Panorama Administrator's Guide

Version 9.1

paloaltonetworks.com/documentation

techDOCS paloalto NETWORKS

Contact Information

Corporate Headquarters:
Palo Alto Networks
3000 Tannery Way
Santa Clara, CA 95054
www.paloaltonetworks.com/company/contact-support

About the Documentation

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

Copyright

Palo Alto Networks, Inc.
www.paloaltonetworks.com

Last Revised

October 28, 2021

Table of Contents

Panorama Overview

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

About Panorama

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

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

Five Panorama Models are available: the Panorama virtual appliance, M-600 appliance, M-500 appliance, M-200 appliance, and M-100 appliance (M-100 appliances are supported in PAN-OS 9.1 only if they have been upgraded to 32 GB memory from the default 16 GB). Panorama Centralized Management illustrates how you can deploy Panorama in a high availability (HA) configuration to manage firewalls.

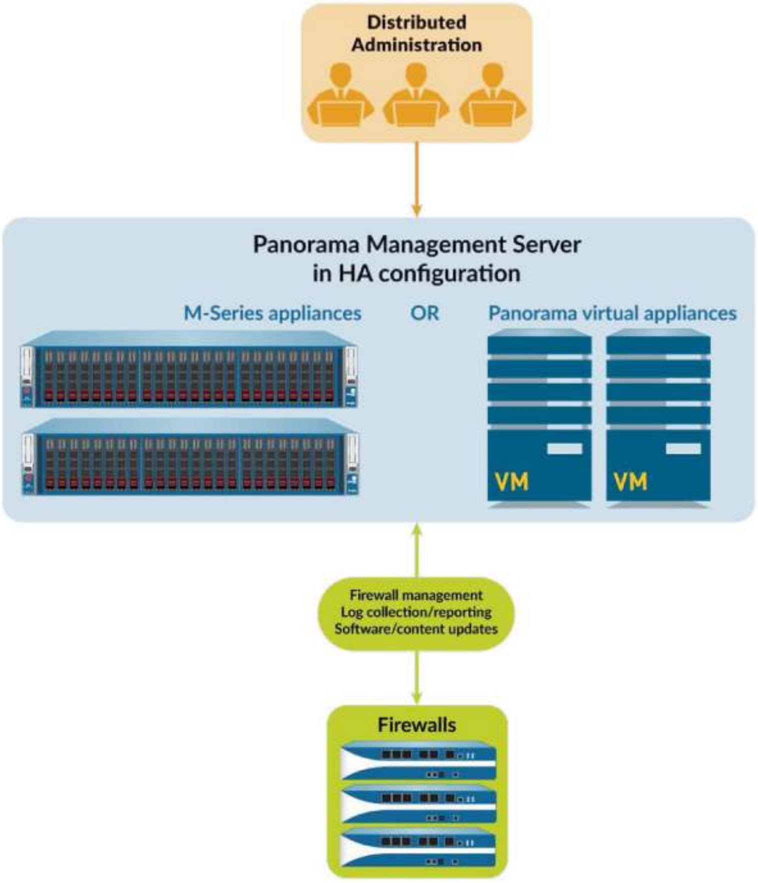

Figure 1: Panorama Centralized Management

Panorama Models

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

- **Panorama virtual appliance**—This model provides simple installation and facilitates server consolidation for sites that need a virtual management appliance. You can install Panorama on Amazon Web Services (AWS), AWS GovCloud, Microsoft Azure, Google Cloud Platform (GCP), KVM, Hyper-V, a VMware ESXi server, or on VMware vCloud Air. The virtual appliance can collect firewall logs locally at rates of up to 20,000 logs per second and can manage Dedicated Log Collectors for higher logging rates. The virtual appliance can function as a dedicated management server, a Panorama management server with local log collection capabilities, or as a Dedicated Log Collector. For the supported interfaces, log storage capacity, and maximum log collection rates, see the Setup Prerequisites for the Panorama Virtual Appliance. You can deploy the virtual appliance in the following modes:
  - **Panorama mode**—In this mode, the Panorama virtual appliance supports a local Log Collector with 1 to 12 virtual logging disks (see Deploy Panorama Virtual Appliances with Local Log Collectors). Each logging disk has 2TB of storage capacity for a total maximum of 24TB on a single virtual appliance and 48TB on a high availability (HA) pair. Only Panorama mode enables you to add multiple virtual logging disks without losing logs on existing disks. Panorama mode also provides the benefit of faster report generation. In Panorama mode, the virtual appliance does not support NFS storage.

    As a best practice, deploy the virtual appliance in Panorama mode to optimize log storage and report generation.
  - **Legacy mode** (ESXi and vCloud Air only)—In this mode, the Panorama virtual appliance receives and stores firewall logs without using a local Log Collector (see Deploy Panorama Virtual Appliances in Legacy Mode with Local Log Collection). By default, the virtual appliance in Legacy mode has one disk partition for all data. Approximately 11GB of the partition is allocated to log storage. If you need more local log storage, you can add one virtual disk of up to 8TB on ESXi 5.5 and later versions or on vCloud Air. Earlier ESXi versions support one virtual disk of up to 2TB. If you need more than 8TB, you can mount the virtual appliance in Legacy mode to an NFS datastore but only on the ESXi server, not in vCloud Air. This mode is only available if your Panorama virtual appliance is in Legacy mode on upgrade to PAN-OS 9.1. On upgrade to PAN-OS 9.0 and later releases, Legacy mode is no longer available if you change to any other mode. If you change your Panorama virtual appliance from Legacy mode to one of the available modes, you will no longer be able to change back into Legacy mode.

    While supported, Legacy mode is not recommended for production environments but may still be used for lab or demo environments.
  - **Management Only mode**—In this mode, the Panorama virtual appliance is a dedicated management appliance for your managed devices and Dedicated Log Collectors. Additionally, an appropriately resourced Panorama virtual appliance can manage up to 2,500 firewalls in this mode. The Panorama virtual appliance has no log collection capabilities except for config and system logs and requires a Dedicated Log Collector to these store logs. By default, the virtual appliance in Management Only mode has only one disk partition for all data so all logs forwarded to a Panorama virtual appliance in Management Only mode are dropped. Therefore, to store the log data from your managed appliances, you must configure log forwarding in order to store the log data from your managed devices. For more information, see Increased Device Management Capacity Requirements.
  - **Log Collector mode**—The Panorama virtual appliance functions as a Dedicated Log Collector. If multiple firewalls forward large volumes of log data, a Panorama virtual appliance in Log Collector mode provides increased scale and performance. In this mode, the appliance does not have a web

interface for administrative access; it has only a command line interface (CLI). However, you can manage the appliance using the web interface of the Panorama management server. CLI access to a Panorama virtual appliance in Log Collector mode is necessary only for initial setup and debugging. For configuration details, see Deploy Panorama with Dedicated Log Collectors.

- **M-Series appliance**—The M-100, M-200, M-500, and M-600 appliances are dedicated hardware appliances intended for large-scale deployments. In environments with high logging rates (over 10,000 logs per second) and log retention requirements, these appliances enable scaling of your log collection infrastructure. For the supported interfaces, log storage capacity, and maximum log collection rates, see M-Series Appliance Interfaces. All M-Series models share the following attributes:

  - RAID drives to store firewall logs and RAID 1 mirroring to protect against disk failures
  - SSD to store the logs that Panorama and Log Collectors generate
  - MGT, Eth1, Eth2, and Eth3 interfaces that support 1Gbps throughput
  - Redundant, hot-swappable power supplies (except for the M-100 appliance)
  - front-to-back airflow

  > *M-100 appliances are supported in PAN-OS 9.0 and later releases only if they have been upgraded to 32GB memory from the default 16GB. See* M-100 Memory Upgrade Guide *for more information.*

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

  - Eth4 and Eth5 interfaces that support 10Gbps throughput

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

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

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

  - **Panorama mode**—The appliance functions as a Panorama management server to manage firewalls and Dedicated Log Collectors. The appliance also supports a local Log Collector to aggregate firewall logs. Panorama mode is the default mode. For configuration details, see Deploy Panorama M-Series Appliances with Local Log Collectors.
  - **Management Only mode**—The Panorama appliance is a dedicated management appliance for your managed devices and Dedicated Log Collectors. The Panorama appliance has no log collection capabilities except for config and system logs and your deployment requires a Dedicated Log Collector to store these logs. By default, the Panorama appliance in Management Only mode has only one disk partition for all data so all logs forwarded to a Panorama virtual appliance in Management Only mode are dropped. Therefore, to store the log data from your managed appliances, you must configure log forwarding in order to store the log data from your managed devices.
  - **Log Collector mode**—The appliance functions as a Dedicated Log Collector. If multiple firewalls forward large volumes of log data, an M-Series appliance in Log Collector mode provides increased scale and performance. IIn this mode, the appliance does not have a web interface for administrative access; it has only a command line interface (CLI). However, you can manage the appliance using the web interface of the Panorama management server. CLI access to an M-Series appliance in Log Collector mode is necessary only for initial setup and debugging. For configuration details, see Deploy Panorama with Dedicated Log Collectors.

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

Centralized Firewall Configuration and Update Management

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

- Context Switch—Firewall or Panorama
- Templates and Template Stacks
- Device Groups

Context Switch—Firewall or Panorama

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

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

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

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

Templates and Template Stacks

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

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

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

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

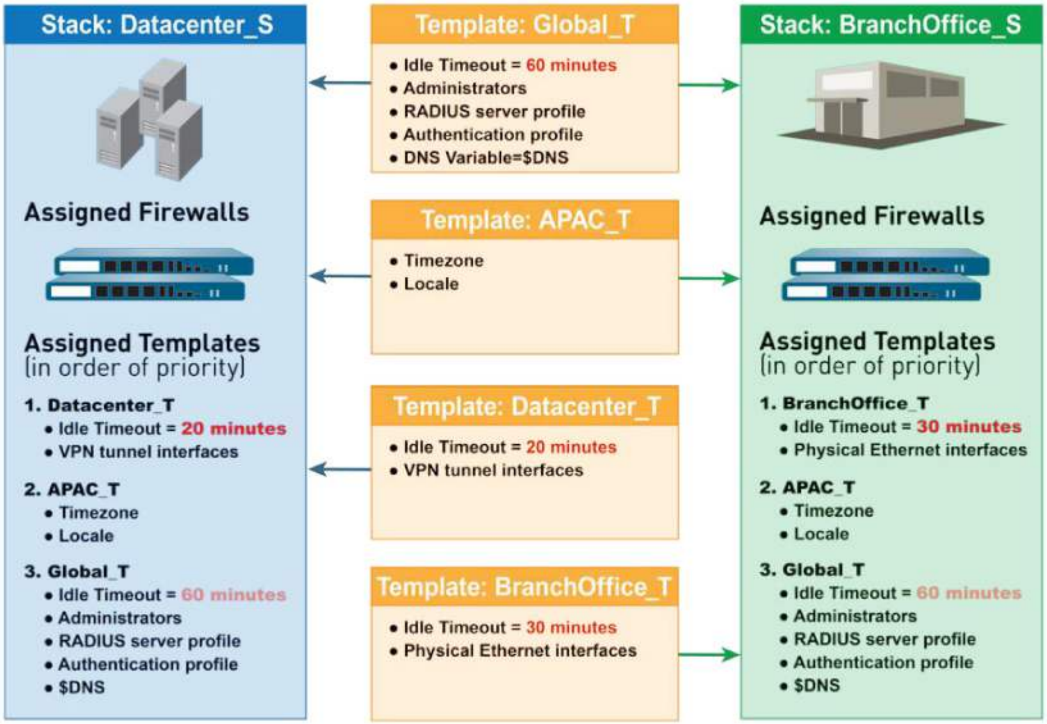

Figure 2: Template Stacks

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

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

Device Groups

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

- Device Group Hierarchy
- Device Group Policies
- Device Group Objects

Device Group Hierarchy

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

Creating a device group hierarchy enables you to organize firewalls based on common policy requirements without redundant configuration. For example, you could configure shared settings that are global to all firewalls, configure device groups with function-specific settings at the first level, and configure device groups with location-specific settings at lower levels. Without a hierarchy, you would have to configure both function- and location-specific settings for every device group in a single level under Shared.

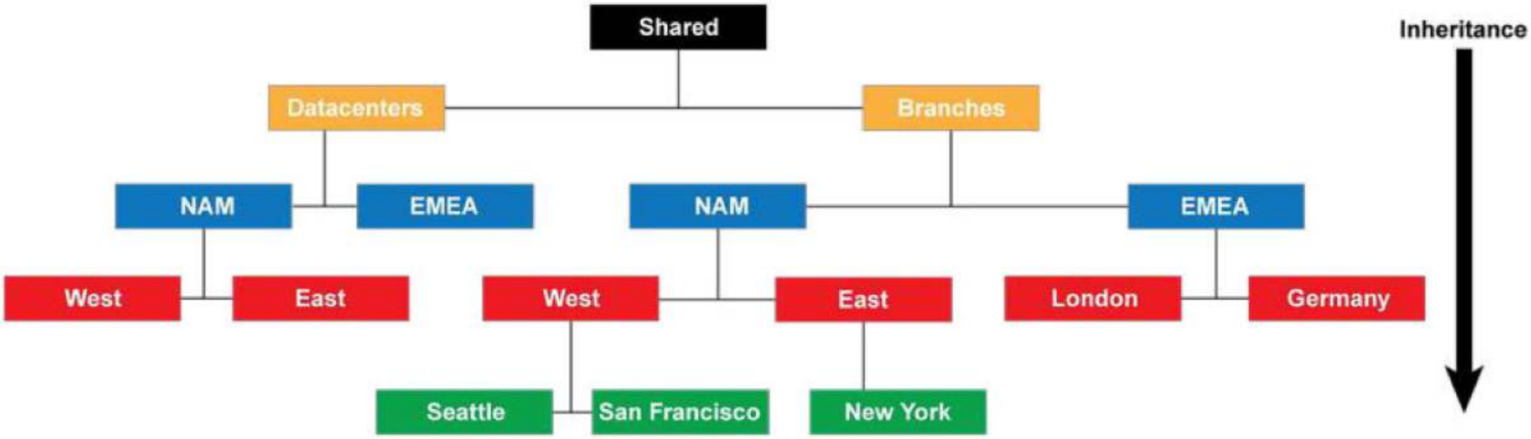

Figure 3: Device Group Hierarchy

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

Device Group Policies

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

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

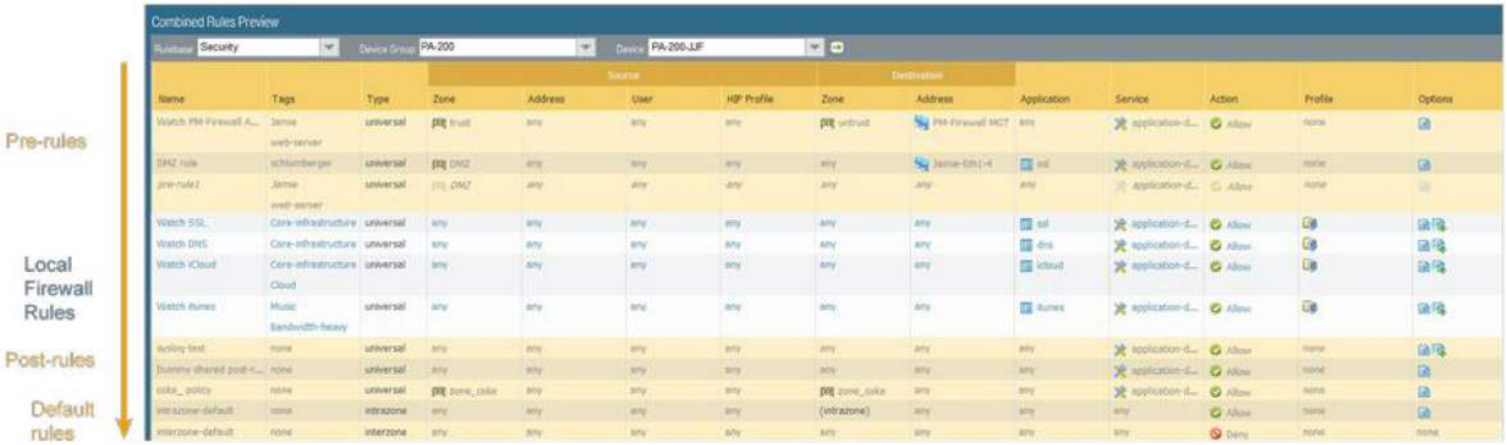

Combined Rules Preview

Rulebase: Security | Device Group: PA-200 | Device: PA-200-JJF

			Source				Destination						
Name	Tags	Type	Zone	Address	User	HIP Profile	Zone	Address	Application	Service	Action	Profile	Options
Watch PM-Firewall A...	Jamie web-server	universal	trust	any	any	any	untrust	PM-Firewall MGT	any	application-d...	Allow	none	
DMZ rule	schlumberger	universal	DMZ	any	any	any	any	Jamie-Eth1-4	ssl	application-d...	Allow	none	
pre-rule1	Jamie web-server	universal	DMZ	any	any	any	any	any	any	application-d...	Allow	none	
Watch SSL	Core-infrastructure	universal	any	any	any	any	any	any	ssl	application-d...	Allow		
Watch DNS	Core-infrastructure	universal	any	any	any	any	any	any	dns	application-d...	Allow		
Watch iCloud	Core-infrastructure Cloud	universal	any	any	any	any	any	any	icloud	application-d...	Allow		
Watch itunes	Music Bandwidth-heavy	universal	any	any	any	any	any	any	itunes	application-d...	Allow		
syslog-test	none	universal	any	any	any	any	any	any	any	application-d...	Allow	none	
Dummy shared post-r...	none	universal	any	any	any	any	any	any	any	application-d...	Allow	none	
coke_policy	none	universal	zone_coke	any	any	any	zone_coke	any	any	application-d...	Allow	none	
intrazone-default	none	intrazone	any	any	any	any	(intrazone)	any	any	any	Allow	none	
interzone-default	none	interzone	any	any	any	any	any	any	any	any	Deny	none	none

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

Evaluation Order	Rule Scope and Description	Administration Device
		to a local firewall context, can edit local firewall rules.
Device group post-rules	Panorama pushes shared post-rules to all the firewalls in all device groups. Panorama pushes device group-specific post-rules to all the firewalls in a particular device group and its descendant device groups. If a firewall inherits rules from device groups at multiple levels in the device group hierarchy, it evaluates post-rules in the order of lowest to highest level. This means the firewall first evaluates the rules of device groups with no descendants and last evaluates shared rules. Post-rules typically include rules to deny access to traffic based on the App-ID™ signatures, User-ID™ information (users or user groups), or service.	These rules are visible on firewalls but you can only manage them in Panorama.
Shared post-rules		
intrazone-default interzone-default	The default rules apply only to the Security rulebase, and are predefined on Panorama (at the Shared level) and the firewall (in each vsys). These rules specify how PAN-OS handles traffic that doesn't match any other rule. The intrazone-default rule allows all traffic within a zone. The interzone-default rule denies all traffic between zones. If you override default rules, their order of precedence runs from the lowest context to the highest: overridden settings at the firewall level take precedence over settings at the device group level, which take precedence over settings at the Shared level.	Default rules are initially read-only, either because they are part of the predefined configuration or because Panorama pushed them to firewalls. However, you can override the rule settings for tags, action, logging, and security profiles. The context determines the level at which you can override the rules: • Panorama—At the Shared or device group level, you can override default rules that are part of the predefined configuration. • Firewall—You can override default rules that are part of the predefined configuration on the firewall or vsys, or that Panorama pushed from the Shared location or a device group.

Device Group Objects

Objects are configuration elements that policy rules reference, for example: IP addresses, URL categories, security profiles, users, services, and applications. Rules of any type (pre-rules, post-rules, default rules, and rules locally defined on a firewall) and any rulebase (Security, NAT, QoS, Policy Based Forwarding, Decryption, Application Override, Captive Portal, and DoS Protection) can reference objects. You can reuse an object in any number of rules that have the same scope as that object in the Device Group Hierarchy.

For example, if you add an object to the Shared location, all rules in the hierarchy can reference that *shared object* because all device groups inherit objects from Shared. If you add an object to a particular device group, only the rules in that device group and its descendant device groups can reference that *device group object*. If object values in a device group must differ from those inherited from an ancestor device group, you can Override inherited object values (see Step Override inherited object values.). You can also Revert to Inherited Object Values at any time. When you Create Objects for Use in Shared or Device Group Policy once and use them many times, you reduce administrative overhead and ensure consistency across firewall policies.

You can configure how Panorama handles objects system-wide:

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

Centralized Logging and Reporting

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

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

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

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

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

Managed Collectors and Collector Groups

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

- **Local Log Collector**—This type of Log Collector runs locally on the Panorama management server. Only an M-600, M-500 appliance, M-200, M-100 appliance, or Panorama virtual appliance in Panorama mode supports a local Log Collector.

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

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

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

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

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

Local and Distributed Log Collection

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

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

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

Dedicated Log Collectors are M-600, M-500, M-200, or M-100 appliances in Log Collector mode. Because they perform only log collection, not firewall management, Dedicated Log Collectors allow for a more robust environment than local Log Collectors. Dedicated Log Collectors provide the following benefits:

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

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

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

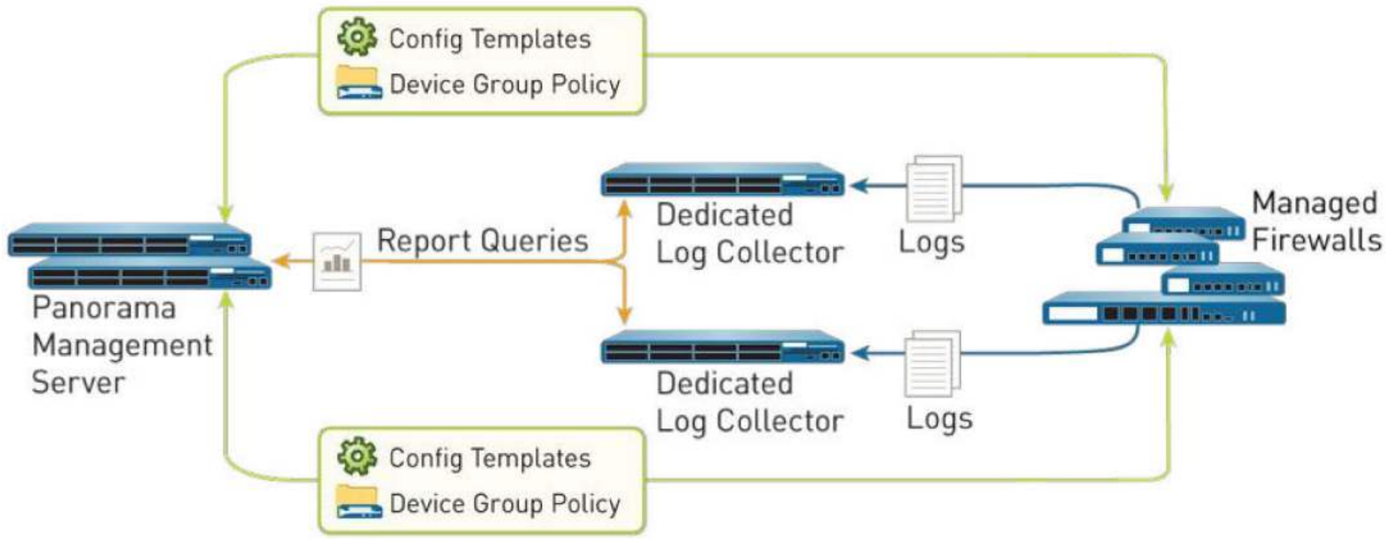

Figure 4: Distributed Log Collection

Caveats for a Collector Group with Multiple Log Collectors

You can Configure a Collector Group with multiple Log Collectors (up to 16) to ensure log redundancy, increase the log retention period, and accommodate logging rates that exceed the capacity of a single Log Collector (see Panorama Models for capacity information). In any single Collector Group, all the Log Collectors must run on the same Panorama model: all M-600 appliances, all M-500 appliances, all, M-200 appliances all M-100 appliances, or all Panorama virtual appliances. For example, if a single managed firewall generates 48TB of logs, the Collector Group that receives those logs will require at least six Log Collectors that are M-100 appliances or two Log Collectors that are M-500 appliances or Panorama virtual appliances.

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

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

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

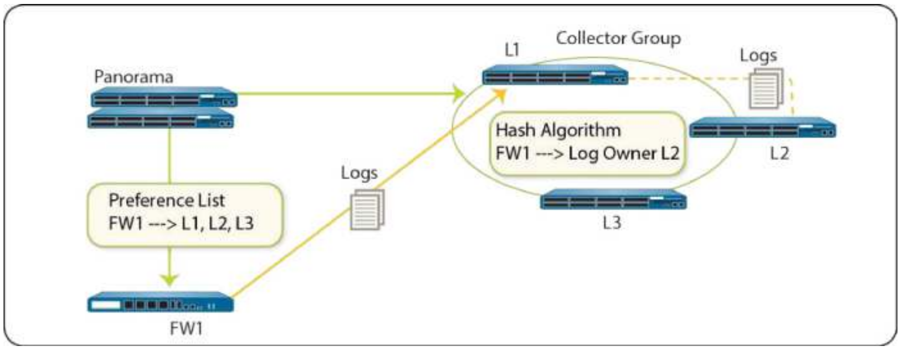

Figure 5: Example - Typical Log Collector Group Setup

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

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

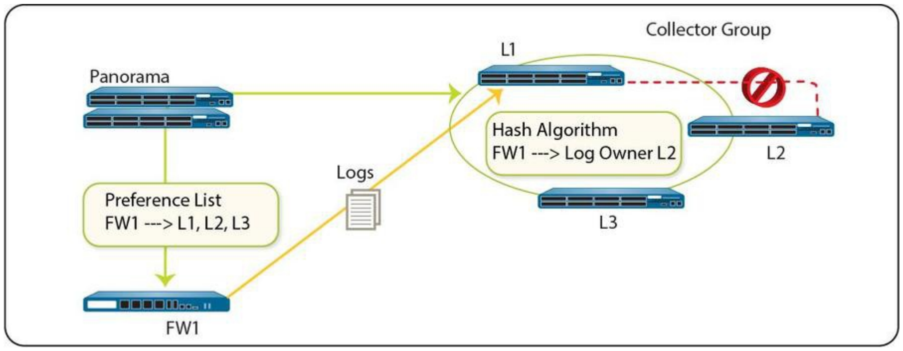

Figure 6: Example - When a Log Collector Fails

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

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

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

  Enabling redundancy doubles the log processing traffic in a Collector Group, which reduces its maximum logging rate by half, as each Log Collector must distribute a copy of each log it receives.
- Obtain an On-Site-Spare (OSS) to enable prompt replacement if a Log Collector failure occurs.
- In addition to forwarding logs to Panorama, configure forwarding to an external service as backup storage. The external service can be a syslog server, email server, SNMP trap server, or HTTP server.

Log Forwarding Options

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

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

- Forward logs from firewalls to Panorama and from Panorama to external services—This configuration is best for deployments in which the connections between firewalls and external services have insufficient bandwidth to sustain the logging rate, which is often the case when the connections are remote. This configuration improves firewall performance by offloading some processing to Panorama.

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

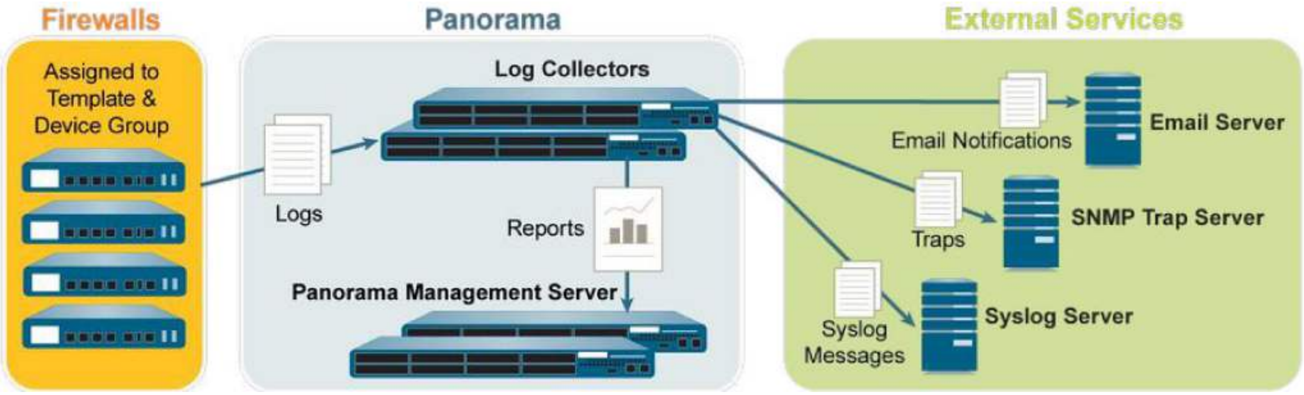

Figure 7: Log Forwarding to Panorama and then to External Services

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

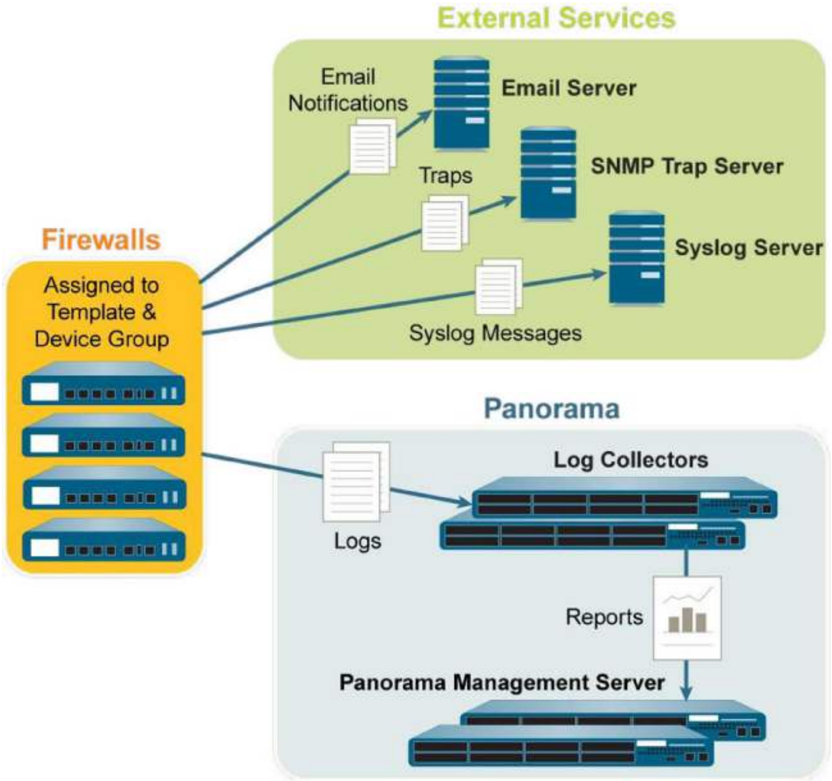

Figure 8: Log Forwarding to External Services and Panorama in Parallel

Centralized Reporting

Panorama aggregates logs from all managed firewalls and enables reporting on the aggregated data for a global view of application use, user activity, and traffic patterns across the entire network. As soon as the firewalls are added to Panorama, the ACC can display all traffic traversing your network. With logging enabled, clicking into a log entry in the ACC provides direct access to granular details about the application.

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

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

For more information, see Monitor Network Activity.

User-ID Redistribution Using Panorama

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

Each firewall, Log Collector, or Panorama management server can receive User-ID information from up to 100 redistribution points. The redistribution points can be Windows-based User-ID agents or other firewalls, Log Collectors, and Panorama management servers. Panorama and Log Collectors as User-ID Redistribution Points illustrates a redistribution sequence where the firewalls perform user mapping by directly monitoring information sources such as directory servers and syslog senders. However, you can also use Windows-based User-ID agents to perform the mapping and redistribute the information to firewalls. Only the firewalls record authentication timestamps when user traffic matches Authentication policy rules.

You can redistribute user mappings collected through any method except Terminal Services (TS) agents. You cannot redistribute username-to-group mapping or HIP match information.

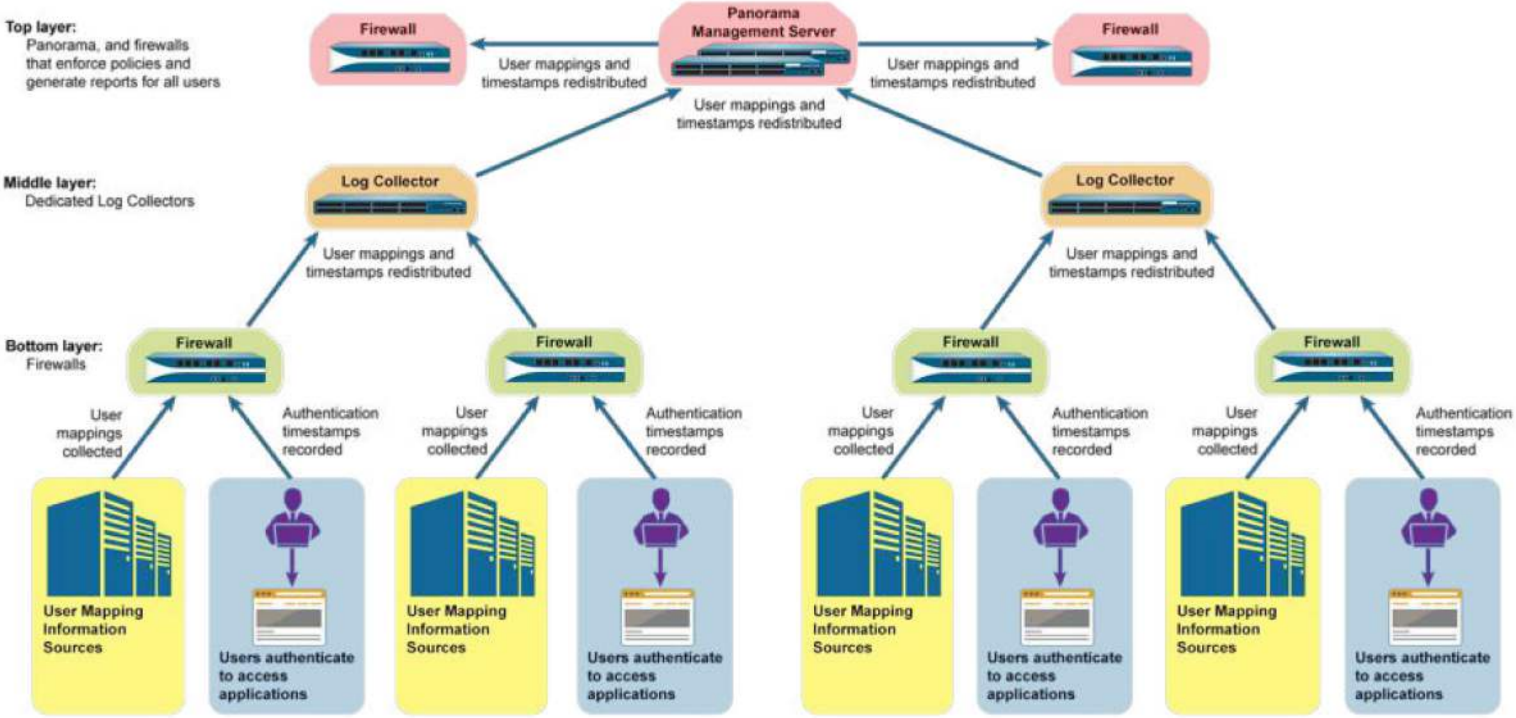
Top layer:
Panorama, and firewalls that enforce policies and generate reports for all users
Firewall
Panorama Management Server
Firewall
User mappings and timestamps redistributed
User mappings and timestamps redistributed
User mappings and timestamps redistributed
Middle layer:
Dedicated Log Collectors
Log Collector
Log Collector
User mappings and timestamps redistributed
User mappings and timestamps redistributed
Bottom layer:
Firewalls
Firewall
Firewall
Firewall
Firewall
User mappings collected
Authentication timestamps recorded
User Mapping Information Sources
Users authenticate to access applications

Figure 9: Panorama and Log Collectors as User-ID Redistribution Points

Role-Based Access Control

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

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

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

Administrative Roles

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

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

Dynamic Role	Privileges
Superuser	Full read-write access to Panorama
Superuser (read-only)	Read-only access to Panorama
Panorama administrator	Full access to Panorama except for the following actions: • Create, modify, or delete Panorama or firewall administrators and roles. • Export, validate, revert, save, load, or import a configuration in the **Device** > **Setup** > **Operations** page.

Dynamic Role	Privileges
	• Configure **Scheduled Config Export** functionality in the **Panorama** tab.

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

Admin Role Profile	Description
Panorama	For these roles, you can assign read-write access, read-only access, or no access to all the Panorama features that are available to the superuser dynamic role except the management of Panorama administrators and Panorama roles. For the latter two features, you can assign read-only access or no access, but you cannot assign read-write access. An example use of a Panorama role would be for security administrators who require access to security policy definitions, logs, and reports on Panorama.
Device Group and Template	For these roles, you can assign read-write access, read-only access, or no access to specific functional areas within device groups, templates, and firewall contexts. By combining these roles with Access Domains, you can enforce the separation of information among the functional or regional areas of your organization. Device Group and Template roles have the following limitations: • No access to the CLI or XML API • No access to configuration or system logs • No access to VM information sources • In the **Panorama** tab, access is limited to: • Device deployment features (read-write, read-only, or no access) • The device groups specified in the administrator account (read-write, read-only, or no access) • The templates and managed firewalls specified in the administrator account (read-only or no access) An example use of this role would be for administrators in your operations staff who require access to the device and network configuration areas of the web interface for specific device groups and/or templates.

Authentication Profiles and Sequences

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

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

Access Domains

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

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

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

For the relevant procedures, see:

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

Administrative Authentication

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

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

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

Panorama Commit, Validation, and Preview Operations

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

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

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

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

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

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

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

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

Plan Your Panorama Deployment

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

  If you already deployed and configured the Palo Alto Networks firewalls on your network, determine whether to transition the firewalls to centralized management. This process requires a migration of all configuration and policies from your firewalls to Panorama. For details, see Transition a Firewall to Panorama Management.
- □ Verify the Panorama and firewall software versions. Panorama can manage firewalls running PAN-OS versions that match the Panorama version or are earlier than the Panorama version. For example, Panorama 8.0 cannot manage firewalls running PAN-OS 8.1. Additionally, Panorama 8.1 cannot manage firewalls running PAN-OS 6.0.0 through 6.0.3 and cannot manage firewalls that run a later PAN-OS version than the Panorama version.
- □ 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 rules 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.
- □ Determine your authentication method between Panorama and its managed devices and high availability peer. By default, Panorama uses predefined certificates to authenticate the SSL connections used for management and inter-device communication. However, you can configure custom certificate-based authentication to enhance the security of the SSL connections between Panorama, firewalls, and log collectors. By using custom certificates, you can establish a unique chain of trust to ensure mutual authentication between Panorama and the devices it manages. You can import the certificates from your enterprise public key infrastructure (PKI) or generate it on Panorama.
- □ Plan to use Panorama in a high availability configuration; set it up as an active/passive high availability pair. See Panorama High Availability.
- □ Plan how to accommodate network segmentation and security requirements in a large-scale deployment. By default, Panorama running on an M-Series appliance uses the management (MGT) interface for administrative access to Panorama and for managing devices (firewalls, Log Collectors, and WildFire appliances and appliance clusters), collecting logs, communicating with Collector Groups, and deploying software and content updates to devices. However, to improve security and enable network segmentation, you can reserve the MGT interface for administrative access and use dedicated M-Series Appliance Interfaces (Eth1, Eth2, Eth3, Eth4, and Eth5) for the other services.
- □ For meaningful reports on network activity, plan a logging solution:

  - Verify the resource allocation for your Panorama virtual appliance deployed in Log Collector mode on AWS or Azure. The Panorama virtual appliance does not retain Log Collector mode if resized. This results in log data loss.
  - Estimate the log storage capacity your network needs to meet security and compliance requirements. Consider such factors as the logging capacities of your Panorama Models, network topology, number of firewalls sending logs, type of log traffic (for example, URL Filtering and Threat logs versus Traffic logs), the rate at which firewalls generate logs, and the number of days for which you want to store logs on Panorama. For details, see Determine Panorama Log Storage Requirements.
  - Do you need to forward logs to external services (such as a syslog server) in addition to Panorama? See Log Forwarding Options.
  - Do you want to own or manage your own log storage on premises, or do you want to leverage the Cortex Data Lake provided by Palo Alto Networks?
  - If you need a long-term storage solution, do you have a Security Information and Event Management (SIEM) solution, such as Splunk or ArcSight, to which you can forward logs?
  - Do you need redundancy in logging?

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

If you deploy Panorama virtual appliances in Legacy mode in an HA configuration, the managed firewalls can send logs to both HA peers so that a copy of each log resides on each peer. This redundancy option is enabled by default (see Modify Log Forwarding and Buffering Defaults).

- Will you log to a Network File System (NFS)? If the Panorama virtual appliance is in Legacy mode and does not manage Dedicated Log Collectors, NFS storage is the only option for increasing log storage capacity beyond 8TB. NFS storage is available only if Panorama runs on an ESXi server. If you use NFS storage, keep in mind that the firewalls can send logs only to the primary peer in the HA pair; only the primary peer is mounted to the NFS and can write to it.

□ Determine which role-based access privileges administrators require to access managed firewalls and Panorama. See Set Up Administrative Access to Panorama.

□ Plan the required Device Groups. Consider whether to group firewalls based on function, security policy, geographic location, or network segmentation. An example of a function-based device group is one that contains all the firewalls that a Research and Development team uses. Consider whether to create smaller device groups based on commonality, larger device groups to scale more easily, or a Device Group Hierarchy to simplify complex layers of administration.

□ Plan a layering strategy for administering policies. Consider how firewalls inherit and evaluate policy rules within the Device Group Hierarchy, and how to best implement shared rules, device-group rules, and firewall-specific rules to meet your network needs. For visibility and centralized policy management, consider using Panorama for administering rules even if you need firewall-specific exceptions for shared or device group rules. If necessary, you can Push a Policy Rule to a Subset of Firewalls within a device group.

□ Plan the organization of your firewalls based on how they inherit network configuration settings from Templates and Template Stacks. For example, consider assigning firewalls to templates based on hardware models, geographic proximity, and similar network needs for time zones, a DNS server, and interface settings.

Deploy Panorama: Task Overview

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

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

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

STEP 3 | Register Panorama and Install Licenses.

STEP 4 | Install Content and Software Updates for Panorama.

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

STEP 6 | Add a Firewall as a Managed Device.

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

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

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

Set Up Panorama

For centralized reporting and cohesive policy management across all the firewalls on your network, you can deploy the Panorama™ management server as a virtual appliance or as a hardware appliance (the M-100, M-200, M-500 or M-600 appliance).

> *M-100 appliances are supported in PAN-OS 9.1 only if they have been upgraded to 32GB memory from the default 16GB. See M-100 Memory Upgrade Guide for more information.*

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

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

Determine Panorama Log Storage Requirements

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

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

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

Factors that affect log retention requirements include:

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

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

STEP 2 | Determine the average daily logging rates.

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

1. Display the current log generation rate in logs per second:

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

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

   - If Panorama is already collecting logs, run the following command at the CLI of each appliance that receives logs (Panorama management server or Dedicated Log Collector) and calculate the total rates. This command gives the average logging rate for the last five minutes.

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

   You can also use an SNMP manager to determine the logging rates of Log Collectors (see the panLogCollector MIB, OID 1.3.6.1.4.1.25461.1.1.6) and firewalls (see the panDeviceLogging, OID 1.3.6.1.4.1.25461.2.1.2.7).
2. Calculate the average of the sampled rates.
3. Calculate the daily logging rate by multiplying the average logs-per-second by 86,400.

STEP 3 | Estimate the required storage capacity.

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

Use the formula:

<required_storage_duration> x <average_log_size> x <average_logging_rate>

The average log size varies considerably by log type. However, you can use 500 bytes as an approximate average log size.

For example, if Panorama must store logs for 30 days and the average total logging rate for all firewalls is 21,254,400 logs per day, then the required log storage capacity is: 30 x 500 x 21,254,400 = 318,816,000,000 bytes (approximately 318GB).

STEP 4 | Next steps...

If you determine that Panorama requires more log storage capacity:

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

Manage Large-Scale Firewall Deployments

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

- Determine the Optimal Large-Scale Firewall Deployment Solution
- Increased Device Management Capacity for M-600 and Panorama Virtual Appliance

Determine the Optimal Large-Scale Firewall Deployment Solution

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

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

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

Increased Device Management Capacity for M-600 and Panorama Virtual Appliance

The M-600 appliance in Management Only mode can manage up to 5,000 firewalls or a Panorama virtual appliance in Management Only mode can manage up to 2,500 firewalls in order to reduce the management footprint of your large-scale firewall deployment.

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

Increased Device Management Capacity Requirements

You can manage up to 5,000 firewalls using a single M-600 appliance in Management Only mode or manage up to 2,500 firewalls using a single Panorama virtual appliance in Management Only mode. Managing such large deployments from a single Panorama management server alleviates the operational complexity of configuration management and reduces the security and compliance risk of managing multiple Panorama management servers.

For log collection, a single Panorama management server is ideal because it provides a centralized location to view and analyze log data from managed devices rather than requiring you to access each individual Panorama management server. To provide redundancy in the event of system or network failure, Palo

Alto Networks recommends deploying two Panorama management servers in a high availability (HA) configuration.

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

To manage up to 5,000 firewalls, the Panorama management server must meet the following minimum requirements:

Requirement	M-Series Appliance	Panorama Virtual Appliance
Model	M-600	All supported Panorama hypervisors. For more information, see Panorama Models.
Panorama Mode	Management Only	Management Only
Number of managed firewalls	5,000	2,500
System Disk	240GB SSD—Used to store the operating system files and system logs.	• 81GB—Used to store the operating system files and system logs. • Additional disk with a minimum 90GB capacity.
Cores	28 (with Hyper-Threaded)	28 (with Hyper-Threaded)
Memory	256GB	250GB
Log Collection	Local log collection is not supported. See Deploy Panorama with Dedicated Log Collectors to set up log collection.	Local log collection is not supported. See Deploy Panorama with Dedicated Log Collectors to set up log collection.
Logging and Reporting	Enable the **Use Panorama Data for Pre-Defined Reports** setting (**Panorama** > **Setup** > **Management** > **Logging and Reporting Settings** > **Log Export and Reporting**)	Enable the **Use Panorama Data for Pre-Defined Reports** setting (**Panorama** > **Setup** > **Management** > **Logging and Reporting Settings** > **Log Export and Reporting**)

Deploy Panorama for Increased Device Management

To deploy Panorama for increased device management, determine your deployment scenario and follow the procedure:

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

Install Panorama for Increased Device Management Capacity

Activate the device management license to manage more than 1,000 firewalls from a single M-600 Panorama™ management server or a single Panorama virtual appliance.

STEP 1 | Contact your Palo Alto Networks sales representative to obtain the Panorama device management license that enables you to manage up to 5,000 firewalls with an M-600 appliance or up to 2,500 firewalls with a Panorama virtual appliance.

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

STEP 2 | Set up the Panorama management server.

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

or

- Set Up the Panorama Virtual Appliance.

STEP 3 | Change the Panorama management server to Management Only mode if Panorama is not already in this mode.

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

STEP 4 | Register your Panorama management server and install licenses.

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

STEP 5 | Select **Panorama** > **Licenses** and verify that the device management license is successfully activated.

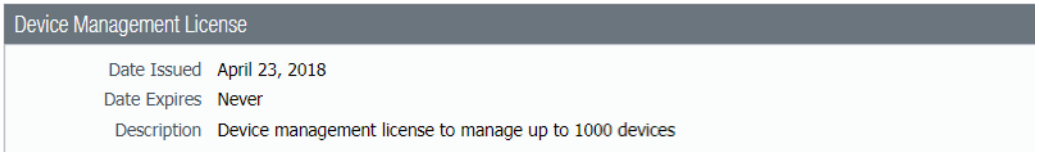

Device Management License

Date Issued	April 23, 2018
Date Expires	Never
Description	Device management license to manage up to 1000 devices

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

Upgrade Panorama for Increased Device Management Capacity

Upgrade to PAN-OS 9.1 to use your existing device management license on your M-600 appliance to manage up to 5,000 firewalls or Panorama™ virtual appliance to manage up to 2,500 firewalls.

STEP 1 | Increase CPUs and Memory on the Panorama Virtual Appliance if the Panorama virtual appliance does not already meet the minimum resource requirements for increased device management.

Review the Increased Device Management Capacity Requirements to verify whether your existing Panorama virtual appliance meets the minimum requirements before upgrading.

STEP 2 | Log in to the Panorama CLI.

STEP 3 | Change the Panorama management server to Management Only if Panorama is not already in this mode.

- (M-600 appliances only) Begin at Step 5 to Set Up an M-Series Appliance in Management Only Mode.

or

- Set up a Panorama Virtual Appliance in Management Only Mode.

STEP 4 | Log in to the Panorama Web Interface.

STEP 5 | Upgrade the Panorama management server.

- Install Updates for Panorama with an Internet Connection.
- Install Updates for Panorama When Not Internet-Connected.
- Install Updates for Panorama in an HA Configuration.

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

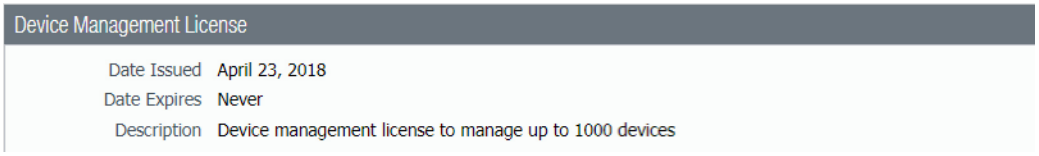

Device Management License

Date Issued	April 23, 2018
Date Expires	Never
Description	Device management license to manage up to 1000 devices

If you activated your device management license and then upgraded to PAN-OS 9.1, you can manage up to 5,000 firewalls with an M-600 appliance, or up to 2,500 firewalls with a Panorama virtual appliance, but the Description still displays `Device management license to manage up to 1000 devices`.

Set Up the Panorama Virtual Appliance

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

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

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

Setup Prerequisites for the Panorama Virtual Appliance

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

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

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

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

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

It is recommended to deploy the Panorama management server in Panorama mode for both device management and log collection capabilities. While still supported, Legacy mode is not recommended for production environments. Additionally, you can no longer switch Panorama to Legacy mode. For more information on supported modes, see Panorama Models.

Table 1: System Requirements for the Panorama Virtual Appliance

Requirements	Panorama Virtual Appliance in Management Only Mode	Panorama Virtual Appliance in Panorama Mode	Panorama Virtual Appliance in Log Collector Mode
Virtual hardware version	- **VMware ESXi and vCloud Air**—64-bit kernel-based VMware ESXi 6.0, 6.5, 6.7, or 7.0. The supported version of the virtual hardware family type (also known as the VMware virtual hardware version) on the ESXi server is vmx-10 *The Panorama virtual appliance for ESXi does not support the creation quiesced snapshots. Disable Quiesce guest file system in the vSphere client or set the* **`quiesce`** *flag to 0 or false in the vSphere CLI before creating a snapshot of your virtual Panorama appliance.* - **Hyper-V**—Windows Server 2016 with Hyper-V role or Hyper-V 2016 - **KVM**—Ubuntu version 16.04 or CentOS7 In Panorama mode, the virtual appliance running on any ESXi version supports up to 12 virtual logging disks with 2TB of log storage each, for a total maximum capacity of 24TB. (VMware ESXi and vCloud Air only) In Legacy mode, the virtual appliance supports one virtual logging disk. ESXi 5.5 and later versions supports one disk of up to 8TB. Earlier ESXi versions support one disk of up to 2TB.		
(ESXi and vCloud Air only) Client computer	To install the Panorama virtual appliance and manage its resources, you must install a VMware vSphere Client or VMware Infrastructure Client that is compatible with your ESXi server.		
System disk	- **Default**—81GB - (ESXi and GCP only) **Upgraded**—224GB An upgraded system disk is required for SD-WAN.	- **Default**—81GB - (ESXi and GCP only) **Upgraded**—224GB An upgraded system disk is required for SD-WAN. For log storage, Panorama uses virtual logging disks	81GB For log storage, Panorama uses virtual logging disks instead of the system disk or an NFS datastore.

Requirements	Panorama Virtual Appliance in Management Only Mode	Panorama Virtual Appliance in Panorama Mode	Panorama Virtual Appliance in Log Collector Mode
		instead of the system disk or an NFS datastore.	
CPUs, memory, and logging disks	• Manage up to 500 managed devices • 16 CPUs • 32GB memory • Local log storage not supported • Manage up to 1,000 managed devices • 32 CPUs • 128GB memory • Local log storage not supported • To manage more than 1,000 firewalls, see Increased Device Management Capacity Requirements.	• Up to 10,000 logs/sec: • 16 CPUs • 32GB memory • 4x2TB logging disks • Manage up to 500 managed devices • Up to 20,000 log/sec • 32 CPUs • 128GB memory • 8x2TB logging disks • Manage up to 1,000 managed devices	• Up to 15,000 log/sec • 16 CPUs • 32GB memory • 4x2TB logging disks • Up to 25,000 logs/sec • 32 CPUs • 128GB memory • 8x2TB logging disks
Minimum CPUs and memory	• 16 CPUs • 32GB memory	The minimum resources below do not take LPS into consideration and are only required for the Panorama virtual appliance to function based on the number of logging disks added. Palo Alto Networks recommends you refer to the recommended resources above. For larger Panorama deployments, be aware that you may be under-provisioning your Panorama. This may lead to impacted performance and may cause Panorama to become unresponsive depending on the number of firewalls managed, the configuration size, the number of administrators logged in to Panorama, and the volume of logs ingested. • **2TB to 8TB**—16 CPUs, 32GB memory • **10TB to 24TB**— 16 CPUs, 64GB memory	
Log storage capacity	Panorama in Management Only mode requires log forwarding to a Dedicated Log Collector.	2TB to 24TB	2TB to 24TB

Supported Interfaces

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

Function	Amazon Web Services (AWS) and AWS GovCloud		Microsoft Azure	Google Cloud Platform (GCP)	KVM	Hyper-V	VMware (ESXi, vCloud Air)
Device Management	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported
Device Log Collection	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported
Collector Group Communication	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported	Any interface supported
Licensing and Software Updates	MGT interface only	MGT interface only	MGT interface only	MGT interface only	Any interface supported	Any interface supported	Any interface supported

Install the Panorama Virtual Appliance

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

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

- Install Panorama on VMware
- Install Panorama on AWS
- Install Panorama on AWS GovCloud
- Install Panorama on Azure
- Install Panorama on Google Cloud Platform
- Install Panorama on KVM
- Install Panorama on Hyper-V

Install Panorama on VMware

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

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

Install Panorama on an ESXi Server

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

STEP 1 | Download the Panorama 9.1 base image Open Virtual Appliance (OVA) file.

1. Go to the Palo Alto Networks software downloads site. (If you can't log in, go to the Palo Alto Networks Customer Support web site for assistance.)
2. In the Download column in the Panorama Base Images section, download the latest version of the Panorama release OVA file (`Panorama-ESX-9.1.0.ova`).

STEP 2 | Install Panorama.

1. Launch the VMware vSphere Client and connect to the VMware server.
2. Select **File** > **Deploy OVF Template**.
3. **Browse** to select the Panorama OVA file and click **Next**.
4. Confirm that the product name and description match the downloaded version, and click **Next**.
5. Enter a descriptive name for the Panorama virtual appliance, and click **Next**.
6. Select a datastore location (system disk) on which to install the Panorama image. See the Setup Prerequisites for the Panorama Virtual Appliance for the supported system disk sizes. After selecting the datastore, click **Next**.
7. Select **Thick Provision Lazy Zeroed** as the disk format, and click **Next**.
8. Specify which networks in the inventory to use for the Panorama virtual appliance, and click **Next**.
9. Confirm the selected options, click **Finish** to start the installation process, and click **Close** when it finishes. Do not power on the Panorama virtual appliance yet.

STEP 3 | Configure resources on the Panorama virtual appliance.

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

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

   This step is required in the following scenarios:

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

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

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

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

      When adding multiple logging disks, it is a best practice to select the same Disk Provisioning format for all disks to avoid any unexpected performance issues that may arise.
   5. Select **Specify a datastore or datastore structure** as the location, **Browse** to a datastore that has sufficient storage, click **OK**, and click **Next**.

    6. Select a SCSI **Virtual Device Node** (you can use the default selection) and click **Next**.

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

    7. Verify that the settings are correct and click **Finish**.
5. Click **OK** to save your changes.

STEP 4 | Power on the Panorama virtual appliance.

1. In the vSphere Client, right-click the Panorama virtual appliance and select **Power** > **Power On**. Wait for Panorama to boot up before continuing.
2. Verify that the virtual appliance is running in the correct mode:
    1. Right-click the Panorama virtual appliance and select **Open Console**.
    2. Enter your username and password to log in (default is **`admin`** for both).
    3. Display the mode by running the following command:

```
> show system info
```

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

STEP 5 | Register the Panorama virtual appliance and activate the device management license and support licenses.

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

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

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

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

STEP 6 | Increase the System Disk for Panorama on an ESXi Server if you intend to use the Panorama virtual appliance for the following:

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

STEP 7 | Complete configuring the Panorama virtual appliance for your deployment needs.

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

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

2. Begin at Step 6 to switch to Log Collector mode.

Enter the Public IP address of the Dedicated Log Collector when you add the Log Collector as a managed collector to the Panorama management server. You cannot specify the IP Address, Netmask, or Gateway.

- For Panorama in Panorama mode.

  1. Add a Virtual Disk to Panorama on an ESXi Server.

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

  1. Set up a Panorama Virtual Appliance in Management Only Mode.
  2. Configure a Managed Collector to add a Dedicated Log Collector to the Panorama virtual appliance.

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

  1. Increase the System Disk for Panorama on an ESXi Server

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

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

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

Install Panorama on vCloud Air

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

STEP 1 | Download the Panorama 9.1 base image Open Virtual Appliance (OVA) file.

1. Go to the Palo Alto Networks software downloads site. (If you can't log in, go to the Palo Alto Networks Customer Support web site for assistance.)
2. In the Download column in the Panorama Base Images section, download the Panorama 8.1 release OVA file (`Panorama-ESX-9.1.0.ova`).

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

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

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

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

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

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

STEP 3 | Install Panorama.

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

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

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

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

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

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

Outbound traffic is allowed by default.

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

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

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

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