enable forwarding of Config logs, you must select forwarding to Panorama, Email, SNMP Traps, and a syslog server, if required.	
HIP Match Logs	<p>To enable forwarding of HIP Match logs, you must select forwarding to Panorama, Email, SNMP Traps, and a syslog server, if required.</p> <p>The Host Information Profile (HIP) Match logs are used to compile information on GlobalProtect clients. A HIP Match log is generated when a device sends a HIP Report and a HIP profile is configured with HIP objects such as OS version, patch level, disk encryption, antivirus version, and so on, that happen to match on the device.</p>	

Refer to the [PAN-OS Getting Started Guide](#) for details on performing these tasks directly on the firewall.

For information on forwarding the logs that are locally generated by Panorama itself, see [Chapter 6, Administer Panorama](#).

## SET UP LOG FORWARDING

- |   |   |
|---|---|
| <p><b>Step 1.</b> (Optional) Create a Server Profile that contains the information for connecting to the external service/Syslog server(s).</p> | <ol style="list-style-type: none"> <li>1. Select a template or create a new template. See <a href="#">Step 1 in Add a New Template</a>.</li> <li>2. Verify that you have selected a template from the <b>Template</b> drop-down.</li> <li>3. Select <b>Device &gt; Server Profiles &gt; Syslog</b>.</li> <li>4. Click <b>Add</b> and then enter a <b>Name</b> for the profile.</li> <li>5. Click <b>Add</b> to add a new Syslog server entry and enter the information required to connect to the Syslog server (you can add up to four Syslog servers to the same profile):             <ul style="list-style-type: none"> <li>• <b>Name</b>—Unique name for the server profile.</li> <li>• <b>Server</b>—IP address or fully qualified domain name (FQDN) of the Syslog server.</li> <li>• <b>Port</b>—The port number on which to send Syslog messages (default is 514); you must specify the same port number on Panorama and the Syslog server.</li> <li>• <b>Facility</b>—Select one of the Syslog standard values, which is used to calculate the priority (PRI) field in your Syslog server implementation. You must select the value that maps to how you use the PRI field to manage your Syslog messages.</li> </ul> </li> <li>6. (Optional) To customize the format of the Syslog messages the firewall sends, select the <b>Custom Log Format</b> tab. For details on how to create custom formats for the various log types, refer to the <a href="#">Common Event Format Configuration Guide</a>.</li> <li>7. Click <b>OK</b> to save the server profile.</li> </ol> |
|---|---|

SET UP LOG FORWARDING (CONTINUED)	
<p><b>Step 2</b> Set up a log forwarding profile for traffic and threat logs.</p> <p><b>Note</b> Threat logs include URL Filtering, Data Filtering, and WildFire logs; the logs are forwarded based on the severity levels for which you enable notification.</p>	<ol style="list-style-type: none"> <li>1. Create a new Device Group or select one. To create a new device group, see <a href="#">Create Device Groups</a>.</li> <li>2. Select <b>Objects &gt; Log Forwarding</b>.</li> <li>3. Click <b>Add</b> and then enter a <b>Name</b> for the Log Forwarding Profile.</li> <li>4. (Optional) Select the <b>Shared</b> check box, to apply these settings to all managed devices.</li> <li>5. Select the <b>Panorama</b> check box for the severity levels for which you would like to enable log forwarding.</li> <li>6. (Optional) Select the server profile for forwarding to a syslog server.</li> </ol> <p><b>Note</b> Make sure that the device (or virtual system) is included in the Device Group and that the template defined in <a href="#">Step 1</a> is applied to the device (or virtual system).</p> <ol style="list-style-type: none"> <li>7. Click <b>OK</b>.</li> </ol>
<p><b>Step 3</b> Enable log forwarding for System, Config, and HIP Match logs.</p>	<p>With the same template selected, optionally, select the log types that you would like to forward.</p> <ul style="list-style-type: none"> <li>• For System logs, select <b>Device &gt; Log Settings &gt; System</b> and select the link for each <b>Severity</b> and enable forwarding to <b>Panorama</b> and the select the server profile to use for forwarding to the <b>Syslog</b> server.</li> <li>• For Config logs, select <b>Device &gt; Log Settings &gt; Config</b> and edit the Log Settings - Config section to enable forwarding to <b>Panorama</b> and the select the server profile to use for forwarding to the <b>Syslog</b> server.</li> <li>• For HIP Match logs, select <b>Device &gt; Log Settings &gt; HIP Match</b> and edit the Log Settings - HIP Match section to enable forwarding to <b>Panorama</b> and the select the server profile to use for forwarding to the <b>Syslog</b> server.</li> </ul>

**SET UP LOG FORWARDING (CONTINUED)**

<p><b>Step 4</b> (Optional) Schedule log export to an SCP or an FTP server.</p> <p><b>Note</b> If you plan to use SCP, you must log in to each managed device and click the <b>Test SCP server connection</b> button after the template is pushed. The connection is not established until the firewall accepts the host key for the SCP server.</p>	<p>For Traffic, Threat, URL Filtering, Data Filtering, and HIP Match logs, you can schedule log export using Panorama templates.</p> <ol style="list-style-type: none"> <li>1. Select <b>Device &gt; Scheduled Log Export</b> tab.</li> <li>2. Select the template in the <b>Template</b> drop-down.</li> <li>3. Click <b>Add</b> and then enter a <b>Name</b> for the Log Forwarding Profile.</li> <li>4. Select the <b>Enable</b> check box to enable log export.</li> <li>5. Select the type of log that you would like to export. To schedule the export of more than one log type, you must create a log export profile for each type of log.</li> <li>6. Enter the time of day (hh:mm) to start the export, using a 24-hour clock (00:00 - 23:59).</li> <li>7. Select the protocol you want to use to export logs from the firewall to a remote host. You can use <b>SCP</b> (secure), or <b>FTP</b>. To enable Passive FTP, select the <b>Enable FTP Passive Mode</b> check box.</li> <li>8. Define the details required to connect to the server. <ol style="list-style-type: none"> <li>a. Enter the hostname or IP address for the server.</li> <li>b. If required for your server, enter the port number (by default, FTP uses port 21 and SCP uses port 22), path or directory in which to save the exported logs, access credentials (username and password).</li> </ol> </li> <li>9. Click <b>OK</b>.</li> </ol>
<p><b>Step 5</b> Save all the configuration changes.</p>	<ol style="list-style-type: none"> <li>1. Click <b>Commit</b>, and select <b>Panorama</b> as the <b>Commit Type</b> to save the changes to the running configuration on Panorama.</li> <li>2. Click <b>Commit</b>, and select <b>Template</b> as the <b>Commit Type</b> to push the changes to the devices included in the selected template.</li> <li>3. Click <b>Commit</b>, and select <b>Device Groups</b> as the <b>Commit Type</b> to push the changes to the devices included in the selected device group.</li> </ol>

## Commit Changes on Panorama

When you edit the configuration on Panorama, you are making changes to the candidate configuration file. The candidate configuration is a copy of the running configuration along with the modifications that you have saved using the **Save** option. The Panorama web interface displays all the configuration changes immediately, however the changes are not implemented until you commit the changes. The commit process validates the changes in the candidate configuration file and saves it as the running configuration on Panorama.

Options on Panorama	Description
<b>Panorama</b>	Commits the changes on the current candidate configuration to the running configuration on Panorama. You must first commit your changes on Panorama, before committing any configuration updates (templates or device groups) to the managed devices or collector groups.
<b>Template</b>	Commits template changes from Panorama to the selected devices.
<b>Device Group</b>	Commit policies and objects configured from Panorama to the selected device/virtual system(s).
<b>Collector Group</b>	Commit changes to the specified collector groups managed by Panorama.

When a commit completes, a result displays. On a successful commit the **Commit succeeded** message displays, if there are warnings, the **Commit succeeded with warnings** message displays.

Some of the other choices that you have when committing your changes are as follows:

- Include Device and Network Templates**—This option is available when committing a Device Group from Panorama. It allows you to commit both Device Group and Template changes, to the pertinent devices, in a single commit operation.  
If you prefer to commit your changes as separate commit operations, do not select this check box.
- Force Template Values**—When performing a Template commit, the **Force Template Values** option overrides all local configuration and removes objects on the selected devices or virtual systems that do not exist in the template or have been overridden by the local configuration. This is an override that reverts all existing configuration on the managed device, and ensures that the device inherits the settings defined in the template only.
- Merge with Candidate Config**—When enabled, this option allows you to merge and commit the Panorama configuration changes with any pending configuration changes that were implemented locally on the target device. If this option is not enabled, the candidate configuration on the device is not included in the commit operation. As a best practice, leave this option disabled if you allow device administrators to modify the configuration directly on a device and you don't want to include their changes when committing changes from Panorama. Another best practice is to use the configuration audit capability on Panorama to review any locally defined configuration changes prior to issuing a commit from Panorama, see [Compare Changes in Configuration](#).

## Modify the 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 configuration, specify which Panorama peer can receive logs. These options are available on the Logging and Reporting section of the **Panorama > Setup > Management** tab.

- Define the log forwarding mode on the device: The firewalls can forward logs to Panorama in either *Buffered Log Forwarding* mode or in the *Live Mode Log Forwarding* mode.

Logging Options	Description
<b>Buffered Log Forwarding from Device</b>  Default: Enabled	<p>Allows each managed device to buffer logs and send the logs at 30-second intervals to Panorama (not user configurable).</p> <p>Buffered log forwarding is very valuable when the device loses connectivity to Panorama. The device 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 device resumes forwarding logs from where it left off.</p> <p>The disk space available for buffering depends on the log storage quota for the platform and the volume of logs that are pending roll over. In the event that the device 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 device's local hard disk is consumed, the oldest entries are deleted to allow logging of new events.</p>
<b>Live Mode Log Forwarding from Device</b>  This option is enabled when the check box for <b>Buffered Log Forwarding from Device</b> is cleared.	<p>In live mode, the managed device sends every log transaction to Panorama at the same time as it records it on the device.</p>

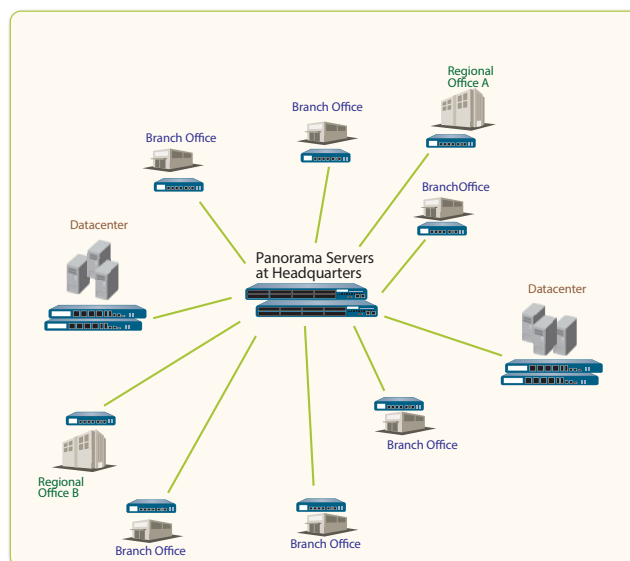
- Define log forwarding preference on a Panorama virtual appliance that is in a high availability configuration:
  - (when logging to a virtual disk) enable logging to the local disk on the Active-Primary Panorama peer only. By default, both Panorama peers in the HA configuration receive logs.
  - (when logging to an NFS) enable the devices to only send newly generated logs to a secondary Panorama peer, which is promoted to primary, after a failover.

Logging Options	Pertains to	Description
<b>Only Active Primary Logs to Local Disk</b>  Default: Disabled	Panorama virtual appliance that is logging to a virtual disk and is set up in a high availability (HA) configuration.	Allows you to configure only the Active-Primary Panorama peer to save logs to the local disk.

Logging Options	Pertains to	Description
<b>Get Only New Logs on Convert to Primary</b> Default: Disabled	Panorama virtual appliance that is mounted to a Network File System (NFS) datastore and is set up in a high availability (HA) configuration	<p>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 devices can only send logs to the primary Panorama peer, which can write to the NFS datastore.</p> <p>When an HA failover occurs, the <b>Get Only New Logs on Convert to Primary</b> option allows an administrator to configure the managed devices to only send 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 devices from sending a large volume of buffered logs when connectivity to Panorama is restored after a significant period of time.</p>

## Use Panorama to Configure Managed Devices: An Example

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 datacenters, and a firewall in each of the two regional head offices.



## Assemble Devices into Device Groups and Templates

The first step in creating your central management strategy is to determine how to group the devices into Device Groups and Templates to efficiently push configurations. You can group the devices in different ways based on the device's business function, geographic location, or administrative domain. In this example, we create two Device Groups and three Templates to administer the devices using Panorama.

### Device Groups

In this example, we decide to define two Device Groups based on the functions the firewalls will perform:

- *DG\_BranchAndRegional* for grouping devices 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 devices with similar functions will require similar policy rulebases.
- *DG\_DataCenter* for grouping the devices that secure the servers at the datacenters.

We can then administer shared policies across both Device Groups as well as administer distinct Device Group policies 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 policies; 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.		

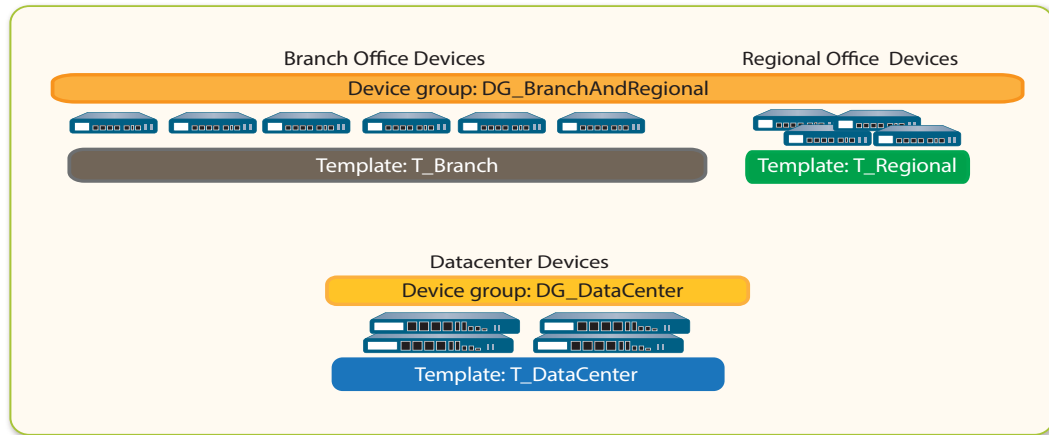
### Templates

When grouping devices 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 devices must be in separate templates. Further, the way the devices are configured to access



network resources might be different because the devices 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, you must place the devices in separate templates as follows:

- *T\_Branch* for the branch office devices
- *T\_Regional* for the regional office devices
- *T\_DataCenter* for the devices at the datacenter



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.

## Plan Your Centralized Configuration and Policies

Let's use the example we started with above and perform the following tasks for centrally deploying and administering the managed devices:

Step 1: Add Managed Devices and deploy content updates and PAN-OS software updates to the managed devices.

Step 2: Use Templates to administer a base configuration.

Step 3: Use Device Groups for managing the policies on your firewalls.

Step 4: Preview your rules and then commit your changes to Panorama, Device Groups, and Templates.

## DEPLOY CONTENT UPDATES AND PAN-OS SOFTWARE UPDATES TO THE MANAGED DEVICES

**Step 1.** [Add Managed Devices](#) and deploy content updates and PAN-OS software updates to the managed devices.

First install the **Applications** or **Applications and Threats** database, then the **Antivirus**, and finally update the **Software** version.

If you have purchased a Threat Prevention subscription, the content and antivirus databases are available to you.

1. Select **Panorama > Device Deployment > Dynamic Updates**.
  - a. Click **Check Now** to check for the latest updates. If the value in the Action column is **Download** it indicates that an update is available.
  - b. Click **Download**. When the download completes, the value in the Action column changes to **Install**.
  - c. Click the **Install** link in the **Action** column. Use the filters or user-defined tags to select the managed devices on which you would like to install this update. Click **OK**.
  - d. Monitor the status, progress, and the result for the content update for each device. The success or failure of the installation displays in the **Results** column.



**Note** To review the status or progress for all tasks performed on Panorama, see [View Task Completion History](#).

2. Select **Panorama > Device Deployment > Software** to deploy software updates.
  - a. Click **Check Now** to check for the latest updates. If the value in the Action column is **Download** it indicates that an update is available.
  - b. Locate the version that you need for each hardware model and then click **Download**. When the download completes, the value in the Action column changes to **Install**.
  - c. Click the **Install** link in the Action column. Use the filters or user-defined tags to select the managed devices on which you would like to install this version.
  - d. Enable the check box for **Upload only to device (do not install)** or **Reboot device after install** and click **OK**. The success or failure of the installation displays in the **Results** column.


## USE TEMPLATES TO ADMINISTER A BASE CONFIGURATION

**Step 2** Use Templates to administer a base configuration.

1. Create templates and assign the appropriate devices to the template. See [Add a new Template](#).
2. Define a DNS server, NTP server, Syslog server, and login banner.
  - a. Select the template in the **Template** drop-down.
  - b. Select **Device > Setup > Services**, and edit the Services section.
    - i. Enter an IP address for the **Primary DNS Server**.
    - ii. Enter an IP address for the **Primary NTP Server**.
  - c. To add a syslog server, select **Device > Server Profiles > Syslog**.
    - i. Enter a **Name** for the profile.
    - ii. Click **Add**, and then click **Add** to add a new Syslog server entry and enter the information required to connect to the Syslog server (you can add up to four Syslog servers to the same profile):
      - **Name**—Unique name for the server profile.
      - **Server**—IP address or fully qualified domain name (FQDN) of the Syslog server.
      - **Port**—The port number on which to send Syslog messages (default is 514); you must use the same port number on Panorama and the Syslog server.
      - **Facility**—Select one of the Syslog standard values, which is used to calculate the priority (PRI) field in your Syslog server implementation. You must select the value that maps to how you use the PRI field to manage your Syslog messages.
    - iii. Click **OK** to save the server profile.
  - d. To add a login banner, select **Device > Setup > Management**, and edit the General Settings section.
    - i. Add the text for the **Login Banner**.
    - ii. Click **OK**.
  - e. Repeat tasks 2a to 2d for each template.
3. Enable HTTPS, SSH, and SNMP access to the management interface of the managed devices.
  - a. Select the template in the **Template** drop-down.
  - b. Select **Device > Setup > Management**, and edit the Management Interface Settings section.
  - c. Select the check box for the **HTTPS**, **SSH**, and **SNMP** under Services.
  - d. Click **OK**.
  - e. Repeat tasks 3a to 3d for each template.

## USE TEMPLATES TO ADMINISTER A BASE CONFIGURATION (CONTINUED)


4. Create and attach a zone protection profile to the untrust zone for the devices in the Datacenter Template (T\_DataCenter).
  - a. Select the template in the **Template** drop-down.
  - b. Select **Network > Network Profiles > Zone Protection**.
  - c. Click **Add** to add a new profile and enter the information required to define the profile. In this example we will enable protection against a SYN Flood and alert for TCP Port Scan, Host Sweep and UDP Port Scan.



- d. To attach the profile to the untrust zone, you must first set up the interface and the zone settings in the template.

**Note** You must have set up the interfaces on the device. At a minimum, you must have defined the interface type, assigned it to a virtual router, if needed, and attached a security zone locally on the device.

- i. Select **Network > Interface**.
- ii. Select the appropriate interface from the table. Click the link to configure the interface.
- iii. Select the Interface **Type** from the drop-down.
- iv. Click the **Virtual Router** link in the **New Virtual Router** drop-down to create a new virtual router. Make sure that the virtual router name matches what is defined on the device.
- v. Click the **New Zone** link in the **Security Zone** drop-down to create a new zone. Make sure that the zone name matches what is defined on the device.
- vi. Click **OK**.
- vii. Select **Network > Zones**, and select the zone that you just created. Verify that the correct interface is attached to the zone.
- viii. In the **Zone Protection Profile** drop-down, select the profile you created above.



- ix. Click **OK**.

5. Commit your template changes.
  - a. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama. Click **OK**.
  - b. Click **Commit**, and select **Template** as the **Commit Type** to push your changes to the devices included in the selected template. Click **OK**.

## USE DEVICE GROUPS TO PUSH POLICIES

**Step 3** Use Device Groups for managing the policies on your firewalls.

1. Create device groups and assign the appropriate devices to each device group. See [Create Device Group\(s\)](#).
2. Create a shared pre-rule to allow DNS and SNMP services.
  - a. Create a shared Application Group for the DNS and SNMP services.
    - i. Select **Objects > Application Group** and click **Add**.
    - ii. Enter a name and select the **Shared** check box to create a shared Application Group object.
    - iii. Click **Add**, and type in DNS and select **dns** from the list. Repeat for SNMP and select **snmp**, **snmp-trap**.
    - iv. Click **OK**. The application group for dns, snmp, and snmp-trap is created.
  - b. Create the shared policy.
    - i. Select the **Shared** Device Group in the **Device Group** drop-down.
    - ii. Select the **Policies** tab, and select **Pre-Rules** in the **Security** policies rulebase.
    - iii. Click **Add** and enter a **Name** for the security policy rule.
    - iv. In the **Source** and **Destination** tabs for the rule, click **Add** and enter a **Source Zone** and a **Destination Zone** for the traffic.
    - v. In the **Applications** tab, click **Add** and type in the name of the applications group object you defined earlier, and then select it from the drop-down.
    - vi. In the **Actions** tab, verify that the Action Setting is **Allow** and click **OK**.

Context		Device Group: Shared						
Security								
Pre Rules								
Post Rules								
NAT								
Pre Rules								
		Source		Destination				
Name	Location	Zone	Address	Zone	Address	Application	Service	
allow_basic_services	Shared	trust_zone	any	untrust_zone	any	basic_apps	any	

## USE DEVICE GROUPS TO PUSH POLICIES (CONTINUED)

### 3. Define the corporate acceptable use policy for all offices.

In this example, we will create shared policy that restricts access to some URL Categories and denies access to peer-to-peer traffic that is of risk level 3, 4, 5.

- a. Select the **Shared** Device Group in the **Device Group** drop-down.
- b. Select the **Policies** tab, and select **Pre-Rules** in the **Security** policies rulebase.
- c. Click **Add** and enter a **Name** for the security policy rule.
- d. In the **Source** and **Destination** tabs for the rule, click **Add** and select **any** for **Source Zone** and **Destination Zone** for the traffic.
- e. To define the application filter, in the **Application** tab:
  - i. Click **Add** and click **New Application Filter**.
  - ii. Enter a **Name**, and select the check box for **Shared**; in the Risk section select levels **3**, **4**, and **5** and in the Technology section select **peer-to-peer**.
  - iii. Click **OK**.
- f. In the **Service/URL Category** tab, click **Add** and select the URL Categories that you would like to block, for example streaming-media, dating, and online-personal-storage.
- g. You can also attach the *default* URL Filtering profile, in the Profile Setting section of the **Actions** tab.
- h. Click **OK**.

Device Group: Shared										
		Source			Destination					
Name	Location	Zone	Address	User	Zone	Address	Application	Service	Action	Options
corp_AUP	Shared	any	any	any	any	any	block-high-risk	any		

### 4. Allow Facebook to all users in the marketing group in the regional offices only.

In order to enable security policy based on user and/or group, you must enable User-ID for each zone that contains users you want to identify. You must have set up User Identification on the firewall (refer to the [PAN-OS Getting Started Guide](#)) and have defined a master device for the Device Group. The master device is the only device in the Device Group that gathers user and group mapping information for policy evaluation.

- a. Select the *DG\_BranchAndRegional* Device Group in the **Device Group** drop-down.
- b. Select the **Policies** tab, and select **Pre-Rules** in the **Security** policies rulebase.
- c. Click **Add** and enter a **Name** for the security policy rule.
- d. In the **User** tab, select **Select**, click **Add** and select the marketing user group in the Source User section.
- e. In the **Application** tab, click **Add** and type in *Facebook* and then select it from the drop-down.
- f. In the **Action** tab, verify that the action is **Allow**.
- g. In the **Target** tab, select the regional office devices and click **OK**.

Device Group: Shared											
(source-user/member eq 'any')											
		Source			Destination						
Name	Location	Zone	Address	User	Zone	Address	Application	Service	Action	Options	Target
Allow-FB	Shared	any	any	companyABC...	any	any	facebook	any			0016060001... PA-4060-2 0006C102148 TSP-B

## USE DEVICE GROUPS TO PUSH POLICIES (CONTINUED)

5. Allow access to the Amazon cloud application for the specified hosts/servers in the datacenter.
  - a. Create an Address Group object for the servers/hosts in the datacenter that need access to the Amazon cloud application.
    - i. Select **Objects > Address Groups**.
    - ii. Select the *DG\_DataCenter* Device Group in the **Device Group** drop-down.
    - iii. Click **Add** and enter a **Name** for the Address Group object.
    - iv. Click **Add** and select **New Address**.
    - v. To define the Address Object, enter a **Name**, and select the **Type** and specify a host IP address, IP Netmask, IP range, or FQDN. Click **OK**.
  - b. Select the *DG\_DataCenter* Device Group in the **Device Group** drop-down.
    - i. Select the **Policies** tab, and select **Pre-Rules** in the **Security** policies rulebase.
    - ii. Click **Add** and enter a **Name** for the security policy rule.
    - iii. In the Source Address section of the **Source** tab, click **Add** to select the Address Group that you defined.
    - iv. In the **Application** tab, click **Add**, type in *amazon* and select the Amazon applications from the list that displays.
    - v. In the **Action** tab, verify that the action is **Allow**.
    - vi. Click **OK**.

Access-EC3	DG_DataCenter	any	any	any	any	any	any	amazon-cloud-drive	amazon-cloud-player	SSL	✓
------------	---------------	-----	-----	-----	-----	-----	-----	--------------------	---------------------	-----	---

6. To enable logging for all Internet-bound traffic on your network, create a rule that matches trust zone to untrust zone.
  - a. Select the **Shared** Device Group in the **Device Group** drop-down.
  - b. Select the **Policies** tab, and select **Pre-Rules** in the **Security** policies rulebase.
  - c. Click **Add** and enter a **Name** for the security policy rule.
  - d. In the **Source** and **Destination** tabs for the rule, click **Add** and select *trust\_zone* as the source zone and *untrust\_zone* as the destination zone.
  - e. In the **Action** tab, verify that the action is **Deny**, and the Log Setting is **Log at Session end**.
  - f. Click **OK**.

**Step 4** Preview your rules and then commit your changes to Panorama, Device Groups, and Templates.

1. Select the **Policies** tab, and click **Preview Rules**. This preview allows you to visually evaluate how your rules are layered for a particular rulebase.
2. Click **Commit** and select Commit Type as **Panorama**. Click **OK**.
3. Click **Commit**, and select Commit Type as **Device Groups**. Verify that the **Include Device and Network Templates** option is enabled. Click **OK**.
4. Switch device context to launch the web interface of a managed device and confirm that the template and policy configurations have been applied.

# Enable Logging

All Palo Alto Networks next-generation firewalls can generate logs that provide an audit trail of the activities and events on the firewall. To centrally monitor the logs and to generate reports, you must forward the logs generated on the managed firewalls to Panorama. If you are deploying a Panorama virtual appliance with a virtual disk or are logging to an NFS, you do not need to perform any additional tasks to enable logging.

If you will be logging to an M-100 appliance—either locally on an M-100 in Panorama mode, or to a dedicated Log Collector that is managed by a Panorama virtual appliance or an M-100 appliance in Panorama mode—you have to perform some additional tasks to enable log collection.

You have to add each Log Collector as a Managed Collector and create Collector Group(s) in order to access, manage and update the Log Collector(s) using Panorama. After you add and configure the Log Collectors on Panorama, Panorama pushes the necessary configuration to the managed devices. You need not explicitly configure the managed devices to forward logs to a Log Collector.

Complete the following tasks to enable logging:

- ▲ [Add a Log Collector to Panorama](#)
- ▲ [Configure Collector Groups](#)
- ▲ [Verify that Log Forwarding is Enabled](#)
- ▲ (Optional) [Modify the Log Forwarding and Buffering Defaults](#)

## Add a Log Collector to Panorama

In order for Panorama—Panorama virtual appliance or an M-100 appliance in Panorama mode—to manage a Log Collector, you have to add the Log Collector as a Managed Collector. If the M-100 appliance is not already set up in Log Collector mode, see [Set Up the M-100 Appliance in Log Collector Mode](#).

If you are using an M-100 appliance in Panorama mode, the *default* log collector that is local on Panorama is added during the manufacturing process. However, if you have loaded/migrated the configuration from a Panorama virtual appliance to the M-100 appliance, the default log collector does not display; use the instructions in this section to add the log collector and then [Configure Collector Groups](#).

ADD A MANAGED COLLECTOR	
Step 1. Add a Managed Collector	<ol style="list-style-type: none"><li>1. Select <b>Panorama &gt; Managed Collectors</b>.</li><li>2. Select <b>Add</b>.</li></ol>
Step 2. Add the serial number of the Log Collector on Panorama.	<ol style="list-style-type: none"><li>1. In the <b>General</b> tab, enter the serial number for the Log Collector in the <b>Collector S/N</b> field.<ul style="list-style-type: none"><li>• If the log collector is local on Panorama, enter the serial number that displays on the <b>Dashboard</b>.</li><li>• If you are adding a dedicated Log Collector, enter the serial number for that M-100 appliance.</li></ul></li></ol>



**ADD A MANAGED COLLECTOR (CONTINUED)**

**Step 3** Complete these tasks only if you are adding a dedicated Log Collector.

Because the default Log Collector is located on the same physical appliance as Panorama that is managing it, you do not need to configure management access or authentication preferences for it.

**Note** An M-100 appliance in Log Collector mode only has CLI access; there is no web interface for managing a Log Collector.

1. Configure the network access settings.  
 Although you have already specified these details during initial configuration on the Log Collector, you must re-enter the information on the **Panorama > Managed Collectors** tab.
  - a. In the **General** tab, add the IP address of the Panorama servers that will manage the Log Collector. If you have deployed Panorama in HA, add the IP address for both the primary and secondary peers.
  - b. Configure the DNS server IP addresses.
  - c. (Optional) Set the time zone that will be used for recording log entries.
2. Using Panorama, configure administrative access to the Log Collector.
  - a. In the **Authentication** tab, the default user is **admin**. You can neither modify this username nor add administrative users on the Log Collector.
  - b. Specify the number of **Failed Attempts** to login after which the user is locked out from accessing the Log Collector and the time interval for which the user is locked out in the **Lockout Time**.
  - c. Specify a password. To generate a hashed password, complete the following tasks:
    - i Enter the following command in the CLI:
 

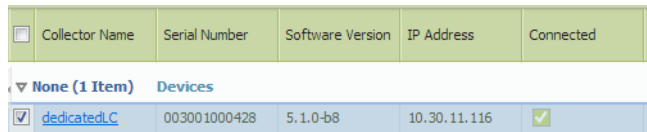
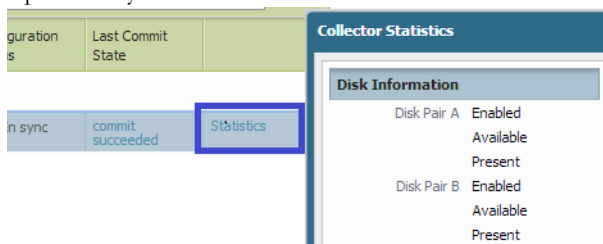
```
request password-hash password <yourpassword>
```

The CLI will display the hashed output for the password you entered.
    - ii Copy the hash value and paste it into the **Password Hash** field.

When you commit the changes to the Collector Group the new hashed password will be pushed to the Log Collector.
3. Specify the management port (MGT) settings that you defined on the Log Collector during initial configuration.
  - a. In the **Management** tab, enter the **IP address**, **Netmask**, and **Default Gateway** IP address defined on the Log Collector.
  - b. (Optional) Enable **SNMP** access for monitoring the Log Collector. By default, you have SSH and ping enabled on the management port.
  - c. (Optional) To restrict access to the Log Collector, click **Add** and enter one or more IP addresses in the **Permitted IP Addresses** list.

**Note** Because only the specified IP addresses can access the Log Collector, make sure to add the IP address for the devices that need to connect and forward logs to the Log Collector.

  - d. Click **OK**.

ADD A MANAGED COLLECTOR (CONTINUED)	
<b>Step 4</b> Save the changes.	Click <b>Commit</b> and in the Commit Type select <b>Panorama</b> . Click <b>OK</b> .
<b>Step 5</b> Verify that the Log Collector has been added and is connected to Panorama.	<p>Select <b>Panorama &gt; Managed Collectors</b> and check that the managed collector you added displays.</p> 
<b>Step 6</b> Enable the disk pairs for logging. To set up the disks in a RAID pair, see <a href="#">Increase Storage Capacity on the M-100 Appliance</a> .	<p>By default, the Disk Pair A is RAID enabled and added to the Log Collector. If you have added additional RAID pairs for increased storage capacity:</p> <ol style="list-style-type: none"><li>1. Click <b>Add</b> in the <b>Disks</b> tab.</li><li>2. Select each additional disk pair from the drop-down.</li><li>3. Click <b>OK</b> to make the new disk pair available for logging.</li><li>4. Click <b>Commit</b> and in the Commit Type select <b>Panorama</b>. Click <b>OK</b>.</li></ol>
<b>Step 7</b> Verify that the disks are enabled, available, and present.	<p>Select <b>Panorama &gt; Managed Collectors</b>, and click the <b>Statistics</b> link. The Collector Statistics window will display the status of the disk pairs that you added.</p> 

## Configure Collector Groups

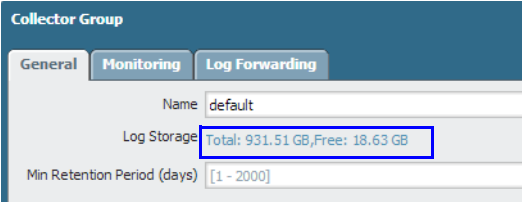
After you add a Log Collector as a Managed Collector, you must assign it to a Collector Group so that it can be managed and configured using Panorama. A Collector Group allows you to assign the managed firewalls to the Log Collector(s) in the Collector Group.

While a Collector Group can include one or more Log Collectors, Palo Alto Networks recommends placing only one Log Collector in a Collector Group. However, if you require more than 4TB of log storage capacity, you will need to add multiple Log Collectors in a Collector Group. To understand how logging works with multiple Log Collectors in a Collector Group, see [Using Multiple Log Collectors in a Collector Group](#).

If you are using an M-100 appliance in Panorama mode, a default Collector Group that contains the default Log Collector is set up for you. Use the instructions in this section to configure the default Collector Group. If you are using a dedicated Log Collector, you must first [Add a Log Collector to Panorama](#) and use the instructions in this section to create a Collector Group.

### CONFIGURE COLLECTOR GROUPS

<p><b>Step 1.</b> Add a Collector Group.</p>	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Collector Groups</b>. A <b>default</b> Collector Group exists for an M-100 appliance in Panorama mode.</li> <li>2. Select the link for the <b>default</b> Collector Group to modify it or click <b>Add</b> to define a new Collector Group.</li> </ol>
<p><b>Step 2</b> Define the members of the Collector Group.</p>	<ol style="list-style-type: none"> <li>1. Select the <b>Log Forwarding</b> tab in the Collector Group window.</li> <li>2. Click <b>Add</b> in the Collector Group Members section, and select the Log Collectors to include in the Collector Group. Only the Log Collectors that you have added as Managed Collectors will display in the selection list.</li> </ol>
<p><b>Step 3</b> Select which devices can forward logs to this Collector Group.</p> <p><b>Note</b> If your network has firewalls running PAN-OS version 4.x and 5.x, you can assign the firewalls running PAN-OS v5.x to forward logs to a dedicated log collector. Firewalls running PAN-OS v4.x cannot be assigned to a Log Collector; they must send logs to a Panorama virtual appliance or an M-100 appliance in Panorama mode.</p>	<ol style="list-style-type: none"> <li>1. In the <b>Log Forwarding</b> tab in the Collector Group window, click <b>Add</b> in the Log Forwarding Preferences section.</li> <li>2. Click <b>Modify</b>, and select the Managed <b>Device</b> from the filtered display options and click <b>OK</b>. Then click <b>Add</b> in the Collectors section, and select the Log Collector. The selected devices can send logs to the assigned Log Collector in the Collector Group.</li> </ol> <div data-bbox="823 976 1250 1134" data-label="Image"> </div> <p><b>Note</b> If you have multiple Log Collectors in the Collector Group, click <b>Add</b> again and select another Log Collector to define a prioritized list. The first Log Collector in the list is the Primary Log Collector assigned to the firewall.</p> <ol style="list-style-type: none"> <li>3. Click <b>OK</b>.</li> </ol>

CONFIGURE COLLECTOR GROUPS	
<p><b>Step 4</b> Allocate the percentage of storage capacity for each log type.</p> <p>If the Log Storage capacity displays as 0MB, you may not have added Log Collectors to the Collector Group. Complete <a href="#">Step 2</a> and then come back to this task (<a href="#">Step 4</a>).</p> <p>If it still reads as 0MB, verify that you have enabled the disk pairs for logging and have committed the changes to the Collector Group. See <a href="#">Step 6</a> in the <a href="#">Add a Log Collector to Panorama</a> section.</p>	<ol style="list-style-type: none"><li>1. Select the <b>General</b> tab in the Collector Group window.</li><li>2. Click the link that displays the <b>Log Storage</b> capacity for the Collector Group.</li></ol> <div>The screenshot shows the 'Collector Group' configuration window with the 'General' tab selected. It displays fields for 'Name' (default), 'Log Storage' (Total: 931.51 GB, Free: 18.63 GB), and 'Min Retention Period (days)' ([1 - 2000]).</div> <ol style="list-style-type: none"><li>3. Modify the <b>Quota</b> allocated for each log type. As you change the value, the screen refreshes to display the corresponding number value (GB/MB) for the percentage allocated based on the total storage on your Collector Group.</li><li>4. (Optional) Click <b>Restore Defaults</b>, if you need to undo your changes and reset the quotas to factory defaults.</li></ol>
<p><b>Step 5</b> Define the minimum log retention time for the Collector Group.</p> <p>The minimum log retention time informs Panorama when to generate an alert if the storage capacity is near full capacity.</p>	<ol style="list-style-type: none"><li>1. Select the <b>General</b> tab in the Collector Group window.</li><li>2. Enter a value between 1-2000 days for the <b>Minimum Retention Period</b>. This value specifies how long you would like to retain logs. A system log is generated on Panorama when the current date minus the oldest log date is less than the defined minimum retention period.</li></ol>



## Remove a Device from a Collector Group

In a distributed log collection deployment, where you have dedicated Log Collectors, if you need a device to send logs to Panorama instead of sending logs to the Collector Group, you must remove the device from the Collector group.

When you remove the device from the Collector Group and commit the change, the device will automatically send logs to Panorama instead of sending it to a Log Collector.

### REMOVE A DEVICE FROM A COLLECTOR GROUP

1. Select the **Panorama > Collector Groups** tab.
2. Click the link for the desired Collector Group, and select the **Log Forwarding** tab.
3. In the Log Forwarding Preferences section, select the device that you would like to remove from the list and click **Delete**.
4. Click **OK**.
5. Click **Commit** and in the Commit Type select **Panorama**. Click **OK**.
6. Click **Commit** and in the Commit Type select **Collector Group**. Click **OK**.



To temporarily remove the log forwarding preference list on the device, you can delete it using the CLI on the device. You must however, remove the assigned firewalls in the Collector Group configuration on Panorama. Otherwise, the next time you commit changes to the Collector Group, the device will be reconfigured to send logs to the assigned Log Collector.

## Verify that Log Forwarding is Enabled

Now that you have added the Log Collector(s) as Managed Collectors, created and configured the Collector Group and assigned the managed devices to forward logs to the specified Collector Group, you can test that your configuration was successful.

### VERIFY LOG FORWARDING

**Step 1.** On the managed device, check that the device has the Log Forwarding Preference list and is forwarding logs to the configured Log Collector.

You cannot view this information from the web interface on the device.

1. Access the CLI on the device.
2. Enter the following commands:

**a. show log-collector preference-list**

If you have assigned only one Log Collector to the Collector Group, the onscreen output will look something like this:

```
Log collector Preference List
Serial Number: 003001000024
IP Address:10.2.133.48
```

**b. show logging-status**

The onscreen output will look something like this:

```
admin@PA-200> show logging-status
```

Type	Last Log Created	Last Log Fwded	Last Seq Num Fwded	Last Seq Num Recd
> CMS 0	Not Sending to CMS 0			
> CMS 1	Not Sending to CMS 1			
> Log Collector	Log Collector log forwarding agent is active and connected to 10.2.133.48			
config	2012/07/13 18:39:34	2012/10/04 17:03:20		531
system	2012/07/13 18:40:07	2012/10/04 17:03:20		3434
threat	2012/10/11 17:23:48	2012/10/11 17:24:08		94343
traffic	2012/10/11 17:24:01	2012/10/11 17:24:08		1063
hipmatch	Not Available	Not Available		0

**Step 2** On Panorama, verify the log collection rate.

Click the **Statistics** link in the **Panorama > Managed Collectors** tab to view the average logs/second being received by Panorama.

Collector Name	Serial Number	Software Version	IP Address	Conn...	Configuration Status	Last Commit State	
▼ DedicatedLCG (1 Item) Devices							
<input type="checkbox"/> dedicated.C	003001000428	5.1.0-b8	10.30.11.116	<input checked="" type="checkbox"/>	In sync	commit succeeded	<a href="#">Statistics</a>

## Deploy Software Updates and Manage Licenses

As an administrator, you can use Panorama to centrally track and manage licenses, handle software updates and dynamic content updates on the managed devices and managed collectors. Panorama checks in with the Palo Alto Networks licensing server or update server, verifies the validity of the request and then allows retrieval and installation of the license/software version on the managed device or log collector. This capability allows ease of deployment because you do not need to repetitively perform the tasks on each device/log collector. It is particularly useful for managed devices that do not have direct Internet access or for managing updates to the M-100 appliance configured in Log Collector mode, which does not support a web interface.

Use this centralized deployment capability to qualify new content updates or software updates on select devices before you perform the update on all managed devices. Or, to retrieve new licenses using an authorization code and push the license keys to the managed device.

Depending on which subscriptions are active on each device, the content updates can include the latest application update/application and threat signature updates, antivirus signatures, WildFire updates, and GlobalProtect data file updates. The software updates that you can manage from Panorama include: PAN-OS, SSL VPN client, and GlobalProtect client.



You must activate the support subscription directly from each firewall; Panorama cannot be used to deploy the support subscription.



## DEPLOY SOFTWARE AND LICENSES ON MANAGED DEVICES USING PANORAMA

- Deploy dynamic updates.

Based on the subscriptions you have purchased, you might need to install **Antivirus** updates, **Applications** or **Applications and Threats** updates and **WildFire** updates, GlobalProtect Data File, and BrightCloud **URL Filtering** database updates.

- Select **Panorama > Device Deployment > Dynamic Updates**.
- Check for the latest updates. Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates. The link in the **Action** column indicates whether an update is available. If a version is available, the **Download** link displays; for the BrightCloud URL Filtering database, the link displays as **Upgrade**.



- Click **Download** to download a selected version. After successful download, the link in the **Action** column changes from **Download** to **Install**.
- Click **Install** and select the devices on which you want to install the update. When the installation completes, a check mark displays in the **Currently Installed** column.

Deploy software updates.

For a managed collector, use the image that corresponds to platform name **m**; for a managed device find the image that corresponds to the hardware model, for example, **5000**.

This example shows you how to install a PAN-OS software update. The SSL VPN client (**Panorama > Device Deployment > SSL VPN Client**) and the GlobalProtect client (**Panorama > Device Deployment > GlobalProtect Client**) use the same mechanism. However, you do not *install* the software on the firewall, instead you *activate* it on the firewall so that it can be downloaded onto client systems.

- Select **Panorama > Device Deployment > Software**.
- Check for the latest updates. Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates. The link in the **Action** column indicates whether an update is available.
- Review the **File Name** and click **Download**. Verify that the software versions that you download match the firewall models deployed on your network. After successful download, the link in the **Action** column changes from **Download** to **Install**.
- Click **Install** and select the devices on which you want to install the software version. The result of the installation attempt displays onscreen.

**Note** You can download a maximum of five versions of software per category to Panorama. After five versions, when a new download is initiated, the oldest image is automatically deleted.

Verify the software and content update version running on each managed device.

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

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

DEPLOY SOFTWARE AND LICENSES ON MANAGED DEVICES USING PANORAMA (CONTINUED)	
Verify the software and content update version running on each managed collector.	<div><div>1.</div><div>To verify the version on a managed collector, you must access the CLI of the managed collector. See <a href="#">Log in to the CLI</a>.</div></div> <div><div>2.</div><div>Enter the command <b>show system info</b></div></div> <div>The following details must display:<div><div>sw-version: 5.1.0-b10</div><div>app-version: 366-1738</div><div>app-release-date: 2013/03/29 15:46:03</div><div>av-version: 1168-1550</div><div>av-release-date: 2013/04/21 14:31:27</div><div>threat-version: 366-1738</div><div>threat-release-date: 2013/03/29 15:46:03</div></div></div>

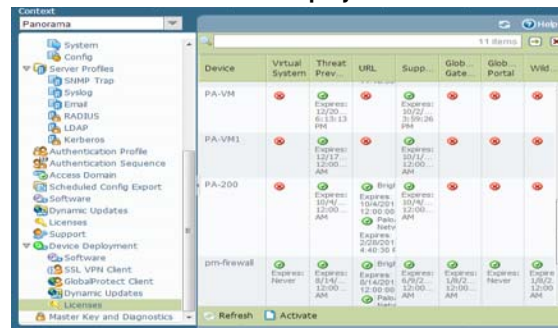
**DEPLOY SOFTWARE AND LICENSES ON MANAGED DEVICES USING PANORAMA (CONTINUED)**

- Deploy licenses.

Each entry on the **Panorama > Device Deployment > Licenses** tab indicates whether the license is active or inactive; it also displays the expiration date for active licenses.

**Note** You cannot use Panorama to activate the support license for the managed devices.


1. Select **Panorama > Device Deployment > Licenses**.



2. If you have previously activated the authorization code for the support subscription directly on the firewall, click **Refresh** and select one or more devices from the list. Panorama retrieves the license(s), deploys it to the managed devices and updates the licensing status on the Panorama web interface.

**Note** This ability to retrieve and deploy license using Panorama is particularly useful for license renewals for devices that do not have direct Internet access.

3. To activate new licenses:
  - a. Click **Activate**. This option allows you to activate a newly purchased subscription, for example, a Threat subscription.
  - b. Find or filter for the managed devices and enter the authentication code(s) that Palo Alto Networks provided for the device in the **Auth Code** column.



- c. Click **Activate**.

## Replace a Managed Device with a New Device

To minimize the effort required to restore the configuration on a managed device on a Return Merchandise Authorization (RMA), you can replace the serial number of the old device with that of the new/replacement device on Panorama. To then restore the configuration on the replacement device, you can either import a device state that you have previously generated and exported from the device or you can use Panorama to generate a partial *device state* for managed devices running PAN-OS v5.0 and later versions. The partial device state that you create replicates the configuration of the managed devices with a few exceptions for Large Scale VPN (LSVPN) setups. It is created by combining two facets of the configuration on a managed device:

- Centralized configuration managed by Panorama: Panorama maintains a snapshot of the shared policies and templates pushed from Panorama.
- Local configuration on the device: When a configuration change is committed, each device sends a copy of its local configuration file to Panorama. This file is stored on Panorama and is used 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 can export using the **Export device state** operation from the **Device > Setup > Operations** tab on the firewall. If you have manually run the device state export or have scheduled an XML API script to export the file to a remote server, you can use the exported device state in your device replacement workflow below.

If you have not exported the device state, the device state that you generate in this workflow will not include the dynamic configuration information, such as the certificate details and registered devices, that is required to restore the complete configuration of a device functioning as an LSVPN portal. See [Before you Begin](#) for more information.

The device state is not stored on Panorama; it is generated on request using the CLI commands listed in [Restore the Configuration on the New Device](#). By replacing the serial number and importing the device state, you can resume managing the device using Panorama.

### Before you Begin

- The managed device (that was replaced) must have been on PAN-OS v5.0.4 and later version. Panorama cannot generate the *device state* for devices running older PAN-OS versions. If you need to restore the configuration for a device running a PAN-OS version earlier than 5.0.4, refer to this article: [Configuration Recovery with Panorama](#).
- Make note of the following details on the old device:
  - **Serial number:** You will need to enter the serial number on the Support portal to transfer the licenses from the old device to your replacement device. You will also enter this information on Panorama, to replace all references to the older serial number with the serial number of the replacement device.
  - (Recommended) **PAN-OS version and the content database version:** Installing the same software and content database versions, including the URL database vendor allows you to create the same state on the replacement device. If you decide to install the latest version of the content database, you may notice differences because of updates and additions to the database. To verify the versions installed on the device, access the device system logs stored on Panorama.

- Prepare the replacement device for deployment. Before you import the device state bundle and restore the configuration, you must:
  - Verify that the replacement device is of the same model and is enabled for similar operational capability. Consider the following operational features: does it need to be enabled for multi- virtual systems, support jumbo frames, or enabled to operate in CC or FIPS mode?
  - Configure network access, transfer the licenses, and install the appropriate PAN-OS version and the content database version.
- You must use the Panorama CLI to complete this device replacement process. This CLI-based workflow is available for the *superuser* and *panorama-admin* user roles.
- If you have an LSVPN configuration, and are replacing a Palo Alto Networks firewall deployed as a satellite device 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 have been following the recommendation to frequently generate and export the device state for devices in an LSVPN configuration, use the device state that you have previously exported from the device itself instead of generating one on Panorama.

If you have not manually exported the device state from the device, and need to generate a partial device state on Panorama, the missing dynamic configuration impacts the device replacement process as follows:

- **If the device you are replacing is a portal device** that is explicitly configured with the serial number of the satellite devices (**Network > GlobalProtect > Portals > Satellite Configuration**), when restoring the device configuration, although the dynamic configuration is lost, the portal device will be able to authenticate the satellite devices successfully. The successful authentication will populate the dynamic configuration information and LSVPN connectivity will be reinstated.
- **If you are replacing a satellite device**, the satellite device will not be able to connect and authenticate to the portal. This connection failure occurs either because the serial number was not explicitly configured on the device (**Network > GlobalProtect > Portals > Satellite Configuration**) or because although the serial number was explicitly configured, the serial number of the replaced device does not match that of the old device. To restore connectivity, after importing the device state bundle, the satellite administrator must log in to the device and enter the credentials (username and password) for authenticating to the portal. When this authentication occurs, the dynamic configuration required for LSVPN connectivity is generated on the portal.

However, if the device was configured in a high availability configuration, after restoring the configuration, the device will automatically synchronize the running configuration with its peer and attain the latest dynamic configuration required to function seamlessly.

## Restore the Configuration on the New Device

Use the following workflow to restore the device configuration.

RESTORE THE DEVICE CONFIGURATION	
<p>▲ Tasks on the new device</p> <p>Use the CLI for a more streamlined workflow.</p>	
<p><b>Step 1.</b> Perform initial configuration and verify network connectivity.</p>	<p>Use a serial port connection or an SSH connection to add an IP address, a DNS server IP address, and to verify that the device can access the Palo Alto Networks updates server.</p> <p>For instructions, refer to the <i>Palo Alto Networks Getting Started Guide</i>.</p>
<p><b>Step 2</b> (Optional) Set the operational mode to match that on the old device. A serial port connection is required for this task.</p>	<ol style="list-style-type: none"> <li>1. Enter the following CLI command to access maintenance mode on the device: <b>debug system maintenance-mode</b></li> <li>2. To boot into the maintenance partition, enter <b>maint</b> during the boot sequence.</li> <li>3. Select the operational mode as <b>Set FIPS Mode</b> or <b>Set CCEAL 4 Mode</b> from the main menu.</li> </ol>
<p><b>Step 3</b> Retrieve the license(s).</p>	<p>Enter the following command to retrieve your licenses: <b>request license fetch</b></p>
<p><b>Step 4</b> (Optional) Match the operational state of the new device with that of the old device. For example, enable multi-virtual system (multi-vsyz) capability for a device that was enabled for multi-vsyz capability.</p>	<p>Enter the commands that pertain to your device settings: <b>set system setting multi-vsyz on</b></p> <p><b>set system setting jumbo-frame on</b></p>
<p><b>Step 5</b> Upgrade the PAN-OS version on the device.</p> <p>You must upgrade to the same OS and content database version that installed on the old device.</p>	<p>Enter the following commands:</p> <ol style="list-style-type: none"> <li>1. To upgrade the content database version: <b>request content upgrade download &lt;xxx-xxxx&gt;</b></li> <li>2. To install the content database version that you downloaded: <b>request content upgrade install version &lt;xxx-xxxx&gt;</b></li> <li>3. To upgrade the PAN-OS software version: <b>request system software download version 5.x.x</b></li> <li>4. To install the content database version that you downloaded: <b>request system software install version 5.x.x</b></li> </ol>



**RESTORE THE DEVICE CONFIGURATION (CONTINUED)**

## ▲ Tasks on the Panorama CLI.

You cannot perform these tasks on the Panorama web interface.

<p><b>Step 6</b> Replace the serial number of the old device with that of the new replacement device on Panorama.</p> <p>By replacing the serial number on Panorama you allow the new device to connect to Panorama after you restore the configuration on the device.</p>	<ol style="list-style-type: none"> <li>1. Enter the following command in operational mode: <b>replace device old &lt;old SN#&gt; new &lt;new SN#&gt;</b></li> <li>2. Go in to configuration mode and commit your changes. <b>configure</b> <b>commit</b></li> <li>3. Exit configuration mode.</li> </ol>
<p>(Skip this step if you have manually exported the device state from your firewall, and go to <a href="#">Step 8</a> below)</p> <p><b>Step 7</b> Export the device state bundle to a computer using SCP or TFTP.</p> <p>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).</p>	<p>Enter one of the following commands:</p> <pre>scp export device-state device &lt;new serial#&gt; to &lt;login&gt; @ &lt;serverIP&gt;: &lt;path&gt;</pre> <p>or,</p> <pre>tftp device-state device &lt;new serial#&gt; to &lt;login&gt; @ &lt;serverIP&gt;: &lt;path&gt;</pre>

## ▲ Tasks on the new device.

<p><b>Step 8</b> Import the device state and commit the changes on the device.</p>	<ol style="list-style-type: none"> <li>1. Access the web interface of the device.</li> <li>2. Select <b>Device &gt; Setup &gt; Operations</b> and click the <b>Import Device State</b> link in the Configuration Management section.</li> <li>3. Browse to locate the file and click <b>OK</b>.</li> <li>4. Click <b>Commit</b> to save your changes to the running configuration on the device.</li> <li>5. To confirm that the device state restored includes the references to Panorama pushed policies and objects, look for the little green icon .</li> </ol> 
--	---

RESTORE THE DEVICE CONFIGURATION (CONTINUED)


▲ Tasks on Panorama

You can now use the Panorama web interface to access and manage the replaced device.

Step 9 Verify that the device configuration was successfully restored.

- 1. Access the Panorama web interface.
- 2. Verify that the new device is connected to Panorama.
- 3. Perform a Template and Device Groups commit to keep the device synchronized with Panorama.

Device Name	Virtual System	Tags	Serial Number	IP Address	Template	Connected
900-1-inkd			0006C1061396	10.3.84.98	900-Template	<input checked="" type="checkbox"/>



After replacing the device, if you need to generate reports for a period that spans the duration when the old device was functional and after you installed the replacement device, you must generate a separate query for each device's serial number because replacing the serial number on Panorama does not overwrite the information in the logs.



## Transition a Device to Central Management

If you have already deployed Palo Alto Networks firewalls and configured them locally, but now want to start using Panorama for centrally managing them, you have pre-migration planning, implementation and post-migration verification tasks. This high-level overview does not address all the critical tasks required to plan, implement, and validate the transition to centralized administration. Here are the high-level planning and configuration activities.

- On Panorama, add the devices and create device groups to logically assemble firewalls or virtual systems that perform a similar role, or function or that have similar characteristics.
- Create common zones for each device group. Decide on the common zone-naming strategy for all devices and virtual systems in a device group. For example, if you have two zones called Branch LAN and WAN, Panorama can centrally push policies that reference those zones without being aware of the variations in port/media type, platform or the logical addressing schema. You must create the zones on each managed device before you can commit the changes to the device group or template. Panorama cannot poll the devices for zone name or configuration.
- Configure each device to communicate with Panorama. You must define the Panorama IP addresses (primary and secondary Panorama) on each device.
- Use device groups to create common policies for devices with similar functionality and use templates to define a common base configuration for the managed device.
- Determine how you will manage local rules and device-specific exceptions to common policies and configuration settings. If you plan to use locally configured rules on the devices, make sure that the names of the rules are unique. A good way to ensure this would be to add a suffix or a prefix to all existing rules.
- Consider removing all “deny rules” in local security policy and use Panorama post-rules. This approach allows you to temporarily disable local rules and test the shared post-rules pushed from Panorama. You can then test the post-rules, make adjustments as necessary and eliminate local administration on the device.
- Validate that the firewalls are functioning efficiently with Panorama-pushed configuration as they did with local configuration.

For detailed information on using the XML-based REST API to complete the transition, refer to the document: [Panorama Device Migration](#). To access this document, you must have an account on the partner portal. If you do not have an account on the partner portal, contact Palo Alto Networks Professional Services to learn about the device migration process.





## 4 Monitor Network Activity

---

Panorama provides a comprehensive, graphical view of network traffic. Using the visibility tools on Panorama—the Application Command Center (ACC), logs, and the report generation capabilities—you can centrally analyze, investigate and report on all network activity, identify areas with potential security impact, and translate them into secure application enablement policies.

This section covers the following topics:

- ▲ [Use Panorama for Visibility](#)
- ▲ [Use Case: Monitor Applications Using Panorama](#)
- ▲ [Use Case: Use Panorama to Respond to an Incident](#)

# Use Panorama for Visibility

In addition to its central deployment and device 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), App-Scope, the log viewer, and the standard and 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. You can for example, 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.

This section provides 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 Reports](#)

## Monitor the Network with the ACC and AppScope

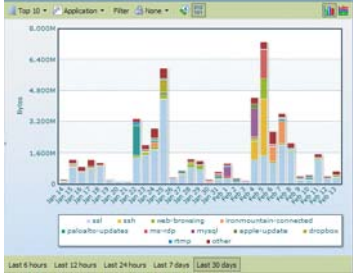
Both the ACC and the AppScope allow you to monitor and report on the data recorded from traffic that traverses your network.

The ACC on Panorama shows a summary of network traffic. Panorama can dynamically query data from all the managed devices 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, data filtering, file blocking, HIP match for GlobalProtect—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.

Use the ACC and the AppScope to answer questions such as:

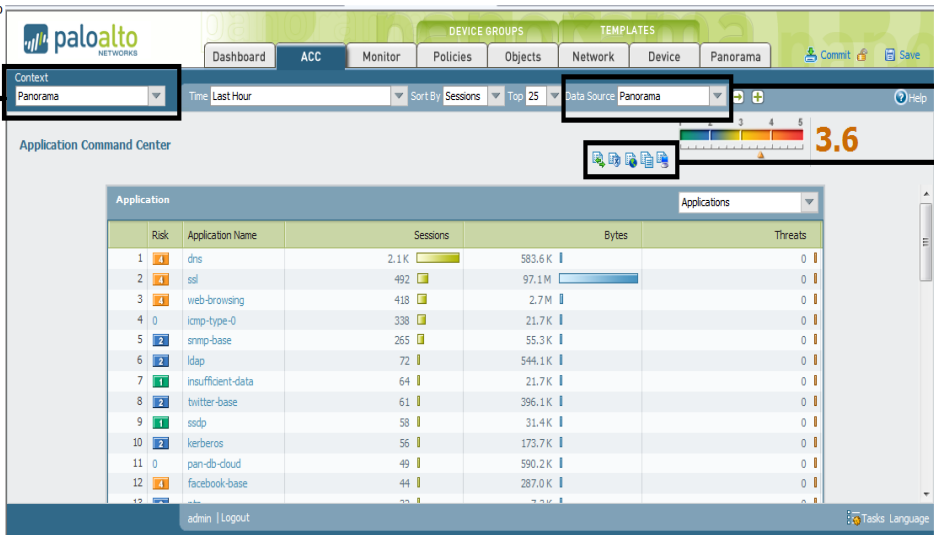
ACC	Monitor > AppScope
<ul style="list-style-type: none"><li>• What are the top applications used on the network and how many are high-risk applications? Who are the top users of high-risk applications on the network?</li><li>• What are the top URL categories being viewed in the last hour?</li></ul>	<ul style="list-style-type: none"><li>• 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?</li><li>• How has user activity changed over the current week as compared to last week or last month?</li></ul>

ACC	Monitor > AppScope
<ul style="list-style-type: none"> <li>What are the top bandwidth-using applications? Who are the users/hosts that consume the highest bandwidth?</li> <li>What content or files are being blocked and are there specific users who trigger this file blocking/data filtering policy?</li> <li>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?</li> </ul>	<ul style="list-style-type: none"> <li>Which users and applications take up most of the network bandwidth? And how has this consumption changed over the last 30 days?</li> </ul>  <ul style="list-style-type: none"> <li>What are the threats on the network, and how are these incoming and outgoing traffic threats distributed geographically?</li> </ul>

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:

Use the **Context** switch to access the web interface of any managed device from Panorama.



Switch the **Data Source** to:

- access the logs stored on Panorama (default).
- access the data from the managed firewalls.

Panorama will query the devices for data.

Access the logs directly. The log details that display match the information you are viewing on this page.

Application	Risk	Application Name	Sessions	Bytes	Threats
1	1	dns	2.1 K	583.6 K	0
2	1	ssl	492	97.1 M	0
3	1	web-browsing	418	2.7 M	0
4	0	icmp-type-0	338	21.7 K	0
5	2	snmp-base	265	55.3 K	0
6	2	ldap	72	544.1 K	0
7	1	insufficient-data	64	21.7 K	0
8	2	twitter-base	61	396.1 K	0
9	1	osdp	58	31.4 K	0
10	2	kerberos	56	173.7 K	0
11	0	pan-db-cloud	49	590.2 K	0
12	1	facebook-base	44	287.0 K	0

- Switch from a Panorama view to a Device view:** Panorama allows access to the web interface of any managed device using the **Context** menu. The context switch is a toggle that provides direct firewall access; it provides the ability to manage device-specific settings, such as device-specific policy, and/or override network configuration pushed from a template on a specific device.

- **Change Data Source:** The default source used to display the statistics on the charts in the ACC is the Panorama local data. With the exception of the data that displays in the **Application** chart, all other charts require you to enable log forwarding to Panorama.

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**. When configured to use Remote Device Data, instead of using the local Panorama data, Panorama will poll all the managed devices and present an aggregated view of the data. The onscreen display shows the total number of devices being polled and the number of devices that have responded to the query for information.

- **Select the Charts to View:** The ACC includes an array of charts in the areas of Application, URL Filtering, Threat Prevention, Data Filtering, and HIP Match. With the exception of the Application charts and HIP Match, all the other charts only display if the corresponding feature has been licensed on the device, and you have enabled logging.
- **Tweak Time Frame and Sort 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. You can sort the data by sessions, bytes, or threats and filter to view from 5-500 items.

## 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 alarms.

Based on the log forwarding configuration on the managed devices, 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, such as the source and destination ports, zones, and addresses, the user who initiated the session, and the action (allow, deny) taken by the firewall on that session; the system and configuration logs can inform you of a configuration change or an alarm that was triggered when a configured threshold was exceeded.

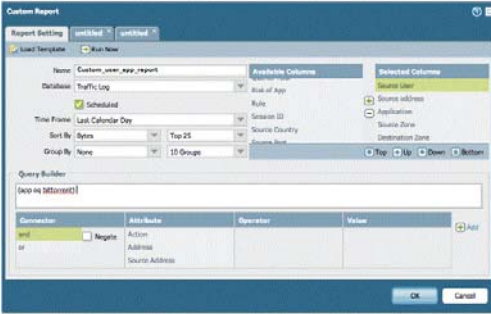
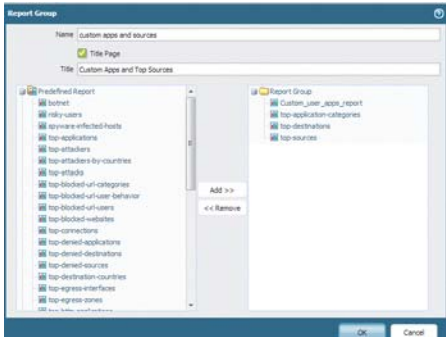
## Generate Reports

Panorama allows you to generate reports manually as needed, or schedule reports to run at specific intervals. You can save and export reports, or you can configure Panorama to email reports to specific recipients. The ability to share reports using email is particularly useful if you want to share reporting information with administrators who do not have access to Panorama.

You can create the following types of reports:

- **Predefined Reports:** A suite of predefined reports are available in four categories —Applications, Threats, URL Filtering, and Traffic—in the **Monitor > Reports** tab.
- **User-activity Reports:** The user activity report is a predefined report that is used to create an on-demand report to document the application use and URL activity broken down by URL category for a specific user with estimated browse time calculations. This report is available in the **Monitor > PDF Reports > User Activity Reports** tab.
- **Custom Reports:** Create and schedule custom reports that show exactly the information you want to see by filtering on conditions and columns to include. You can generate reports to query data from a summary database on Panorama or on the remote devices (that is the managed firewalls), or use the detailed reports on Panorama or on the remote devices. To view the databases available for generating these reports, see the **Monitor > Manage Custom Reports** tab. You can also create Report Groups (**Monitor > PDF Reports > Report Groups** tab) to compile predefined reports and custom reports as a single PDF.
- **PDF Summary Reports:** Aggregate up to 18 predefined reports, graphs, and custom reports into one PDF document.

The following table provides step-by-step instructions for creating and scheduling reports:

GENERATE, SCHEDULE, AND EMAIL REPORTS	
<p><b>Step 1.</b> Generate reports.</p> <p><b>Note</b> You must set up a <b>Report Group</b> to email report(s).</p> <div></div> <div></div>	<ul style="list-style-type: none"><li>• Create a custom report.<ul style="list-style-type: none"><li>a. Select <b>Monitor &gt; Manage Custom Reports</b>.</li><li>b. Click <b>Add</b> and then enter a <b>Name</b> for the report.</li><li>c. Select the database, <b>Panorama</b> or <b>Remote Device Data</b>, that you would like to use for the report. You can use the summary database or the detailed logs on Panorama or on the managed devices.</li><li>d. Select the <b>Scheduled</b> check box.</li><li>e. Define your filtering criteria. Select the <b>Time Frame</b>, the <b>Sort By</b> order, <b>Group By</b> preference, and select the columns that must display in the report.</li><li>f. (Optional) Select the <b>Query Builder</b> attributes, if you want to further refine the selection criteria.</li><li>g. To test the report settings, select <b>Run Now</b>. Modify the settings as required to change the information that is displayed in the report.</li><li>h. Click <b>OK</b> to save the custom report.</li></ul></li><li>• Run a <b>PDF Summary Report</b>.<ul style="list-style-type: none"><li>a. Select <b>Monitor &gt; PDF Reports &gt; Manage PDF Summary</b>.</li><li>b. Click <b>Add</b> and then enter a <b>Name</b> for the report.</li><li>c. Use the drop-down list for each report group and select one or more of the elements to design the PDF Summary Report. You can include a maximum of 18 report elements.</li><li>d. Click <b>OK</b> to save the settings.</li></ul></li><li>• Define the <b>Report Group</b>. It can include predefined reports, PDF Summary reports, and custom reports. Panorama compiles all the reports included into a single PDF.<ul style="list-style-type: none"><li>a. Select <b>Monitor &gt; Report Group</b>.</li><li>b. Click <b>Add</b> and then enter a <b>Name</b> for the report group.</li><li>c. (Optional) Select <b>Title Page</b> and add a <b>Title</b> for the PDF output.</li><li>d. Select from the Predefined Report, PDF Summary Report and the Custom Report lists; click <b>Add</b> to include the selected report(s) to the report group.</li><li>e. Click <b>OK</b> to save the settings.</li></ul></li></ul>




**GENERATE, SCHEDULE, AND EMAIL REPORTS (CONTINUED)**

**Step 2** Set up Panorama to email reports.

1. Select **Panorama > Server Profiles > Email**.
2. Click **Add** and then enter a **Name** for the profile.
3. Click **Add** to add a new email server entry and enter the information required to connect to the Simple Mail Transport Protocol (SMTP) server and send email (you can add up to four email servers to the profile):
  - **Server**—Name to identify the mail server (1-31 characters). This field is just a label and does not have to be the host name of an existing SMTP server.
  - **Email Display Name**—The name to show 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 host name of the SMTP gateway to use to send the emails.
4. Click **OK** to save the server profile.
5. Click **Commit** and select **Panorama** as the **Commit Type** to save the changes to the running configuration.

**Step 3** Schedule the report for delivery by email.



1. Select **Monitor > PDF Reports > Email Scheduler**.
2. Click **Add** and then enter a **Name** for the email scheduler profile.
3. Select the **Report Group**, the **Email Profile**, and the **Recurrence** for the report.
4. To verify that the email settings are accurate, select **Send test email**.
5. Click **OK** to save your settings.

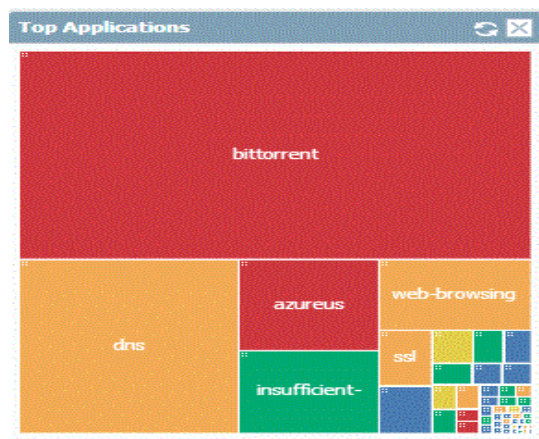
**Step 4** Save the configuration changes.

Click **Commit** and select **Panorama** as the **Commit Type** to save the changes to the running configuration.

## 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. You can either glance over the list of top applications and mouse over each application block that you want to review the details for, or you can navigate to 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**.



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.

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 remote devices varies by the time period for which you choose to view data and the volume of traffic that is generated on your network.

Going back to the list of top applications, we can see that bittorrent is very popular. If you now click into the link for the bittorrent application, the **ACC** view filters the display to show information on the application, its behavior, risk level, and the associated URL categorization details.

	Risk	Application	Sessions	Bytes
1	5	bittorrent	10.5 K	1.6 G
2	4	dns	5.9 K	1.4 M
3	4	web-browsing	1.0 K	21.4 M
4	4	ssl		
5	2	twitter-base		
6	4	facebook-base		
7	1	paloalto-wildfire-cloud		
8	2	ntp		
9	4	ssh		
10	2	idoud-base		
11	3	itunes-base		
12	1	pan-db-cloud		

Application: bittorrent

Application Information

Name: bittorrent

Description: BitTorrent is a peer-to-peer file sharing (P2P) communications protocol. BitTorrent is a method of distributing large amounts of data widely without the original distributor incurring the entire costs of hardware, hosting and bandwidth resources. Instead, when data distributed using the BitTorrent protocol, each recipient supplies pieces of the data to newer recipients, reducing the cost and burden on any given individual source, providing redundancy against system problems, and reducing dependence on the original distributor. The protocol is the brainchild of programmer Bram Cohen, who designed it in April 2001 and released a first implementation on 2 Jul 2001. It is now maintained by Cohen's company BitTorrent, Inc. Usage of the protocol accounts for significant traffic on the Internet but the precise amount has proven difficult to measure. There are numerous compatible BitTorrent clients, written in a variety of programming languages, and running on a variety of computing platforms, include uTorrent, BitComet, Deluge, TurboBT, and Transmission.

Standard Ports: tcp/dynamic, udp/dynamic

Capable of File Transfer: yes

Used by Malware: yes

Excessive Bandwidth Use: yes

Evasive: yes

Category: general-internet

Subcategory: file-sharing

Technology: peer-to-peer

Risk: 5

Widely Used: yes

Has Known Vulnerabilities: yes

In the **Top Sources** table, you can also see how many users are using bittorrent and the volume of traffic being generated. If you have enabled User-ID, you will be able to view the names of the users who are generating this traffic. You can now click on a source user and drill down to review all the activity for that user.

Using the **ACC** view to filter for bittorrent traffic generated by the specific source address or user allows us to verify the source and destination country for this traffic, the device that is processing this traffic, the ingress and egress zones and the security rule that is letting this connection through.

Top Security Rules					
	Virtual System	Device	Rule	Bytes	Sessions
1	vsys1	PA-200_GT	L3L4Inet	1.5 G	5.5 K

Top Ingress Zones					
	Device	Virtual System	Source Zone	Bytes	Sessions
1	PA-200_GT	vsys1	Trust_DC	1.5 G	5.5 K

Top Egress Zones					
	Device	Virtual System	Destination Zone	Bytes	Sessions
1	PA-200_GT	vsys1	untrust	1.5 G	5.5 K

For more detailed information, drill down 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 as compared 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 decide to create a *pre-rule* on Panorama to block all bittorrent traffic. Then use Panorama Device Groups to create and push this policy rule to one or more devices.
- Enforce bandwidth use limits and create a *QoS profile and policy* that de-prioritizes non-business traffic. Then use Panorama templates to push this policy to one or more devices. Refer to the article [Panorama Templates](#) for defining QoS policy using templates.
- Reduce risk to your network assets and create an *application filter* that blocks all file sharing applications that are peer-to-peer technology with a risk factor of 4 or 5. Make sure to verify that the bittorrent application is included in that application filter, and will therefore be blocked.
- 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.

Now, with the combination of the IP addresses of the servers (destination IP in the logs), 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 [Unknown Applications](#) for more information on identifying unknown application and [Custom Application Signature](#) for information on developing custom signatures to discern the application.

## Use Case: Use Panorama to Respond to an Incident

Network threats can originate from different vectors, including malware and spyware infections due to drive-by downloads, phishing attacks, under secured and 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. Used alone or in conjunction with the reports and logs generated from a Security Information Event Manager (SIEM), the correlated attack information can be used to investigate how an attack was triggered and how to prevent future attacks and loss of damage to your network.

The questions we will probe in this section 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?

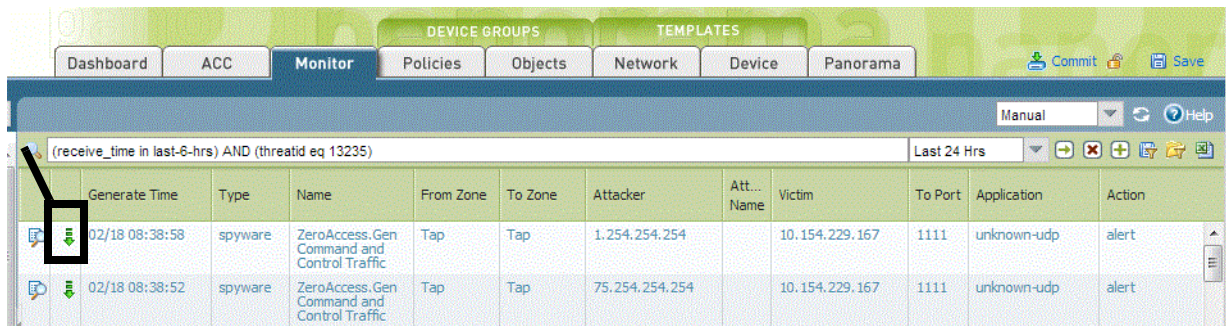
In this use case we will trace a specific incident and show you how the visibility tools on Panorama can help you respond to the report.

There are several ways that you could be alerted to an incident depending on how you've configured the Palo Alto Networks devices 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 an **Zero Access gent.Gen Command And Control Traffic** that matched against a spyware signature. Also listed in the email is the IP address of the source and destination for the session, a threat ID and the timestamp of when the event was logged.

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, has a green arrow 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.



	Generate Time	Type	Name	From Zone	To Zone	Attacker	Att... Name	Victim	To Port	Application	Action
	02/18 08:38:58	spyware	ZeroAccess.Gen Command and Control Traffic	Tap	Tap	1.254.254.254		10.154.229.167	1111	unknown-udp	alert
	02/18 08:38:52	spyware	ZeroAccess.Gen Command and Control Traffic	Tap	Tap	75.254.254.254		10.154.229.167	1111	unknown-udp	alert

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.




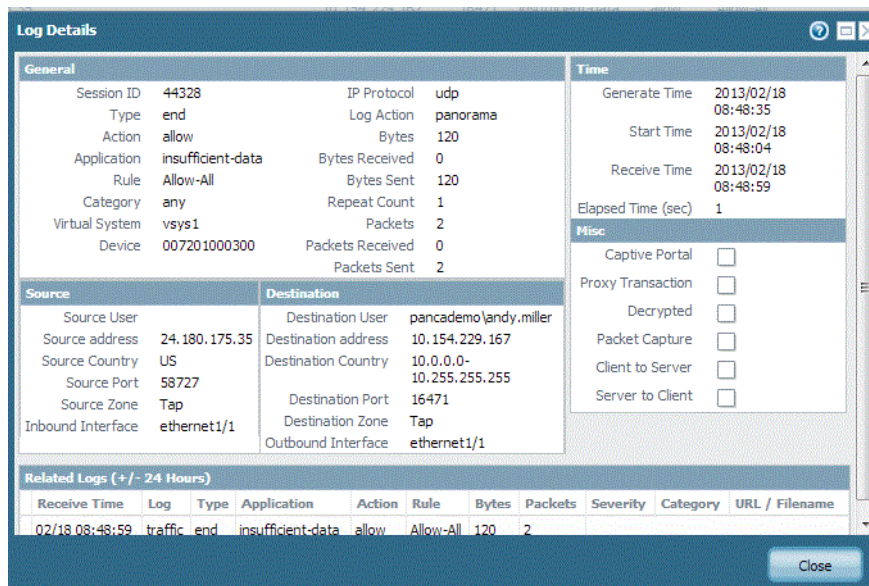
If you see a DNS-based spyware signature, the IP address of your local DNS server might display as the **Victim** IP address. Often this is because the firewall is located north of the local DNS server, and so all DNS queries show the local DNS server as the source IP rather than showing the IP address of the client that originated the request. If you see this issue, consult the local DNS logs to find the client from which the query originated.

For the future, consider routing client DNS requests directly to the firewall so that the firewall can record the victim IP address accurately.

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 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.



General		Time	
Session ID	44328	Generate Time	2013/02/18 08:48:35
Type	end	Start Time	2013/02/18 08:48:04
Action	allow	Receive Time	2013/02/18 08:48:59
Application	insufficient-data	Elapsed Time (sec)	1
Rule	Allow-All	Misc	
Category	any	Captive Portal	<input type="checkbox"/>
Virtual System	vsys1	Proxy Transaction	<input type="checkbox"/>
Device	007201000300	Decrypted	<input type="checkbox"/>
IP Protocol	udp	Packet Capture	<input type="checkbox"/>
Log Action	panorama	Client to Server	<input type="checkbox"/>
Bytes	120	Server to Client	<input type="checkbox"/>
Bytes Received	0		
Bytes Sent	120		
Repeat Count	1		
Packets	2		
Packets Received	0		
Packets Sent	2		

Source		Destination	
Source User		Destination User	pancademo\andy.miller
Source address	24.180.175.35	Destination address	10.154.229.167
Source Country	US	Destination Country	10.0.0.0-10.255.255.255
Source Port	58727	Destination Port	16471
Source Zone	Tap	Destination Zone	Tap
Inbound Interface	ethernet1/1	Outbound Interface	ethernet1/1

Related Logs (+/- 24 Hours)										
Receive Time	Log	Type	Application	Action	Rule	Bytes	Packets	Severity	Category	URL / Filename
02/18 08:48:59	traffic	end	insufficient-data	allow	Allow-All	120	2			

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 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/write 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 on the application that caused the infection (web-browsing, SMTP, FTP), make more stringent URL filtering policies, restrict some applications/actions such as file downloads to specific user groups.

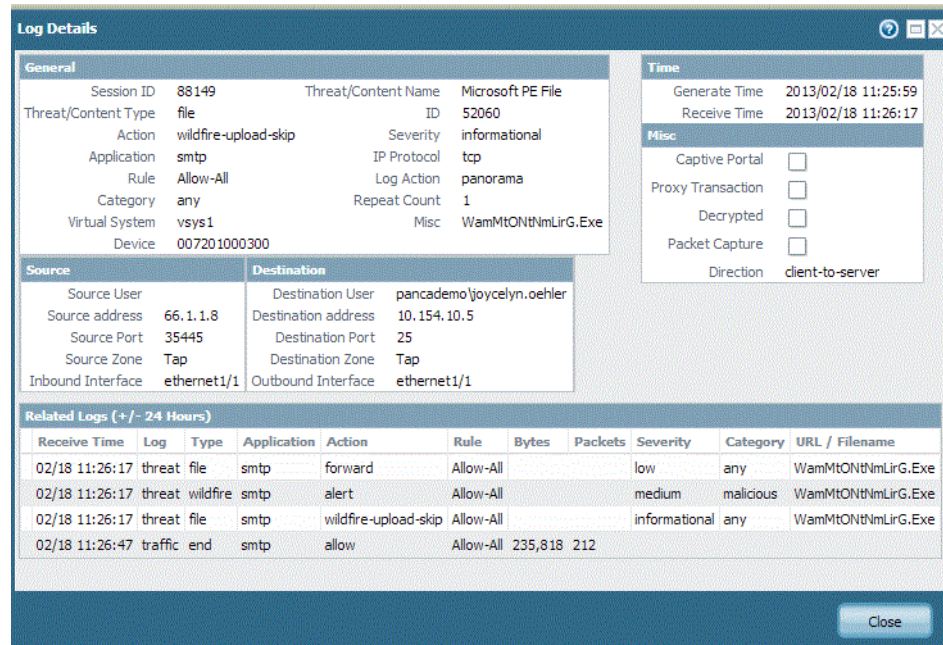


Access to the WildFire logs from Panorama requires the following: a WildFire subscription, a file blocking profile that is attached to a security policy, and threat log forwarding to Panorama.

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.



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 devices, 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 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.

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 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 they can access the web server and import the list at a defined frequency. After creating a dynamic block list object, define a security policy 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 policies or device group policies to block specific applications that caused the infection (web-browsing, SMTP, FTP), make more stringent URL filtering policies, restrict some applications/actions such as file downloads to specific user groups.
- On Panorama, you can also switch to the device context and configure the firewall for botnet reports that identify potential botnet-infected hosts on the network.



## 5 Panorama High Availability

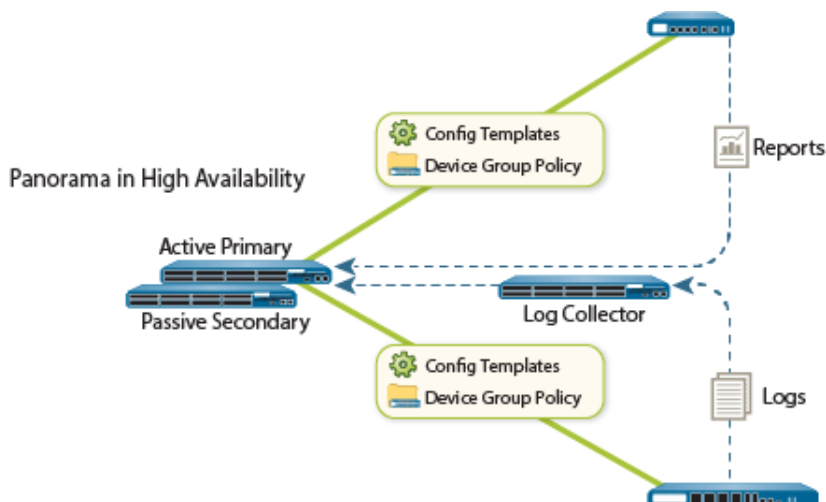
---

Panorama High availability (HA) is a configuration in which two Panorama servers are placed in a group (two-device cluster) to provide redundancy in the event of a system or network failure. Panorama in HA provides continuity in the task of centrally administering and monitoring the firewalls to secure your network. This section covers the following topics:

- ▲ [High Availability Overview](#)
- ▲ [Configure a Panorama High Availability Pair](#)
- ▲ [Upgrade Panorama in High Availability](#)

## High Availability Overview

Panorama offers a centralized pane for configuring, monitoring and reporting on the Palo Alto Networks firewalls. Panorama in HA provides redundancy in the central management and reporting functions in your deployment.



To configure Panorama in HA, you require a pair of identical Panorama servers with the following requirements on each:

- **The same form factor**—Must both be hardware-based appliances (M-100 appliances), or virtual appliances. For HA, the M-100 appliances must be in Panorama mode; HA is not supported on a pair of M-100 appliances configured as log collectors.
- **The same Panorama OS version**—Must be running the same version of Panorama in order to synchronize configuration information and maintain parity for a seamless failover.
- **The same set of licenses**—Must purchase and install the same device management capacity license for each Panorama.
- **(Panorama virtual appliance only) Unique serial number**—Must have a unique serial number for each Panorama virtual appliance; if the serial number is duplicated, both instances of Panorama will be placed in a suspended mode until you resolve the issue.

The Panorama servers in the HA configuration are peers and either peer can be used to centrally manage the devices with a [few exceptions](#). The HA peers use the management port to synchronize the configuration elements pushed to the managed devices and to maintain state information. Typically, Panorama HA peers are geographically located in different sites, so you need to make sure that the management port IP address assigned to each peer is routable through your network. HA connectivity uses TCP port 28 with encryption enabled and 28769 when encryption is not enabled.

Each device in the HA pair is assigned a *priority* value. The priority value of *primary* or *secondary* determines which Panorama peer 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*. The active device handles all the configuration changes and pushes them to the managed firewalls; the passive device cannot not make any configuration changes or push configuration to the managed devices. However, either peer can be used to run reports or to perform log queries.

The passive peer is synchronized and ready to transition to the active state, should a path, link, system, or network failure occur on the active device.

## Failover Triggers

When a failure occurs on the active device and the passive device 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 device 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 [Heartbeat Polling and Hello Messages](#).

When the Panorama peers cannot communicate with each other, the active peer monitors whether the devices are still connected to it 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 [Path Monitoring](#).

In addition to the failover triggers listed above, a failover also occurs when the administrator places the device in a suspended state or if *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 device is the preferred active device.

## Heartbeat Polling and Hello Messages

The HA peers use hello message 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 devices are connected and responsive. By default, the interval for the heartbeat is 1000 milliseconds and 8000ms for hello messages.

## Path Monitoring

Path monitoring checks for network connectivity and link state for a specified IP address. The active peer uses ICMP pings to verify that one or more destination IP addresses can be reached. You can, for example, monitor the availability of an 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 5000ms. An IP address is considered unreachable when three consecutive pings (the default value) fail, and a device failure is triggered when any or all of the IP addresses monitored become unreachable. By default, if any one of the IP addresses becomes unreachable, the HA state transitions to *non-functional*.

## Logging Considerations in HA

Setting up Panorama in an HA configuration provides redundancy for log collection. Because the managed devices are connected to both Panorama peers over SSL, when a state change occurs, each Panorama sends a message to the managed devices. The devices are notified of the Panorama HA state and can forward logs accordingly.

The logging options on the hardware-based Panorama and on the Panorama virtual appliance differ.

### Logging Failover on a Panorama Virtual Appliance

On the Panorama virtual appliance, you have the following options:

- **Logging to virtual disk:** By default, the managed devices send logs to both peers in the HA pair; logs are sent as independent log streams to each Panorama HA peer. If a peer becomes unavailable, by default, the managed devices buffer the logs and when the peer reconnects it resumes sending logs from where it had last left off (subject to disk storage capacity and duration of the disconnection).

Logging to a virtual disk provides redundancy in logging, however, the log storage capacity is limited to a maximum of 2TB.



The option to forward logs to the active peer only is configurable. However, log aggregation is not supported across the HA pair. So, if you are logging to virtual disk or to local disk, for monitoring and reporting you must query the Panorama peer that collects the logs from the managed devices.

- **Logging to a Network File Share (NFS):** When configured to use an NFS, only the *active-primary* device 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 devices, but it does not receive the logs and it cannot write to the NFS. In order 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 to Resume NFS Logging](#).

## Logging Failover on an M-100 Appliance

If you are using a pair of M-100 appliances (must be in Panorama mode), the managed devices can send logs to only one peer in the HA pair, either the active or the passive peer. Unlike the virtual Panorama deployment, you cannot configure the devices to send logs to both peers, however, the RAID-enabled disks on the M-100 appliance protect against disk failure and loss of logs.

If you have a distributed log collection set up where the managed devices are sending logs to a dedicated log collector, the Panorama peers in HA will query all the managed log collectors for aggregated log information.
















## Priority and Failover







When a failover occurs, only the state (active or passive) of the device 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 device deployment functions such as license updates or software upgrades on the managed firewalls.
- It cannot log to an NFS until you manually change its priority to primary. (Panorama virtual appliance only)

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)			

Capability	active-primary	passive-primary passive-secondary	active-secondary
Deploy software and licenses			
Export Panorama configuration			

## What Settings are Not Synchronized Between the 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 policies, device group objects and policies, template configuration, and administrative access configuration, are synchronized between the Panorama HA peers.

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 port IP address, FQDN settings, login banner, NTP server, time zone, geographic location, DNS server, permitted IP addresses for accessing Panorama, and SNMP system settings
- NFS partition configuration and all disk quota allocation for logging
- 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 used on Panorama, the same master key must be used to encrypt the private keys and certificates on both peers in the HA pair. If the master keys are different, the HA configuration will not synchronize between the peers.




## Configure a Panorama High Availability Pair

Make sure to review the prerequisites in the [High Availability Overview](#) section before you set up a Panorama HA pair:

- [Set Up High Availability on Panorama](#)
- [Verify Failover](#)
- (Panorama virtual appliance only) [Switch Priority to Resume NFS Logging](#)

### Set Up High Availability on Panorama

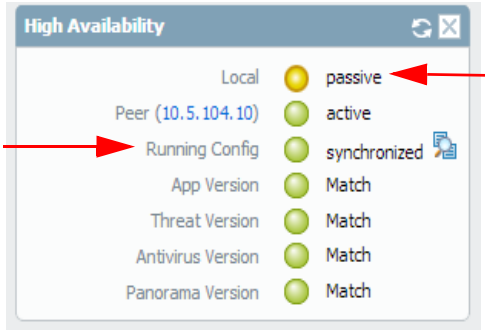
CONNECT AND CONFIGURE THE DEVICES	
<p><b>Step 1</b> Set up connectivity between the MGT ports on the HA peers.</p>	<p>The Panorama peers communicate with each other using the MGT port. Make sure that the IP addresses you assign to the MGT port on the Panorama servers in the HA pair is routable and that the peers can communicate with each other across your network. To set up the MGT port, see <a href="#">Chapter 2, Set Up Panorama</a>.</p>
<p>▲ Pick a device in the pair and complete these tasks:</p>	
<p><b>Step 2</b> Enable HA and enable encryption for the HA connection.</p> 	<ol style="list-style-type: none"> <li>1. Select the <b>Panorama &gt; High Availability</b> tab. Edit the <b>Setup</b> section.</li> <li>2. Select <b>Enable HA</b>.</li> <li>3. Enter the IP address assigned to the peer device in <b>Peer HA IP Address</b>.</li> <li>4. (Optional) To enable encryption, select <b>Encryption Enabled</b> and complete the following tasks:             <ol style="list-style-type: none"> <li>a Select <b>Panorama &gt; Certificate Management &gt; Certificates</b>.</li> <li>b Select <b>Export HA key</b>. Save the HA key to a network location that the peer device can access.</li> <li>c On the peer device, navigate to <b>Panorama &gt; Certificate Management &gt; Certificates</b>, and select <b>Import HA key</b> to browse to the location that you saved the key and import it.</li> </ol> </li> </ol>

CONNECT AND CONFIGURE THE DEVICES (CONTINUED)															
<div>Step 3</div> <div>Set the priority.</div>	<div><div>1.</div><div>In <b>Panorama &gt; High Availability</b>, edit the <b>Election Settings</b> section.</div></div> <div><div>2.</div><div>Define the <b>Device Priority</b> as <b>Primary</b> or <b>Secondary</b>. Make sure to set one peer as primary and the other as secondary.</div></div> <div><div>Note</div><div>If both peers have the same priority setting, the peer with the higher serial number will be placed in a suspended state.</div></div> <div><div>3.</div><div>Define the <b>Preemptive</b> behavior. By default preemption is enabled. The preemption selection— enabled or disabled— must be the same on both peers.</div></div> <div><div>Note</div><div>If you are using an NFS for logging and you have disabled preemption, to resume logging to the NFS see <a href="#">Switch Priority to Resume NFS Logging</a>.</div></div>														
<div>Step 4</div> <div>To configure path monitoring, define one or more path groups.</div> <div>The path group lists the destination IP addresses (nodes) that Panorama must ping to verify network connectivity.</div>	<div><div>1.</div><div>Select <b>Panorama &gt; High Availability</b> tab, and click <b>Add</b> in the Path Group section.</div></div> <div><div>2.</div><div>Enter a <b>Name</b> for the path group.</div></div> <div><div>3.</div><div>Click <b>Add</b> and enter the destination IP addresses that you would like to monitor.</div></div> <div><div>4.</div><div>Select a <b>Failure Condition</b>— <b>all</b> or <b>any</b>— for this group.</div><div><div>• The failure condition <b>any</b> triggers a link monitoring failure if any one of the IP addresses becomes unreachable</div><div>• <b>all</b> triggers a link monitoring failure only when none of the IP addresses can be reached.</div></div><div>The path group is added to the <b>Path Group</b> section.</div><div><table><tr><th></th><th>Name</th><th>Enabled</th><th>Group Failure Condition</th><th>Ping Interval</th><th>Ping Count</th><th>Destination IP</th></tr><tr><td><input checked="" type="checkbox"/></td><td>Access to Corp</td><td><input checked="" type="checkbox"/></td><td>all</td><td>5000</td><td>3</td><td>10.2.1.200 10.2.1.101 10.2.1.102</td></tr></table><div><div> Add</div><div> Delete</div></div></div></div> <div><div>5.</div><div>Repeat steps 1 through 4 above to add more path groups that include the nodes that you want to monitor.</div></div>		Name	Enabled	Group Failure Condition	Ping Interval	Ping Count	Destination IP	<input checked="" type="checkbox"/>	Access to Corp	<input checked="" type="checkbox"/>	all	5000	3	10.2.1.200 10.2.1.101 10.2.1.102
	Name	Enabled	Group Failure Condition	Ping Interval	Ping Count	Destination IP									
<input checked="" type="checkbox"/>	Access to Corp	<input checked="" type="checkbox"/>	all	5000	3	10.2.1.200 10.2.1.101 10.2.1.102									
<div>Step 5</div> <div>(Optional) Select the failure condition for path monitoring on Panorama.</div>	<div>Select <b>Panorama &gt; High Availability</b>, and select a <b>Failure Condition</b> in the Path Monitoring section.</div> <div><div>• The failure condition <b>all</b> triggers a failover only when all monitored path groups fail.</div><div>• The default setting is <b>any</b>; a failover is triggered when any monitored path group fails.</div></div>														
<div>Step 6</div> <div>Save your configuration changes.</div>	<div>Click <b>Commit</b>, select <b>Panorama</b> in the <b>Commit Type</b> option, and click <b>OK</b>.</div>														
<div>Step 7</div> <div>Configure the other Panorama peer.</div>	<div>Repeat <a href="#">Step 2</a> through <a href="#">Step 6</a> on the other peer in the HA pair.</div>														

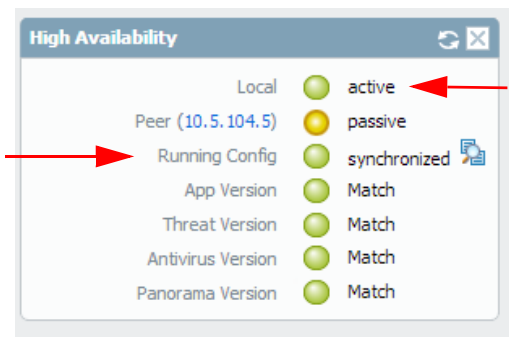
CONNECT AND CONFIGURE THE DEVICES (CONTINUED)

<p>Step 8</p> <p>Verify that the Panorama servers are paired in HA.</p>	<p>After you finish configuring both Panorama servers for HA:</p> <ol style="list-style-type: none"><li>1. Access the <b>Dashboard</b> on each Panorama, and view the <b>High Availability</b> widget.</li><li>2. Confirm that the Panorama servers are paired and synced, as shown below:</li></ol>
---	--

On the passive Panorama: The state of the local peer must display **passive** and the configuration must be **synchronized**.

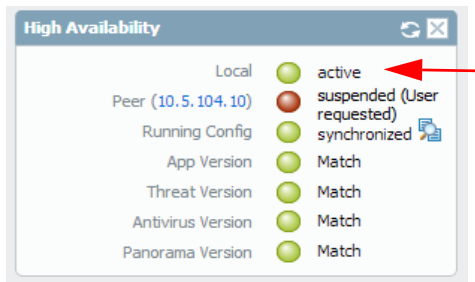


On the active Panorama: The state of the local peer must display **active** and the configuration must be **synchronized**.



## Verify Failover

To test that your HA configuration works properly, trigger a manual failover and verify that the peer transitions states successfully.

VERIFY FAILOVER	
<b>Step 1</b> 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.
<b>Step 2</b> Suspend the active Panorama peer.	Select <b>Panorama &gt; High Availability</b> , and then click the <b>Suspend local Panorama</b> link in the Operational Commands section.
<b>Step 3</b> Verify that the passive Panorama peer has taken over as active.	<p>On the Panorama <b>Dashboard</b>, verify that the state of the passive server changes to <b>active</b> in the <b>High Availability</b> widget. Also verify that the state of the peer changed to suspended.</p> 
<b>Step 4</b> Restore the suspended peer to a functional state. Wait for a couple minutes, and then verify that preemption has occurred, if preemptive is enabled.	<p>On the Panorama you previously suspended:</p> <ol style="list-style-type: none"> <li>1. In the Operational Commands section of the <b>Device &gt; High Availability</b> tab, click the <b>Make local Panorama functional</b> link.</li> <li>2. In the <b>High Availability</b> widget on the <b>Dashboard</b>, confirm that this (Local) Panorama has taken over as the active peer and that the other peer is now in a passive state.</li> </ol>

## Switch Priority to Resume NFS Logging



Support for a Network File Share (NFS) based logging mechanism is only available on the Panorama virtual appliance.

When a Panorama HA pair is configured to use a Network File Share (NFS) based logging mechanism, 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 is *active-secondary*. Although a secondary Panorama peer can actively manage the devices, it cannot receive logs or write to the NFS because it does not own the NFS partition. When the managed device cannot forward logs to the *primary* Panorama peer, the logs are written to the local disk on each device. The devices maintain a pointer for the last set of log entries that were 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 devices, 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.

SWITCH PRIORITY ON PANORAMA	
<b>Step 1</b> Power off the currently passive-primary Panorama.	Select <b>Panorama &gt; Setup &gt; Operations</b> , and click <b>Shutdown Panorama</b> in the Device Operations section.
<b>Step 2</b> Change the priority on the active-secondary Panorama.	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; High Availability</b>.</li> <li>2. Edit the Election Settings section to change the <b>Priority</b> to <b>Primary</b></li> </ol>
<b>Step 3</b> Save your configuration changes.	Click <b>Commit</b> , and select <b>Panorama</b> in the <b>Commit Type</b> option. On commit, you will be prompted to reboot; do not reboot yet.
<b>Step 4</b> Log in to the CLI and change the ownership of the NFS partition to this peer.	In the CLI, enter the following command: <b>request high-availability convert-to-primary</b>
<b>Step 5</b> Reboot Panorama.	Select <b>Panorama &gt; Setup &gt; Operations</b> and click <b>Reboot Panorama</b> in the Device Operations section.  When Panorama reboots, it will dynamically mount the NFS. As the <i>active-primary</i> peer, this Panorama can now write to the NFS.
<b>Step 6</b> Power on the Panorama peer that you powered off in Step1.  This peer will now be in a passive-secondary state.	

## Upgrade Panorama in High Availability

To ensure a seamless failover, the primary and secondary Panorama peers in an HA pair must have the same Panorama version and the same versions of the Applications and Threat databases.


The following example shows how to upgrade an HA pair with an active-primary peer named Primary\_A and the passive-secondary peer named Secondary\_B.

UPGRADE PANORAMA	
<p><b>Step 1</b> Upgrade the Panorama software version on Secondary_B, the passive-secondary peer.</p>	<p>For upgrade instructions, see <a href="#">Install Content and Panorama Software Updates</a>.</p> <p>On upgrade this Panorama will transition to a non-functional state because the OS version does not match that of its peer.</p>
<p><b>Step 2</b> Suspend Primary_A to trigger a failover.</p>	<p>In the <b>Panorama &gt; High Availability</b> tab on Primary_A:</p> <ol style="list-style-type: none"> <li>Click the <b>Suspend local Panorama</b> link in the <b>Operational Commands</b> section to suspend this peer.</li> </ol>  <ol style="list-style-type: none"> <li>Verify that the state displays as suspended; the state displays in the bottom-right corner of the web interface.</li> </ol> <p>Placing Primary_A in a suspended mode triggers a failover and Secondary_B transitions to active-secondary state.</p>
<p><b>Step 3</b> Upgrade the Panorama software version on Primary_A.</p>	<p>For upgrade instructions, see <a href="#">Install Content and Panorama Software Updates</a>.</p> <p>On reboot, Primary_A first transitions to the passive-primary state. Then, because preemption is enabled by default, Primary_A will automatically transition to the active-primary state and Secondary_B will revert to the passive-secondary state.</p> <p>If you have disabled preemption, see <a href="#">Restore the Primary Server to the Active State</a> for restoring Primary_A to the active state.</p>
<p><b>Step 4</b> Verify that the Panorama software version and other content database versions are the same on both peers.</p>	<p>On the <b>Dashboard</b> of each Panorama peer, verify that the Panorama software version, the Threat version, and the Application versions are a <b>Match</b> and that the running configuration is <b>synchronized</b> with the peer.</p>

## Restore the Primary Server 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 in to a suspended state, it transfers the candidate configuration to the passive device so that all your uncommitted configuration changes are saved and can be accessed on the other peer.

SUSPEND PANORAMA	
<b>Step 1</b> Suspend Panorama.	<ol style="list-style-type: none"> <li>1. Log in to the Panorama peer that you want to place in a suspended state.</li> <li>2. Select <b>Panorama &gt; High Availability</b>, and click the <b>Suspend local Panorama</b> link in the Operational Commands section.</li> </ol>
<b>Step 2</b> Verify that the status displays that the device was suspended at user request.	<p>On the <b>High Availability</b> widget on the <b>Dashboard</b>, verify that the state displays as suspended.</p>  <p>A failover is triggered when you suspend a peer, and the other Panorama takes over as the active peer.</p>
RESTORE PANORAMA TO A FUNCTIONAL STATE	
To restore the suspended Panorama to a functional state.	<ol style="list-style-type: none"> <li>1. Click the <b>Make local Panorama functional</b> link on the Operational Commands section of the <b>Panorama &gt; High Availability</b> tab.</li> <li>2. In the <b>High Availability</b> widget on the <b>Dashboard</b>, confirm that the device has transitioned to either the active or passive state.</li> </ol>







## 6 Administer Panorama

---

This section describes how to administer and maintain Panorama. It includes the following topics:

- ▲ Manage Configuration Backups
- ▲ Compare Changes in Configuration
- ▲ Restrict Access to Configuration Changes
- ▲ Add Custom Logos
- ▲ View Task Completion History
- ▲ Reallocate Log Storage Quota
- ▲ Monitor Panorama
- ▲ Reboot or Shutdown Panorama
- ▲ Generate Diagnostic Files
- ▲ Configure Password Profiles and Password Complexity
- ▲ Replace the Virtual Disk on a Panorama Virtual Appliance



For instructions on completing initial set up including defining network access settings, licensing, upgrading the Panorama software version, and setting up administrative access to Panorama, see Chapter 2, Set Up Panorama.

## Manage Configuration Backups

A configuration backup is a snapshot of the system configuration. In case of a system failure or a misconfiguration, a configuration backup allows you to restore Panorama to a previously saved version of the configuration. On Panorama, you can manage configuration backups of the managed firewalls and that of Panorama:

- **Manage configuration backups of the managed devices:** Panorama automatically saves every configuration change that is committed to a managed firewall running PAN-OS version 5.0 or later. By default, Panorama stores up to 100 versions for each device. This value is configurable.
- **Manage Panorama configuration backups:** You can manually export the running configuration of Panorama, as required.
- **Export a configuration file package:** In addition to its own running configuration, Panorama saves a backup of the running configuration from all managed devices. You can generate a gzip package of the latest version of the configuration backup of Panorama and that of each managed device either on-demand or schedule an export using the **Scheduled Config Export** capability. The package can be scheduled for daily delivery to an FTP server or an Secure Copy (SCP) server; the files in the package are in an XML format, and each file name references the device serial number for easy identification.

You can perform the following tasks to manage configuration backups:

- ▲ [Schedule Export of Configuration Files](#)
- ▲ [Manage Panorama Configuration Backups](#)
- ▲ [Configure the Number of Backups Stored on Panorama](#)
- ▲ [Load a Configuration Backup on a Managed Device](#)

## Schedule Export of Configuration Files

Use these instructions to schedule the export of the configuration file package that contains the backup of the running configuration of Panorama and the managed devices, at a specified time everyday. Superuser privileges are required to configure the export.

---

### SCHEDULE THE EXPORT OF CONFIGURATION FILES AT A SPECIFIED TIME DAILY

---

1. Select **Panorama > Scheduled Configuration Export**.
2. Click **Add**, and enter a **Name** and **Description** for the file export process.
3. Select **Enable** to allow the configuration file export.
4. Enter a time or select one from the drop-down for daily export of the configuration files. A 24-hour clock is used.
5. Select the protocol.
6. Enter the details for accessing the server. Provide the hostname or IP address, port, path for uploading the file, and authentication credentials.
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. The connection is not established until the host key is accepted.

If Panorama can successfully connect to the SCP server, it creates and uploads the test file named **ssh-export-test.txt**.

8. Save the changes. Click **Commit**, select **Panorama** as the **Commit Type** and click **OK**.
-

## Manage Panorama Configuration Backups

Use these instructions to validate, revert, save, load, export, or import a Panorama configuration version.

### MANAGE BACKUPS: VALIDATE, REVERT, SAVE, LOAD, EXPORT OR IMPORT

1. Select **Panorama > Setup > Operations**.
2. In the Configuration Management section, choose from the following options:
  - **Validate candidate Panorama configuration**—Verifies that the candidate configuration has no errors; validating the configuration file allows you to resolve errors before you commit the changes.
  - **Revert to last saved Panorama configuration**—Overwrites the current candidate configuration and restores the last saved candidate configuration from disk.
  - **Revert to running Panorama configuration**—Reverts all changes saved to the candidate configuration; it effectively allows you to undo all configuration changes that were made since the last commit operation.
  - **Save named Panorama configuration snapshot**—Saves the candidate configuration to a file. Enter a file name or select an existing file to overwrite. Note that the current active configuration file (running-config.xml) cannot be overwritten.
  - **Save candidate Panorama configuration**—Saves the candidate configuration to disk; it is the same as using the **Save** link at the top of the page to save the changes to the candidate configuration file.
  - **Load Panorama configuration version**—Loads a configuration file from a list of previously committed versions.
  - **Load named Panorama configuration snapshot**—Loads a selected candidate configuration; you can select a previously imported or saved configuration. The current candidate configuration is overwritten.
  - **Export named Panorama configuration snapshot**—Exports the active configuration (running-config.xml) or a previously saved or imported configuration. Select the configuration file to be exported. You can open the file and/or save it in any network location.
  - **Export Panorama configuration version**—Exports a previously committed version of the configuration file. Select the version to export.
  - **Export Panorama and devices config bundle**—This option is used to manually generate and export the latest version of the configuration backup of Panorama and that of each managed device. To automate the process of creating and exporting the configuration bundle daily to an SCP or FTP server, see [Schedule Export of Configuration Files](#).
  - **Import named Panorama configuration snapshot**—Imports a previously exported configuration file. Click **Browse** to locate the saved file and click **OK** to import.

## Configure the Number of Backups Stored on Panorama

Specify the number of Panorama backups that are stored.

### CONFIGURE THE NUMBER OF BACKUPS STORED ON PANORAMA

1. Select **Panorama > Setup > Management**, and click the edit button in the Logging and Reporting Settings section.

Number of Versions for Config Audit	100
Number of Versions for Config Backups	100

2. Enter a value between 1 and 1048576. The default is 100.
3. Save the changes. Click **Commit**, select **Panorama** as the **Commit Type** and click **OK**.

## Load a Configuration Backup on a Managed Device

Use Panorama to load a configuration backup on a managed device. You can choose to revert to a previously saved or committed configuration on the device. Panorama pushes the selected version to the managed device, and the current candidate configuration on the device is overwritten.

---

### LOAD A CONFIGURATION BACKUP ON A MANAGED DEVICE

---

1. Select **Panorama > Managed Devices**.
  2. Select the **Manage...** link in the **Backups** column.
  3. Select from the **Saved Configurations** or the **Committed Configurations**.
    - Click the link in the **Version** column to view the contents of the selected version.
    - Click **Load** to load a chosen configuration version.
  4. Save the changes. Click **Commit** and select **Panorama** as the **Commit Type**.
-

## Compare Changes in Configuration

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 allows you to:

- Preview the changes in configuration 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.

COMPARE CHANGES IN CONFIGURATION	
<ul style="list-style-type: none"> <li>• View and compare configuration files. To easily compare versions, the changes are highlighted:</li> </ul> <div> <span>Added</span> <span>Modified</span> <span>Deleted</span> </div>	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Config Audit</b>.</li> <li>2. For each drop-down, select a configuration for the comparison.</li> <li>3. Select the number of lines that you want to include for <b>Context</b>, and click <b>Go</b>.</li> </ol>
<ul style="list-style-type: none"> <li>• Configure the number of versions stored on the Panorama for a configuration audit.</li> </ul>	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Setup &gt; Management</b>, and click the edit icon in the Logging and Reporting section.</li> <li>2. Enter a value between 1 and 1048576 in <b>Number of Versions for Config Audit</b>. The default is 100.</li> <li>3. Save the changes. Click <b>Commit</b> and select <b>Panorama</b> as the <b>Commit Type</b>.</li> </ol>

PREVIEW CHANGES IN CONFIGURATION	
<ul style="list-style-type: none"> <li>• View and compare configuration files before you commit your changes.</li> </ul>	<ol style="list-style-type: none"> <li>1. Select <b>Commit</b>.</li> <li>2. Select <b>Preview Changes</b> and choose the number of lines of context that you would like to see.</li> <li>3. Click <b>OK</b>.</li> </ol>

## Restrict Access to Configuration Changes

Use locks to prevent multiple administrative users from making configuration changes or committing changes on Panorama, shared policies, or to selected Templates and/or Device Groups.

- ▲ [Types of Locks](#)
- ▲ [Locations for Taking a Lock](#)
- ▲ [Take a Lock](#)
- ▲ [View Current Lock Holders](#)
- ▲ [Enable Automatic Acquisition of the Commit Lock](#)
- ▲ [Remove a Lock](#)

### Types of Locks

The two types of locks available are:

- **Config Lock**—Blocks other administrators from making changes to the configuration. This type of lock can be set globally or for a virtual system. It can be removed only by the administrator who set it or by a superuser. The configuration lock is not released automatically.
- **Commit Lock**—Blocks other administrators from committing changes until all of the locks have been released. The commit lock ensures that partial configuration changes are not inadvertently committed to the device or to Panorama when two administrators are making changes at the same time and the first administrator finishes and commits changes before the second administrator has finished. The lock is released automatically when the administrator who applied the lock commits the changes; the lock can be removed manually by the administrator who took the lock or by the superuser.

If a commit lock is held on a device, and an administrator commits configuration changes or shared policies to a template or device group that includes that device, the commit will fail with an error message indicating that there is an outstanding lock on a device.



Read-only administrators who cannot make configuration changes to the device or Panorama will not be able to take either lock.

Role-based administrators who cannot commit changes can take the config lock and save the changes to the candidate configuration. They cannot, however, commit the changes themselves. Because they cannot commit the changes, the lock is not automatically released on commit; the administrator must manually remove the config lock after making the required changes.


### Locations for Taking a Lock

The administrator can take a lock for any of the following categories, or *locations*:

- **Device Group**— Restricts changes to the selected Device Group

- **Template**—Restricts changes to the devices included in the selected Template
- **Shared**—Restricts changes to the shared policies
- **Panorama**—Restricts access to changes on Panorama

## Take a Lock

TAKE A LOCK	
Step 1	Select the lock icon  on the top right corner of the web interface.
Step 2	Select <b>Take Lock</b> .
Step 3	Based on your role/permissions, select <b>Commit</b> or <b>Config</b> as <b>Type</b> .
Step 4	Select the category for which you want to take the lock.
Step 5	As a best practice, add a <b>Comment</b> to describe the reasons for taking the lock.
Step 6	Click <b>OK</b> .

## View Current Lock Holders

Before changing a particular area of the configuration, check whether another administrator has taken the lock for the area.

VIEW LOCK HOLDERS	
•	Select the lock icon on the top right corner of the web interface and review the details. The lock icon displays the total number of locks taken. It also includes information on the username of the lock holder, type of lock, the category in which the lock is held, when it was taken, the last activity by the administrator, and whether or not the administrator is still logged in.

## Enable Automatic Acquisition of the Commit Lock

By default, you must manually take a lock before you start making changes on Panorama. If you would like to enable automatic acquisition of the commit lock, use the following procedure.

ACQUIRE AN AUTOMATIC COMMIT LOCK ON PANORAMA	
Step 1	Select <b>Panorama &gt; Setup &gt; Management</b> tab and click the edit icon in the General Settings section.
Step 2	Select the check box for <b>Automatically Acquire Commit Lock</b> .
Step 3	Click <b>OK</b> .
Step 4	To save the changes, click <b>Commit</b> and select <b>Panorama</b> as the <b>Commit Type</b> .




## Remove a Lock

---

### REMOVE A LOCK

---

**Step 1** Select the lock icon  on the top right corner of the web interface.

**Step 2** Select the lock that you want to release and click **Remove Lock**.

**Note** Unless you are a superuser, you can only remove the lock that you have previously taken.

**Step 3** Click **OK**.

---

## Add Custom Logos

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 supported include the following: .jpg, .gif, 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.

---

### ADD CUSTOM LOGOS


---

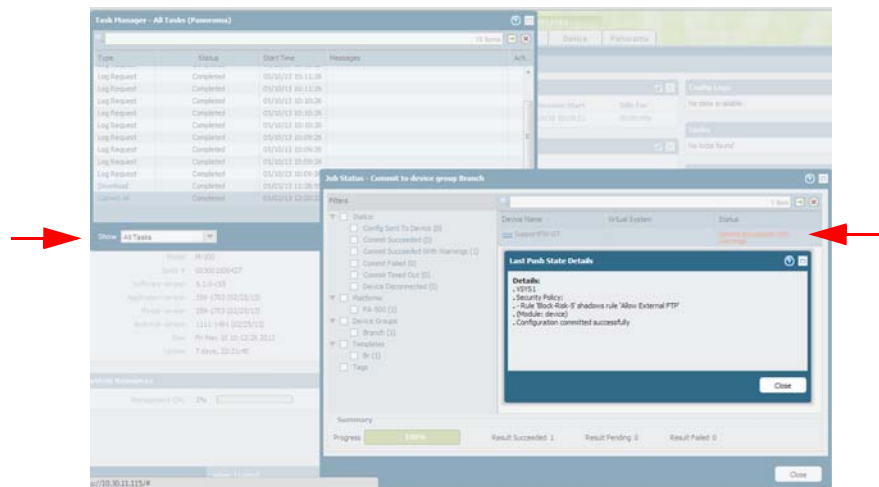
1. Select **Panorama > Setup > Operations**.
  2. Click **Custom Logos** in the Miscellaneous section.
  3. Click the Upload logo icon , and select an image for any of the following options: the login screen, the left corner of the main user interface, the PDF report title page and the PDF report footer.
  4. Click **Open** to add the image. To preview the image, click the preview logo icon .
  5. (Optional) To clear the green background header on the Panorama web interface, select the check box for **Remove Panorama background header**.
  6. Click **Close** to save your changes.
  7. Click **Commit** and select **Panorama** as the **Commit Type**.
-

## View Task Completion History

The historical data on all tasks or currently-running tasks on Panorama can be viewed using the **Tasks** link on the Panorama web interface. It displays information on the success or failure of the event and lists errors, if any.

### VIEW TASK HISTORY

1. Click **Tasks** . The link displays on the bottom right corner of the web interface.
2. Select the list of tasks to review. By default **All Tasks** are displayed. You can filter by **All** or **Running** tasks and select **Jobs, Reports, or Log Requests**.



- **Jobs:** Lists commits, auto commits, downloads and installs for software and dynamic updates performed on locally on Panorama or centrally pushed to the managed devices from Panorama. Each job is a link; click the link in the Type column to view details on the devices, status, and review errors, if any.
- **Reports:** Displays the status and start time for scheduled reports.
- **Log Requests:** Lists the log queries triggered from the **Monitor > Log Viewer** tab or the **Dashboard**. For example, to display the logs in the URL Filtering widget or the Data Filtering widget on the Dashboard, log requests are generated on Panorama.

# Reallocate Log Storage Quota

To redistribute the available log storage capacity on Panorama, you can increase or decrease the log storage quota for each log type. The log quota reallocation process is different on the Panorama virtual appliance and the M-100 appliance as follows:

- **On the Panorama virtual appliance:** All logs are written to the storage space that is assigned to the server: the default 10 GB disk that is created on install, or the virtual disk added to the server, or the NFS partition that is mounted on Panorama.
- **On the M-100 appliance:** Logs are saved to two locations: the internal SSD and the RAID-enabled disks. The internal SSD on the M-100 appliance is used to store the configuration logs, the system logs, and the application statistics that Panorama automatically receives at 15 minute intervals from all managed devices. All the other logs are stored on the RAID-enabled disks. To reallocate the storage capacity for logs stored on the RAID disks, you must modify the Collector Group configuration.

Use the following instructions for:

- Reallocating log storage quota on the Panorama virtual appliance
- Reallocating log storage capacity for the system logs, configuration logs and the Application statistics (App Stats) on the M-100 appliance
- Reallocating storage capacity for the logs stored on the RAID-enabled disks

## REALLOCATE LOG STORAGE QUOTA ON THE PANORAMA VIRTUAL APPLIANCE AND THE M-100 APPLIANCE

**Step 1** Apportion the log storage capacity among the various log types.



1. Select **Panorama > Setup > Management**.
  2. Select the edit icon on the **Logging and Reporting Settings** section of the **Management** tab.
  3. Modify the **Quota (%)** for the log types to which you want to add or reduce storage space.
- As you change the values, the screen refreshes to display the corresponding number value (GB/MB) for the percentage allocated based on the total storage on the virtual disk/NFS/HDD, as applicable.

**Note** On the M-100 appliance, you can allocate the available capacity among the **Config**, **System** and **App Stats** logs; all the other logs are written to the RAID-enabled disks and are not stored on the HDD.

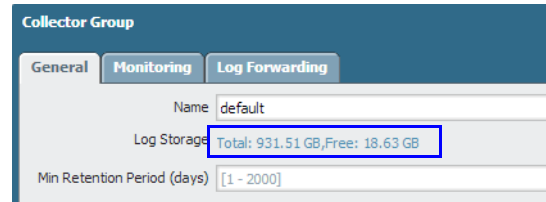
**REALLOCATE THE LOG STORAGE QUOTA FOR THE RAID-ENABLED DISKS**

**Step 1** Allocate the percentage of storage capacity for each log type on the M-100 appliance.

If the Log Storage capacity displays as 0MB, you may not have added the Log Collector to the Collector Group. Complete [Step 2](#) and then come back to this task.

If it still reads as 0MB, verify that the disk pairs are enabled for logging and have committed the changes to the Collector Group. See [Step 6](#) in the [Add a Log Collector to Panorama](#) section.

1. Select **Panorama > Collector Groups**, and click the link for the Collector Group.
2. Click the link for the **Log Storage** capacity of the Collector Group.



The screenshot shows the 'Collector Group' configuration page with the 'Log Forwarding' tab selected. The 'Log Storage' section displays 'Total: 931.51 GB, Free: 18.63 GB'. The 'Min Retention Period (days)' is set to '[1 - 2000]'.

3. Modify the **Quota** allocated for each log type.  
As you change the value, the screen refreshes to display the corresponding number value (GB/MB) for the percentage allocated based on the total storage in the Collector Group.
4. (Optional) Click **Restore Defaults** to undo your changes and reset the quotas to factory defaults, if necessary.

## Monitor Panorama

You can configure Panorama to send notifications if a system event occurs or anytime a configuration change is made. By default, every configuration change is logged to the configuration log. On the system log, each event has a severity level associated with it. The level indicates the urgency and the impact of the event, and you can choose to record all or selected system events depending on the severity levels that you want to monitor.



This section covers Panorama logs only. For information on forwarding logs from the managed devices, see [Configure the Firewalls to Forward Logs to Panorama](#).

- **Config Logs**—Enable forwarding of Configuration logs by specifying a server profile in the log settings configuration (**Panorama > Log Settings > Config Logs**).
- **System Logs**—Enable forwarding of System logs by specifying a server profile in the log settings configuration (**Panorama > Log Settings > System Logs**). Select a server profile for each severity level you want to forward. The following table summarizes the system log severity levels:

Severity	Description
<b>Critical</b>	Indicates a failure and signals the need for immediate attention, such as a hardware failure, including HA failover and link failures.
<b>High</b>	Serious issues that will impair the operation of the system, including disconnection of a Log Collector or a commit failure.
<b>Medium</b>	Mid-level notifications, such as antivirus package upgrades, or a Collector Group commit.
<b>Low</b>	Minor severity notifications, such as user password changes.
<b>Informational</b>	Notification events such as log in/log off, any configuration change, authentication success and failure notifications, commit success, and all other events not covered by the other severity levels.

The Configuration logs and System logs are stored on the HDD of the M-100 appliance; on the Panorama virtual appliance they are stored on the assigned storage volume. If you need longer-term storage of logs, for auditing, you can also configure Panorama to forward the logs to a syslog server.

To monitor Panorama, you can either periodically view the logs on Panorama or configure SNMP traps and/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 require your attention.



Panorama only forwards its own locally generated system and config logs. You cannot use Panorama to forward any logs sent from the managed devices to an external SIEM or syslog server. You cannot, for example, use Panorama to forward Traffic or Threat logs to an external storage server.

To set up email alerts and SNMP access, see the following tasks:

- ▲ [Set up Email Alerts](#)
- ▲ [Set up SNMP Access](#)

## Set up Email Alerts

SET UP EMAIL ALERTS	
<p><b>Step 1.</b> Create a server profile for your email server.</p>	<ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Server Profiles &gt; Email</b>.</li> <li>2. Click <b>Add</b> and then enter a <b>Name</b> for the profile.</li> <li>3. Click <b>Add</b> to add a new email server entry and enter the information required to connect to the Simple Mail Transport Protocol (SMTP) server and send email (you can add up to four email servers to the profile): <ul style="list-style-type: none"> <li>• <b>Server</b>—Name to identify the mail server (1-31 characters). This field is just a label and does not have to be the host name of an existing SMTP server.</li> <li>• <b>Display Name</b>—The name to show in the From field of the email.</li> <li>• <b>From</b>—The email address where notification emails will be sent from.</li> <li>• <b>To</b>—The email address to which notification emails will be sent.</li> <li>• <b>Additional Recipient(s)</b>—To send notifications to a second account, enter the additional address here.</li> <li>• <b>Gateway</b>—The IP address or host name of the SMTP gateway to use to send the emails.</li> </ul> </li> <li>4. Click <b>OK</b> to save the server profile.</li> </ol>
<p><b>Step 2</b> (Optional) Customize the format of the logs Panorama sends.</p>	<p>Select the <b>Custom Log Format</b> tab. For details on how to create custom formats for the various log types, refer to the <a href="#">Common Event Format Configuration Guide</a>.</p>
<p><b>Step 3</b> Save the server profile and commit your changes.</p>	<ol style="list-style-type: none"> <li>1. Click <b>OK</b> to save the profile.</li> <li>2. Click <b>Commit</b>, and select <b>Panorama</b> as the <b>Commit Type</b>.</li> </ol>
<p><b>Step 4</b> Enable email notification for specific events in the system and config logs.</p>	<ol style="list-style-type: none"> <li>1. Enable email notification. <ul style="list-style-type: none"> <li>• For system events: <ol style="list-style-type: none"> <li>a. Select <b>Panorama &gt; Log Settings &gt; System</b>.</li> <li>b. Click the link for each severity level for which to enable notification, and then select the email server profile you created in <a href="#">Step 1</a> from the <b>Email</b> drop-down.</li> </ol> </li> <li>• For configuration changes: <ol style="list-style-type: none"> <li>a. Select <b>Panorama &gt; Log Settings &gt; Config</b>.</li> <li>b. Click the edit icon in the Log Settings - Config section and then select the email server profile you created in <a href="#">Step 1</a> from the <b>Email</b> drop-down.</li> </ol> </li> </ul> </li> <li>2. Click <b>Commit</b> and select <b>Panorama</b> as the <b>Commit Type</b>.</li> </ol>

## Set up SNMP Access

The SNMP protocol allows access to specific Object Identifiers (OIDs) or ranges of OIDs contained in the Palo Alto Networks MIBs from an SNMP Management station. It can be used to query the state of Panorama (SNMP GETs) and to trigger an alert (SNMP TRAPs) when an event occurs. Panorama supports SNMP v2c and v3. For a complete list see [Palo Alto Networks MIBs](#).

SNMP Traps alert on a failure or a change, such as a system fan failure, addition of disk drives, or HA failovers. Traps are initiated by Panorama and sent to the SNMP manager; they are not sent on a regular schedule.

SNMP GETs allow pro-active monitoring. You can, for example, poll Panorama for trending graphs that help identify potential system issues before a fault occurs for the following issues:

- Monitor the incoming log rate on an M-100 appliance or the capacity of the logging disks on the appliance to determine if a log collector is close to maximum capacity. This information will help you evaluate whether you need to expand log storage capacity or add additional log collectors.
- Monitor system information such as high availability mode and state of Panorama, currently installed content update versions—antivirus version, application and threat database version, and/or Panorama version.



**SET UP SNMP**

**Step 1** Configure the management interface to listen for the SNMP service.

1. Select the **Panorama > Setup > Management**.
2. In the Management Interface Settings section, verify that SNMP is enabled in **Services**. If SNMP is not enabled, click the Edit icon in the Management Interface Settings section, and select the check box for the **SNMP** service, and click **OK** to save the changes.

**Step 2** Configure Panorama for SNMP monitoring.

This screen shot is for SNMP v3.

**SNMP Setup**

Physical Location:

Contact:

☒ Use Event-specific Trap Definitions

Version: ☐ V2c ☒ V3

Name	View
<input checked="" type="checkbox"/> 1.3.6.1.4.1.25461.2.3.30	andrew: 1.3.6.1: include: 0x00

Users	View	Auth Password	Priv Password
<input checked="" type="checkbox"/> andrew	1.3.6.1.4.1.2546...	*****	*****

1. Select **Panorama > Setup > Operations**.
2. In the Miscellaneous section, select **SNMP Setup**.
3. Enter a text string to specify the physical **Location** of Panorama.
4. Add the email address of one or more administrative **Contact**.
5. Select the SNMP **Version** and then enter the configuration details as follows (depending on which SNMP version you are using) and then click **OK**:
  - **V2c**—Enter the **SNMP Community String** that will allow the SNMP manager access to the SNMP agent on Panorama. The default value is **public**, however because this is a well-known community string, it is a best practice to use a value that is not easily guessed.
  - **V3**—You must create at least one View and one User in order to use SNMPv3. The view specifies which management information the manager has access to. If you want to allow access to all management information, just enter the top-level **OID** of .1.3.6.1 and specify the **Option** as **include** (you can also create views that exclude certain objects). Use **0xf0** as the **Mask**. Then when creating a user, select the **View** you just created and specify the **Auth Password** and **Priv Password**. The authentication settings (the community string for V2c or the username and passwords for V3) configured on Panorama must match the values configured on the SNMP manager.
6. Click **OK** to save the settings.
7. Click **Commit**, and select **Panorama** as your **Commit Type** to save the changes to the running configuration.

**SET UP SNMP (CONTINUED)**

<p><b>Step 3</b> Create a Server Profile that contains the information for connecting and authenticating to the SNMP manager(s).</p> <ol style="list-style-type: none"> <li>1. Select <b>Panorama &gt; Server Profiles &gt; SNMP Trap</b>.</li> <li>2. Click <b>Add</b> and then enter a <b>Name</b> for the profile.</li> <li>3. Specify the version of SNMP you are using (V2c or V3).</li> <li>4. Click <b>Add</b> to add a new <b>SNMP Trap Receiver</b> entry (you can add up to four trap receivers per server profile). The required values depend on whether you are using SNMP V2c or V3 as follows: <ul style="list-style-type: none"> <li><b>On SNMP V2c</b> <ul style="list-style-type: none"> <li>• <b>Server</b>—Name to identify the SNMP manager (1-31 characters). This field is just a label and does not have to be the hostname of an existing SNMP server.</li> <li>• <b>Manager</b>—The IP address of the SNMP manager to which to send traps.</li> <li>• <b>Community</b>—The community string required to authenticate to the SNMP manager.</li> </ul> </li> <li><b>On SNMP V3</b> <ul style="list-style-type: none"> <li>• <b>Server</b>—Name to identify the SNMP manager (1-31 characters). This field is just a label and does not have to be the hostname of an existing SNMP server.</li> <li>• <b>Manager</b>—The IP address of the SNMP manager to which to send traps.</li> <li>• <b>User</b>—The username required to authenticate to the SNMP manager.</li> <li>• <b>EngineID</b>—The engine ID of Panorama. This is a hexadecimal value from 5 to 64 bytes with a 0x prefix. Each Panorama has a unique engine ID. In order to find out the engine ID, configure the server for SNMP v3 and send a GET message from the SNMP Manager or MIB browser to Panorama.</li> <li>• <b>Auth Password</b>—The password to be used for authNoPriv level messages to the SNMP manager. This password will be hashed using Secure Hash Algorithm (SHA-1), but will not be encrypted.</li> <li>• <b>Priv Password</b>—The password to be used for authPriv level messages to the SNMP manager. This password will be hashed using SHA and will be encrypted using Advanced Encryption Standard (AES 128).</li> </ul> </li> </ul> </li> <li>5. Click <b>OK</b> to save the server profile.</li> </ol>	
<p><b>Step 4</b> Enable SNMP Traps for config log and syslog events.</p>	<ul style="list-style-type: none"> <li>• For system events: <ol style="list-style-type: none"> <li>a. Select <b>Panorama &gt; Log Settings &gt; System</b>.</li> <li>b. Click the link for each severity level for which to enable notification, and then select the server profile you created in <a href="#">Step 3</a> from the <b>SNMP Trap</b> drop-down.</li> </ol> </li> <li>• For configuration changes: <ol style="list-style-type: none"> <li>a. Select <b>Panorama &gt; Log Settings &gt; Config</b>.</li> <li>b. Click the edit icon in the Log Settings - Config section, and then select the server profile you created in <a href="#">Step 3</a> from the <b>SNMP Trap</b> drop-down.</li> </ol> </li> </ul>
<p><b>Step 5</b> Save your changes.</p>	<p>Click <b>Commit</b>, and select <b>Panorama</b> as the <b>Commit Type</b>.</p>
<p><b>Step 6</b> Enable the SNMP manager to interpret an SNMP trap.</p>	<p>To interpret a trap sent by Panorama, you must load the <a href="#">PAN-OS MIB files</a> into your SNMP management software and, if necessary, compile them. The compiled MIBs allow the SNMP Manager to map the object identifier (OID) to the event definition that the trap defines.</p> <p>Refer to the documentation for your SNMP manager for specific instructions on how to do this.</p>

**SET UP SNMP (CONTINUED)**

<b>Step 7</b> Identify the statistics to monitor.	Using a MIB browser, walk the PAN-OS MIB files to identify the object identifiers (OIDs) that correspond to the statistics you want to monitor. For example, suppose you want to monitor on the log collection rate on the M-100 appliance. Using a MIB browser you will see that this statistic corresponds to OID 1.3.6.1.4.1.25461.2.3.16.1.1.0 in the PAN-LC-MIB.
<b>Step 8</b> Configure the SNMP management software to monitor the OIDs you are interested in.	Refer to the documentation for your SNMP manager for specific instructions on how to do this.

## Reboot or Shutdown 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.

---

### REBOOT/SHUT DOWN

---

**Step 1** Select **Panorama > Setup > Operations**.

**Step 2** In the Device Operations section, pick one of the following options:

- To reboot: Select **Reboot Panorama**.
  - To shutdown: Select **Shutdown Panorama**.
-

## Generate Diagnostic Files

Diagnostic files aid in monitoring system activity and in discerning potential causes for issues on Panorama. In order to assist Palo Alto Networks Technical Support in troubleshooting an issue, the support representative might request diagnostic files—tech support file or a stats dump file. The following procedure describes how a diagnostic file and upload it to your support case.

---

### GENERATE DIAGNOSTIC FILES

---

1. Select **Panorama > Support**.
    - Click **Generate Tech Support File**.
    - Click **Generate Stats Dump File**.
  2. Download and save the file(s) to your computer.
  3. Upload the file(s) to your case on the Support Portal.
-

## Configure Password Profiles and Password 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 device-wide and apply to all passwords.

To enforce periodic password updates, create a password profile that defines a validity period for passwords.

## PASSWORD PROFILES AND COMPLEXITY SETTINGS

**Step 1** Configure minimum password complexity settings.

1. Select **Panorama > Setup > Management** and then click the Edit icon in 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)**.

Minimum Length: 10

Minimum Uppercase Letters: 1

Minimum Lowercase Letters: 1

Minimum Numeric Letters: 2

Minimum Special Characters: 1

Block Repeated Characters: 0

☒ 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.

New Password Differs By Characters: 4

☐ Require Password Change on First Login

Prevent Password Reuse Limit: 1

Block Password Change Period (days): 0

Required Password Change Period (days): 90

Expiration Warning Period (days): 15

Post Expiration Admin Login Count: 2

Post Expiration Grace Period (days): 2

Functionality requirements can be overridden by password profiles

**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 then click **Add**.
2. Enter a **Name** for the password profile and define the following:
  - a. **Required Password Change Period:** Frequency, in days, at which the passwords must be changed.
  - b. **Expiration Warning Period:** Number of days before expiration that the administrator will receive a password reminder.
  - c. **Post Expiration Grace Period:** Number of days that the administrator can still log in to the system after the password expires.
  - d. **Post Expiration Admin Login Count:** Number of times that the administrator can log in to the system after the password has expired.

## Replace the Virtual Disk on a Panorama Virtual Appliance

A virtual disk cannot be resized after it is added to a Panorama virtual appliance. Because the Panorama virtual appliance only allows only one log storage location, if you need to add disk space for logging (or lessen, if you allocated more space), you must replace the virtual disk on the ESX(i) server to adjust the log storage capacity.

Perform the following tasks on the ESX(i) server to replace the virtual disk assigned to Panorama:

### REPLACE THE VIRTUAL DISK ASSIGNED TO THE PANORAMA VIRTUAL APPLIANCE

**Step 1** Export the logs before detaching the virtual disk from the Panorama virtual appliance. The logs on the disk will no longer be accessible after the disk is detached.

1. Access the CLI on the Panorama virtual appliance and check the current disk usage:

```
admin@Panorama> show system logdb-quota
```

2. To export the logs, enter the command:

```
admin@Panorama> scp export logdb to <user account>@<IP of SCP server>:  
<directory path with destination filename>
```

For example:

```
admin@Panorama> scp export logdb to sabel@10.236.10.30:/Panorama/log_file_exportMay2013
```

**Note** You must specify a filename; a *.tar* file with that filename is saved to the SCP server. Because the files are compressed during the export process, the size of the exported file will be smaller than the size on disk.

**Step 2** Power off the Panorama virtual appliance.

**Step 3** Edit the settings on the Panorama virtual appliance to add a new virtual disk.

**Step 4** Create a new virtual disk with the desired capacity. The virtual disk type must be IDE and the maximum capacity is 2TB.



**Step 5** Remove the virtual disk you want to replace.

**Step 6** Power on the Panorama virtual appliance. The reboot process might take several minutes and a message `cache data unavailable` will display on screen.

**Step 7** Log in to the Panorama virtual appliance.

Select **Panorama > Setup > Management**, verify that the modified log storage capacity is displayed accurately in the Logging and Reporting section.

Logging and Reporting Settings

Log Storage Total: 101.30 GB



---

**REPLACE THE VIRTUAL DISK ASSIGNED TO THE PANORAMA VIRTUAL APPLIANCE**

---

**Step 8** Import the logs into the new virtual disk on Panorama. To import the files in to the new virtual disk, enter the following command from the CLI:

```
admin@Panorama> scp import logdb from <user account>@<IP of SCP server>:  
<directory path with destination filename>
```

---



## 7 Troubleshooting

---

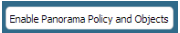
This section addresses the following issues:


- ▲ Why does the Template commit fail?
- ▲ Why is Panorama running a File System Integrity check?
- ▲ Is there a separate connection for forwarding logs to Panorama?
- ▲ Why does the log storage capacity for the Collector Group read 0 MB?
- ▲ Why is Panorama in a suspended state?
- ▲ Where do I view task completion status?

### Why does the Template commit fail?

A template commit could fail because of the following reasons:

- Capability mismatch: When configuring a template, the following options are available: multiple virtual systems capability, VPN mode, and operational mode.
  - If the check box for multiple virtual systems capability is selected, a template commit failure will occur when you push the template to devices that are not capable of or enabled for multiple virtual systems functionality.  
To resolve the error, edit the template in the **Panorama > Templates** tab, and clear the check box for **Virtual systems**.
  - If VPN-related configuration options are pushed to devices that are hard coded to disallow VPN configuration.  
To resolve the error, edit the template in the **Panorama > Templates** tab, and enable the check box for **VPN Disable Mode**.
  - If the operational mode enabled on the device and that on the template are different. For example, if the managed device is enabled for FIPS mode and the template is defined for normal mode.  
To resolve the error, edit the template in the **Panorama > Templates** tab and verify that the **Operational mode** selection is correct.
- The managed device is not enabled for receiving template and device group changes from Panorama. This happens when the ability to receive template and device groups configuration changes has been disabled on the firewall.

To resolve the error, access the web interface of the device and select **Device > Setup**. Edit the Panorama Settings section and select the checkbox for  and .



## Why is Panorama running a 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 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.

To view the progress on the FSCK, set up console access to Panorama and view the status.

## Is there a separate connection for forwarding logs to Panorama?

No, Panorama uses TCP port 3978 for connecting to the firewalls.

For PAN-OS 4.x the SSL connection from the firewall to Panorama connects over TCP port 3978. This is a bi-directional connection where the logs are forwarded from the firewall to Panorama; and configuration changes are pushed from Panorama to the managed devices. Context switching commands are sent over the same connection.

For PAN-OS 5.0 and later, and only in a Distributed Log Collection architecture with dedicated Log Collectors, the firewalls manage two SSL connections. One connection is for Panorama management, and the other connection is to the Log Collector. Both connections use the same port: TCP port 3978.

## Why does the log storage capacity for the Collector Group read 0 MB?

The log storage capacity for the Collector Group might display as 0MB if the disk pairs are not enabled for logging. You must select the Log Collector and enable the disk pairs for logging in the **Panorama > Managed Collectors** tab; for instructions, see [Step 6](#) in the [Add a Log Collector to Panorama](#) section.

To verify that the disks are enabled and available for log storage, select **Panorama > Managed Collectors** tab and verify that the Log Collector displays as **Connected** and that the Configuration Status displays as **In sync**.

## Why is Panorama in a suspended state?

If Panorama is in a suspended state, check for the following conditions:

- 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.
- 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.
- Verify that both Panorama HA peers are running the same Panorama version (major and minor version number).

## Where do I view task completion status?

Use the Task Manager link in the bottom right-side corner of the Panorama web interface to view the success or failure of a task. It also includes a detailed message to help debug an issue. For details, see [View Task Completion History](#).

